Hard/soft traffic isolation in transport networks

A flexible approach for network slicing using P-OTN

White paper

Transport network operators need to support new business models and increase network efficiency and scale so they can lower costs. These needs are driving them to look at new techniques to help future-proof their networks against the growing diversity of service requirements. Network slicing is expected to address these challenges by supporting the creation of multiple network instances using logical network resources. This approach will deliver the required performance for each application while optimizing network resources.

This paper clarifies the three categories of network slicing and discusses the hard and soft traffic isolation mechanisms used to provide physical/logical separation of network instances. It identifies drawbacks with the new Metro Transport Network (MTN) being defined in ITU-T to support hard isolation, then describes how the Nokia 1830 Photonic Service Switch (PSS) can provide flexible hard or soft isolation or a combination of both for traffic and synchronization distribution.

Contents

Introduction

The next-generation transport network will support the aggregation of multiple service types that have different capacity, latency, reliability and synchronization requirements. It will also support new shared-resource business models with a diverse set of applications and use cases. To meet all these demanding requirements simultaneously, service providers need flexible network designs that can optimize performance for each case. In other words, they need network slicing.

Network slicing creates multiple logical networks on top of a common shared physical infrastructure to serve the specific requirements of each service. For example, it could enable one transport network to serve up an ultra-low-latency slice for a 5G mobile network operator (MNO), a high-bandwidth slice for residential broadband services, and a premium performance wholesale slice for multi-tenant enterprise customers. This helps reduce costs by improving network utilization while providing the flexibility and scale needed to address growing demand.

The challenge for transport network operators is to ensure that the traffic on one slice does not interfere with the performance of another slice. This means a traffic isolation mechanism is required. These mechanisms are generally categorized as being either hard or soft. Hard isolation provides the highest level of separation, but it does so by utilizing dedicated bandwidth. Soft isolation allows for statistical multiplexing and a range of bandwidth efficiencies, but with a reduction in traffic isolation guarantees. For multi-service transport networks, service providers need a flexible solution that can provide the appropriate level of isolation to satisfy the requirements of each traffic type.

Network slicing and hard/soft isolation

Categories of network slicing

Network slicing can be broadly categorized into three types: multi-service, multi-tenant and 5G. With multi-service slicing, a service provider aggregates multiple traffic types (e.g., residential broadband, 4G/5G, enterprise private lines) into a converged metro network using a mix of client interfaces (e.g., Ethernet, SDH/SONET, OTN, Fibre Channel). To facilitate performance guarantees, the service provider maps each traffic type to a corresponding transport slice based on its quality of service (QoS) requirements.

With multi-tenant slicing, a service provider's transport infrastructure organization creates a transport slice for each of its business units, for example, residential services, wireless services, and business services. A business unit's transport slice could then be further sliced and offered to different customers, such as a large bank or healthcare network. This is sometimes referred to as a wholesale/retail model.

The third type is 5G slicing, perhaps the most common context, where an MNO uses a horizontal network slice to support a specific traffic type, such as enhanced mobile broadband (eMBB), ultra-reliable low-latency communication (URLLC) or massive machine-type communication (mMTC). A 5G slice is composed of a radio access network (RAN) slice, at least one transport slice and a core slice, from the user equipment (UE) to the application server. There are many references on this topic [1].

In all cases, the network slice resources are dynamically provisioned to accommodate the requirements of the service mix. An orchestrator creates, changes and removes them as needed using Software-Defined Networking (SDN).

Standardization of network slicing

Almost every telecommunications standards development organization (SDO) and forum is involved with some aspect of network slicing. Perhaps the most general interpretation and broadest applications are under development by MEF and the ETSI Zero-touch network and Service Management (ZSM) group. In a 5G mobile network context, there are ongoing efforts by the NGMN, 3GPP, GSMA and O-RAN Alliance. Then there is the combination of fixed–mobile convergence activities by the BBF and ITU-T SG13. Network slicing as it applies to transport networks is being addressed by the IEEE, IETF, ITU-T SG13 and ITU-T SG15 (where a network slice is referred to as a virtual network to avoid confusion with 3GPP slice terminology). Network management approaches for transport slices are being developed in the TM Forum, ONF and open source communities such as ONAP.

Although it is not exhaustive, this list illustrates the broad interest in this topic across the industry. The focus of this paper is on network slicing as it applies to transport networks, so it uses the terminology developed in ITU-T SG15.

Methods for transport network slicing

Traffic isolation goes hand in hand with network slicing because a method is needed to separate the traffic in one slice from the traffic in another slice within the same transport network. ITU-T SG15 uses the term virtual network (VN) to describe a network slice. It defines two basic methods for constraining traffic to a VN:

- Hard isolation is a method for isolating traffic on an interface or link using a circuit-switched connection. The traffic load of one VN has no impact on the traffic in any other VN, including QoS effects. Hard isolation is implemented by providing independent, circuit-switched connections for the exclusive use of one VN. These connections can be provided by, for example, a dedicated wavelength or a dedicated TDM timeslot, such as an ODUk/ODUflex in OTN or an MTN section (FlexE) timeslot. With hard isolation, there is no sharing of resources.
- Soft isolation is a method for isolating traffic on an interface or link at packet level using any Layer 2 or Layer 3 virtual private network (VPN) technique. The traffic load of one VN may have an impact on the QoS provided to the traffic in other VNs. Soft isolation is implemented by statistically multiplexing the traffic from two or more VNs using a packet technology. Traffic engineering can constrain the QoS impact of traffic on other VNs.

Note that Multiprotocol Label Switching (MPLS) tunnels, Segment Routing MPLS tunnels or segments, and hierarchical QoS queues are packet technologies and provide soft isolation only.

In a converged, multi-service metro network, the challenge for the transport network operator is to apply the type of isolation necessary to meet the QoS requirements of its service mix and provide the tools to monitor and protect each service appropriately while optimizing bandwidth efficiency.

MTN hard isolation issues

Some transport network operators have attempted to use a single traffic isolation mechanism, with little success. For example, a large service provider in China initially used soft isolation on its MPLS aggregation network but could not arrive at a single QoS policy that worked across its mix of services. This led the provider to initiate the development of the Slicing Packet Network (SPN), which is based on a hard isolation mechanism, defined by the national China Communications Standards Association (CCSA). The ongoing ITU-T SG15 work on MTN corrects many of the technical deficiencies in the SPN hard isolation mechanism.

MTN is being specified as a new type of transport equipment based on FlexE switching. Besides providing hard isolation for packet streams, it is intended to provide low-latency forwarding. Whereas the OIF FlexE Implementation Agreement defines a point-to-point link technology, MTN extends it to a network technology by creating a 66B block-based path layer for client MAC frames. It adds OAM, which is a hybrid of message and TDM formats, and provides switching using FlexE calendar slot cross-connections and path protection. The first set of base MTN standards may be completed in ITU-T by 2022; standards-compliant products may not be available until 2023.

MTN drawbacks

MTN being based on FlexE, its only client is the Ethernet MAC layer. Support for non-Ethernet clients would require either circuit emulation over Ethernet or the definition of new mappings into the 64B/66B physical coding sublayer (PCS) block structure, which would duplicate existing OTN client mappings.

Consequently, MTN is limited to Ethernet frame-based services only. The user–network interface (UNI) ports terminate the physical layer (PHY), including PHY-based fault signaling, and then encode the Ethernet MAC frames into 66B blocks. There is no support for PHY-based, constant bit rate (CBR) Ethernet services such as the codeword-transparent mappings in OTN, which preserve client timing.

The MTN path layer forwards the 66B blocks through intermediate nodes without decoding to MAC frames, thereby providing a “cut-through” layer to eliminate store-and-forward delay and provide hard isolation. However, this approach is inconsistent with IEEE 802.3 and IEEE 802.1 standards. Further, with modern switching technology, the store-and-forward delay is already very small. When the 66B blocks traverse more than one link and thereby experience more degradation, the PHY performance required by the receiving MAC is no longer guaranteed.

1830 PSS hard/soft isolation using P-OTN

The 1830 PSS Packet-OTN (P-OTN) solution supports both soft and hard isolation using proven, standardized technology to provide network flexibility and efficiency. Figure 1 illustrates how the solution supports soft isolation by separating traffic into different Layer 2 Ethernet virtual connections (EVCs) using Carrier Ethernet bridging or MPLS Transport Profile (MPLS-TP), and hard isolation by separating traffic into different ODUflex channels with a granularity of 1.25 Gbit/s, or wavelengths. Customers can freely allocate the line bandwidth between soft and hard isolated traffic, thereby enabling efficient bandwidth utilization.

Figure 1. 1830 PSS hard/soft isolation

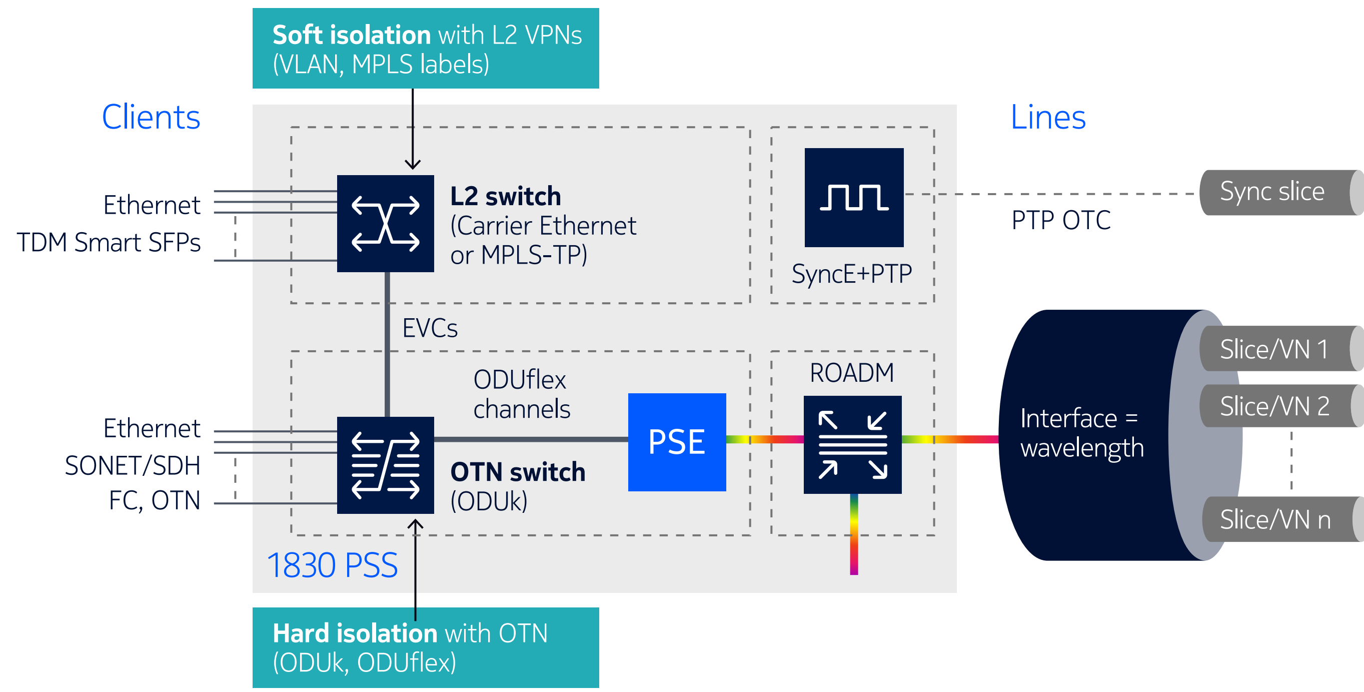

Both the Carrier Ethernet (connectionless) and MPLS-TP (connection-oriented) soft isolation approaches are augmented with transport-grade OAM and protection switching. For Carrier Ethernet, IEEE 802.1ag [2], ITU-T Y.1731 [3], MEF 30.1 [4]/35.1 [5] service OAM and ITU-T G.8032 [6] protection switching (for service assurance) and ITU-T Y.1564 [7] (for service activation) are supported. For MPLS-TP, the OAM functions defined in ITU-T G.8113.2 [8] and the associated suite of IETF RFCs are supported, as are ITU-T G.8131 [9] label-switched path (LSP) protection and IETF RFC 6718 [10] pseudowire (PW) redundancy.

The 1830 PSS supports standardized OTN mappings for Ethernet, SONET/SDH, Fibre Channel and OTN clients into ODUk/ODUflex, including the CBR mappings that preserve client timing, a necessary feature in 4G/5G networks. This approach to hard isolation also ensures the lowest latency. Another fundamental benefit of hard isolation is the inherent high level of security from dedicated connections, a traditional feature of transport networks (e.g., SONET/SDH). Note that while the 1830 PSS supports both hard and soft isolation, routers only support soft isolation.

In multi-service networks, it is not possible to anticipate all service types on Day One, and the service mix and traffic patterns often change over time. The 1830 PSS supports multi-tier, elastic bandwidth adjustment for both hard and soft traffic isolation of MEF 63/64 Layer 1 services [11, 12] and MEF 6.3/51.1 Layer 2 services [13, 14], respectively, as illustrated in Figure 2, allowing customers to adjust the limits as their network evolves. The ability to statistically multiplex Layer 2 services into a soft isolated EVC prior to combining in a hard isolated ODUflex allows for statistical multiplexing gain and ensures that bandwidth is not stranded. Elastic bandwidth adjustment for hard traffic isolation is implemented using ITU-T G.709 [15], while for soft traffic isolation it is implemented using MEF 47.1 [16]. Automated elasticity is provided by SDN, for both hard and soft isolation. This ensures that the network is self-adjustable to traffic changes, which leads to lower OPEX.

Figure 2. 1830 PSS multi-service, multi-tier elastic bandwidth

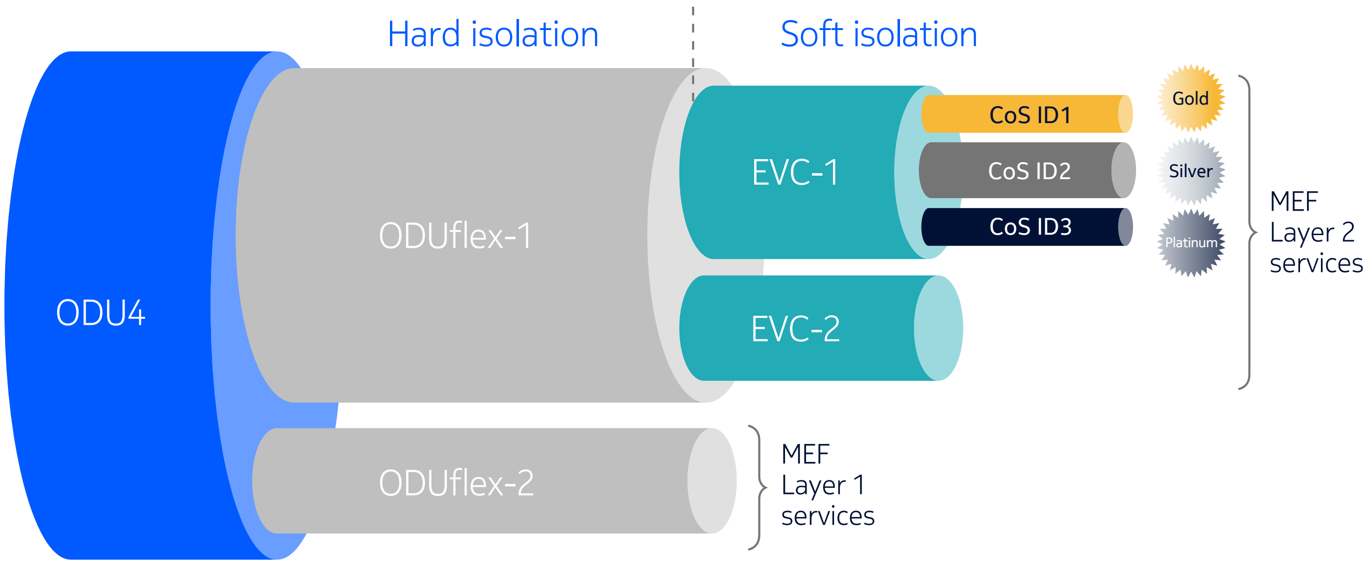

The same hard and soft isolation mechanisms used to ensure flexibility and bandwidth efficiency in multi-service networks can be used for recursive network slicing in multi-tenant, wholesale/retail business models (provided by Nokia WaveSuite Service Enablement) and to support 5G slices.

Hard isolation for synchronization

The concept of hard isolation can be extended to the synchronization plane. Time/phase synchronization has stringent performance requirements, such as uniform, bidirectional latency symmetry, and no interference from other services. Using hard isolation with a dedicated wavelength, or optical timing channel (OTC), to transport synchronization over a WDM network provides the determinism and accuracy the time recovery algorithms require to accurately recover time and phase. Figure 3 illustrates the use of a dedicated transport slice for synchronization.

Figure 3. 1830 PSS dedicated transport slice for synchronization

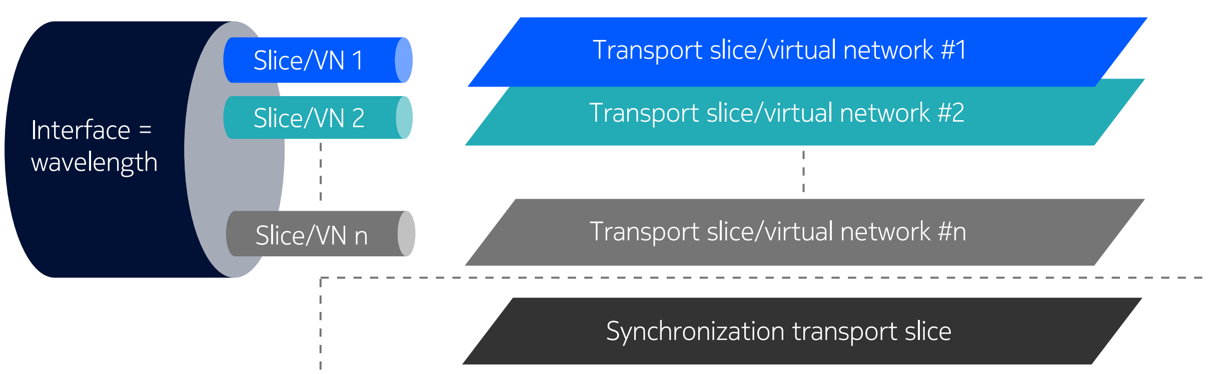

Time/phase synchronization is distributed through a network using the IEEE 1588v2-based Precision Time Protocol (PTP). PTP operates by exchanging time-stamped frames through a chain of PTP clocks that assume a symmetrical delay over a syntonized physical layer (using Synchronous Ethernet for frequency distribution). Fixed asymmetry, such as differences in the lengths of receive and transmit fibers, can generally be compensated by calculating the resulting time error and applying it as an offset. Variable asymmetry caused by the variability in buffering and queuing, OTN mapping/demapping and reconfigurable optical add/drop multiplexer (ROADM) switching from node-to-node and direction-to-direction creates fluctuations in the delay that a PTP frame experiences in the network and is more difficult to correct. The use of a bidirectional OTC eliminates fixed and variable link asymmetry, which is a huge contributor to time inaccuracy in WDM networks.

Delay and offset are important parameters in synchronization and are continuously monitored by OAM tools as described in ITU-T G Suppl. 68 [17]. With a dedicated wavelength to transport time synchronization, the values for these parameters are deterministic and highly representative of actual network conditions. This makes it easier to monitor and resolve any issues with service level specification (SLS) compliance.

Conclusion

Transport network slicing is used in multi-service networks, multi-tenant wholesale/retail business models and 5G slices. Hard and soft traffic isolation are mechanisms used to provide network slicing of resources. The challenge for the transport network operator is to flexibly apply the type of isolation necessary to meet the QoS requirements of its service mix while optimizing bandwidth efficiency. Soft isolation and hard isolation are insufficient on their own.

The 1830 PSS P-OTN solution supports network slices for multi-service, multi-tenant and 5G using a flexible solution of packet for soft isolation and OTN for hard isolation, with self-adjusting elastic bandwidth optimization under SDN control. Furthermore, to transport time and phase synchronization, the 1830 PSS extends hard isolation by using a dedicated wavelength to provide the determinism and accuracy required by the time/phase distribution protocols.

Abbreviations

CBR	constant bit rate
CoS ID	class of service identifier
eMBB	enhanced mobile broadband
EVC	Ethernet virtual connection
ITU-T	International Telecommunication Union-Telecommunication Standardization Sector
LSP	label-switched path
MAC	Medium Access Control
mMTC	massive machine-type communications
MNO	mobile network operator
MPLS	Multiprotocol Label Switching

MPLS-TP	Multiprotocol Label Switching – Transport Profile
MTN	Metro Transport Network
OAM	operations, administration and maintenance
ODUk	Optical Data Unit-k
ODUflex	Optical Data Unit-flexible
OPEX	operational expenditure
OTC	optical timing channel
OTN	Optical Transport Network
PCS	physical coding sublayer
PHY	physical layer entity
P-OTN	packet-optical transport network
PSE	Photonic Service Engine
PSS	Photonic Service Switch
PTP	Precision Time Protocol
PW	pseudowire
QoS	quality of service
RAN	radio access network
ROADM	reconfigurable optical add/drop multiplexer
SDH	Synchronous Digital Hierarchy
SDN	Software-Defined Networking
SDO	standards development organization
SG	Study Group (in ITU-T)
SLS	service level specification
SONET	Synchronous Optical Network
SPN	Slicing Packet Network
UE	user equipment
UNI	user-network interface
URLLC	ultra-reliable low-latency communication
VN	virtual network
VPN	virtual private network
WDM	Wavelength-Division Multiplexing

References

1. 5G network slicing white paper, Nokia
2. IEEE Std 802.1ag, IEEE Standard for Local and Metropolitan Area Networks - Virtual Bridged Local Area Networks Amendment 5: Connectivity Fault Management
3. Recommendation ITU-T G.8013/Y.1731 Operation, administration, and maintenance (OAM) functions and mechanisms for Ethernet-based networks
4. MEF 30.1 Service OAM Fault Management Implementation Agreement: Phase 2
5. MEF 35.1 Service OAM Performance Monitoring Implementation Agreement
6. Recommendation ITU-T G.8032/Y.1344 Ethernet ring protection switching
7. Recommendation ITU-T Y.1564 Ethernet service activation test methodology
8. Recommendation ITU-T G.8113.2/Y.1372.2 Operations, administration and maintenance mechanisms for MPLS-TP networks using the tools defined for MPLS
9. Recommendation ITU-T G.8131/Y.1382 Linear protection switching for MPLS transport profile
10. IETF RFC 6718 Pseudowire Redundancy
11. MEF 63 Subscriber Layer 1 Service Attributes
12. MEF 64 Operator Layer 1 Service Attributes and Services
13. MEF 6.3 Subscriber Ethernet Service Definitions
14. MEF 51.1 Operator Ethernet Service Definitions
15. Recommendation ITU-T G.709/Y.1331 Interfaces for the optical transport network
16. MEF 47.1 Elastic Ethernet Services & Cloud Connectivity
17. ITU-T Series G Supplement 68, Synchronization OAM requirements

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

Document code: (April) CID210871