Beyond 10G PON Technologies and Network Slicing

50G PON Capacity + PON Slicing are Game Changers

A technical paper prepared for presentation at SCTE TechExpo24

Michael Emmendorfer
Vice President of Technology
Calix
mike.emmendorfer@calix.com

Table of Contents

List of Figures

List of Tables

Title **Page Number**

1. Executive Summary

Service providers worldwide have adopted a fiber first strategy for all new build areas and some cable operators, and most telcos have transitioned their legacy cable or copper networks to fiber to the home (FTTH). The use of FTTH is growing tremendously worldwide and in the United States our internal research estimates that fiber households passed may reach 73% by 2026, and over time many of these homes will be connected with fiber using PON technology. Service providers have used fiber for network transport services to businesses and for mobile backhaul for over two (2) decades, with point-to-point active Ethernet (AE) as the primary choice, however this may change with 50G PON and PON slicing.

Next generation PON technologies like ITU-T 50G PON as well as ITU-T PON slicing will be game changers. The use of 50G PON and PON slicing will be part of an end-to-end (E2E) network slicing architecture enabled across multiple network technologies and segments from the end user through the network core. The E2E network slicing architecture will include cross-domain slicing service and network orchestration, SDN domain controllers, network management and analytics platforms, and network elements including mobile, Wi-Fi, PON, and routing platforms that all enable network slicing. These technologies will enable service providers with a highly flexible network to automate E2E network slices of capacity and quality of service (QoS), statically or dynamically, across network technology types, domain controllers, and vendors. The emergence of 50G PON slicing technology will have the capacity to partition bandwidth and QoS for groups of one or more flows associated with one or more ONUs, called a PON slice. The capacity of 50G PON enables multiple PON slices per OLT port and each slice is managed by a dynamic bandwidth assignment (DBA) and a hierarchical scheduler that manages across all PON slices. PON slicing can define guaranteed parameters per-slice and/or per-flow, and surplus bandwidth may be shared within the slice level and even between all slices, thus not wasting non-guaranteed partitioned bandwidth capacity.

Next-gen 50G PON systems should be part of an end-to-end network slicing architecture enabling ITU-T PON slicing on subscriber facing optical line termination (OLT) and optical network unit (ONU) interfaces. On the OLT system WAN-facing interfaces network slicing will be enabled using IETF segment routing (SR). The capacity of 50G PON and the versatility of PON slicing will enable new PON use cases to include business services, 5G mobile backhaul, mid-haul, and fronthaul, mid-band and mmWave fixed wireless access (FWA) backhaul, and Wi-Fi 7 access point backhaul for smart town. In open access models, wholesale service providers could use slicing per retail ISP, grouping ONUs or grouping flows with similar bandwidth and QoS parameters into PON slices. A mobile Wi-Fi offload offering to several mobile operators using community or residential Wi-Fi and PON slicing to offload mobile traffic from the 5G macro radio access network (RAN) is yet another use case. Today, service providers operate parallel networks using PON for residential and some business customers and optical Ethernet for high end businesses and aggregation layer transport. In the future the role of 50G PON and PON slicing will serve as both access layer, connecting end customers, and aggregation layer transport.

This paper examines next generation PON technologies beyond 10 Gbps. It explains the optional support of a 50G PON OLT with a dual line-rate upstream receiver that enables service providers with economic flexibility to support both 50G x 50G ONUs and 50G x 25G upstream ONUs using the same wavelength and the same OLT port. Additionally, the ITU created a separate and optional specification called ITU-T Supplement 74 – PON Slicing. Though optional, OLT and ONU systems that can enable PON slicing functions into 50G PON technology will enable differentiated services and capabilities for the operator. The paper also examines network slicing defined across the telecom industry and in conjunction with PON slicing this could enable an end-to-end network slicing architecture. The paper examines current and future PON technologies to support current and future business drivers and use cases. Residential service tier and traffic growth rates are forecasted to predict the useful life of current and future PON technologies including GPON, XGS-PON, 25GS-PON, and 50G PON.

2. Beyond XGS-PON Technologies

Around the year 2007 the first GPON deployments took place, and approximately nine years later, in 2016, the first XGS-PON deployments took place. This year, 2024, the first commercial launch of PON services using 25GS-PON took place. It is estimated that the first commercial launch of services using 50G PON may take place in 2025 or 2026. These data points mean that the PON industry introduces a next generation PON technology approximately every 8 to 10 years. The market adoption in terms of OLT port and ONU sales exceeding the previous PON generation technology takes much longer. In the year 2021, XGS-PON OLT port sales in North America surpassed GPON OLT port sales and it is predicted by Omdia that in 2027 XGS-PON ONT sales will exceed GPON ONT sales worldwide. As described in the traffic engineering and capacity planning section of this paper, we are projecting that a GPON OLT port can support the peak period traffic of 32 subscribers and a 1 Gbps service tier through the years 2035 or 2036. These data points illustrate that PON technologies are intended to last a long time. It took 14 years for XGS-PON OLT sales to surpass GPON and it took 20 years for XGS-PON ONT sales to surpass GPON. The useful life of GPON may surpass 30 years, of course due to the coexistence with higher capacity PON technology that will enable higher service tiers or speed tiers. The next generation PON technology can extend the useful life of the legacy PON technology by co-existing on the same fiber to enable higher capacity and services not possible on the legacy PON. Considering the last 25 years of PON technology evolution from APON, BPON, GPON, XGS-PON, and 50G PON these had roughly four (4) times increase in capacity. As the industry looks beyond XGS-PON and plans the next generation of PON technology deployments, these historical factors should be considered, such as the capacity increase to enable a long useful life.

2.1. 25GS-PON Multi-Source Agreement (MSA) Group Overview

The 25GS-PON Multi-Source Agreement (MSA) Group is not a standards organization. The MSA group assembled and modified materials from several standard organizations including the IEEE, ITU-T, and the BBF to create a document. The MSA used the IEEE standard 802.3ca™-2020 for PMD layer and FEC and ITU-T G.9807.1 XGS-PON for transmission convergence (TC) layer. The 25G downstream uses the IEEE 802.3ca forward error correction (FEC), however, if the upstream is 10G this uses the FEC from the ITU-T and 25G upstream uses the FEC from the IEEE 802.3ca. The nominal line rates supported in the 25GS-PON MSA include a downstream line rate at 24.8832 Gbps and upstream line rates of 24.8832 and 9.95328 Gbps. The useable data rate after FEC is approximately 21 Gbps.

2.2. ITU-T 50G PON Overview

The ITU-T 50G PON standard is called ITU-T G.9804 HSP G.hsp - Higher Speed PON. The nominal line rates supported include a downstream line rate at 49.7664 Gbps and upstream line rate 49.7664 Gbps, 24.8832 Gbps, and 12.4416 Gbps. The forward error correction (FEC) technology used is Low-density Parity Check (LDPC). The downstream uses a FEC notation of LDPC (17280, 14592) and the upstream LDPC uses (17280, 14592) or (15872, 14592). The useable downstream data rate after FEC overhead is approximately 42 Gbps.

2.3. ITU-T 50G OLT Dual Line-rate Upstream Receiver Overview

The 50G OLT system uses a single wavelength 50G downstream channel and may support the optional dual line-rate upstream receiver. The ITU-T 50G standard allows for an OLT to support dual line-rates with a nominal 25 Gbps or 12.5 Gbps upstream, in addition to the 50 Gbps upstream line rate [G.9804.1]. The 50G OLT dual rate receiver enables service providers with economic flexibility to select ONUs with 50G downstream and 25G upstream as well as 50G downstream and 50G upstream based on cost and customer types. It is likely that 50G ONUs will have an application-specific integrated circuit (ASIC) capable of 50G symmetrical, however the 25G upstream optical technology could be first to market followed by 50G upstream optics. There may also be cost deltas between 25G and 50G upstream initially. Figure 1 is an illustration of the 50G OLT supporting dual line-rate upstream for symmetric and asymmetric ONUs. These ONUs will use the same wavelength and OLT port with the OLT DBA scheduler assigning timeslots per ONU at upstream line rates of 25G or 50G in this example. Service providers with concerns of cost points and availability 50G x 50G ONU optics may start with 50G x 25G and strategically purchase 50x50G ONUs where and when needed. The advantage to the service provider is economic flexibility, that enables the purchase of a 50G OLT and options of asymmetric and symmetric ONUs that terminate on the same OLT port while also using the same wavelength in the ODN.

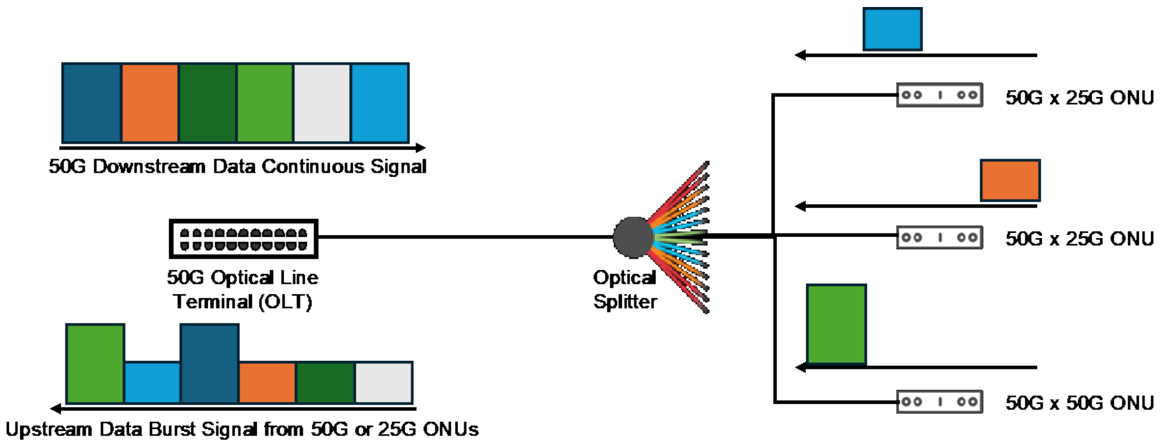

Figure 1: 50G OLT Dual Line-rate Upstream Receiver for 50x25G ONUs and 50x50G ONUs

2.4. PON Wavelength Considerations

The ITU-T and IEEE defined PON wavelengths that overlapped with each other in many cases as shown in Figure 2. The ITU-T 50G PON and 25GS-PON MSA defined the same three upstream wavelengths and two of the wavelengths overlap with GPON and XGS-PON as shown in both Figure 2 and Figure 3. The ITU-T defines the upstream wavelengths as Option 1 with a value of 1260 to 1280 nanometer (nm), or 1270 +/- 10 nm, which overlaps with XGS, Option 2 with a value of 1290 to 1310 nm, or 1300 +/- 10 nm, which partially overlaps GPON, and Option 3 with a value of 1284 to 1288 nm, or 1286 +/- 2 nm, as shown in Figure 3. If a service provider has deployed GPON and XGS-PON on the same fiber, then the operator has one wavelength available, Option 3 1284 to 1288nm. Our recommendation is to use that last wavelength for 50G PON that will meet the capacity and service needs over the expected life of a PON technology.

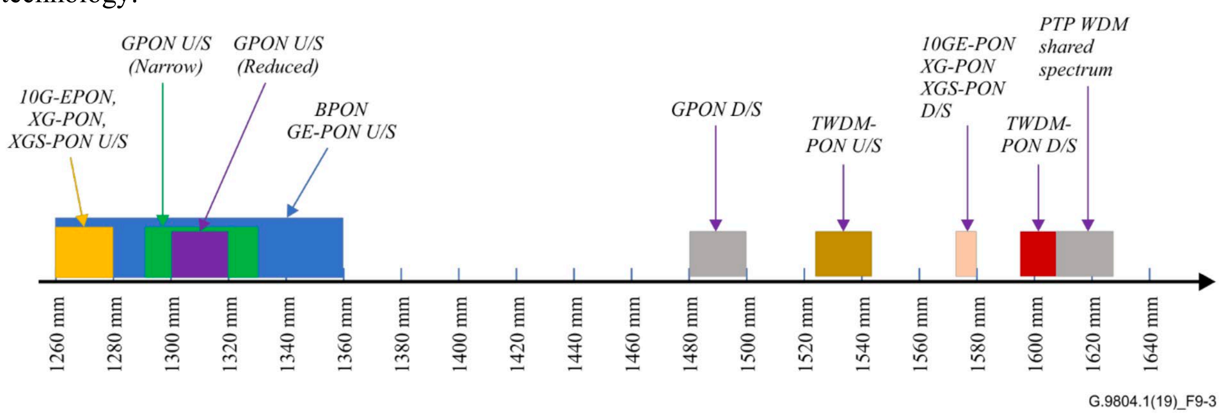

Figure 2: PON Wavelengths Prior to ITU-T 50G [G.9804.1]

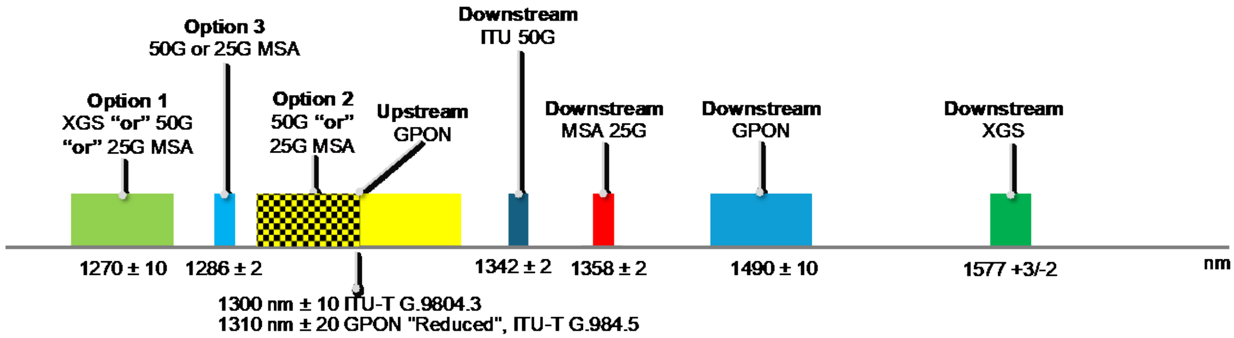

Figure 3: PON Wavelengths for GPON XGS 25GS-PON and 50G PON

3. Network Slicing Overview

The term network slicing is now used across the telecommunications industry from Mobile 3GPP 5G, IEEE Wi-Fi, ITU-T PON, Metro Ethernet Forum (MEF), and the Internet Engineering Task Force (IETF). Each of these organizations have defined use cases and standards to enable network slicing and when combined, this creates an end-to-end network slicing framework. The E2E network slicing framework will be from the user equipment (UE), customer premises equipment (CPE), access layer, provider edge, and through the service provider's distribution and core network.

What is network slicing? Network slicing allows a physical network resource to be logically partitioned into multiple logical networks, called a network slice, with each slice having bandwidth, QoS and latency parameter settings. The network slices may be statically reserved or dynamically allocated with minimum guarantees as well as maximum capacity thresholds. Network slicing can be very flexible, allowing for on-demand slicing whereby the slice could be dynamically created and then removed. Since network slicing is highly flexible, the assignments could even be time based to more efficiently and effectively use shared network resources, where traffic patterns may differ between customer types and during times of the day. An example of traffic patterns being different are business users and residential users that each have different peak periods of traffic utilization. Dynamic network slicing could allow service providers to incentivize business customers to purchase a minimum bandwidth guarantee, while offering a significantly higher non-guaranteed bandwidth class, and if non-guaranteed capacity is unused, it could be shared across other network slices. Dedicated networks, like active Ethernet, allocate a port at the service provider's facility and a wavelength for each customer, which is fixed regardless of the network utilization. The flexibility of PON slicing could enable services on par with dedicated active Ethernet networks, however unlike AE an OLT PON port terminates many customer connections and uses a single wavelength for all customers. The use of PON and dynamic PON slicing for residential and commercial services could take advantage of different traffic utilization periods to more cost effectively and efficiently use network, fiber, space, and power resources. The capacity of 50G PON and the use of PON slicing will enable service providers to use PON for more use cases, thereby reducing, but not eliminating, the use of active Ethernet.

Key network slicing concepts regardless of technology include [SANOG36]:

- Multiple virtualized and independent logical networks on the same shared physical infrastructure with each slice tailored to fulfil diverse requirements
- Partitioning of network resources
- Service guarantee for throughput, latency and jitter without impacting other logical networks
- Slice isolation - performance, traffic separation, security, privacy, and management
- Orchestration and control – end-to-end and multi-domain

Figure 4 illustrates an end-to-end (E2E) network slicing architecture and the industry scope for each network slicing segment to include IEEE Wi-Fi, 3GPP, ITU-T PON Supplement 74, and IETF.

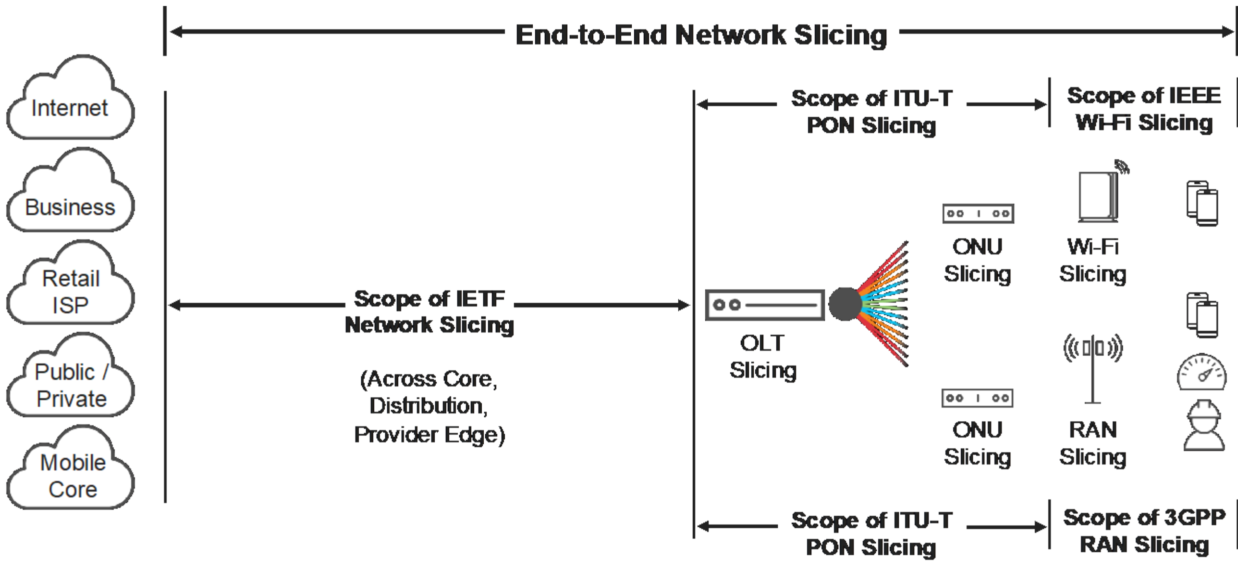

Figure 4: Industry Scope for End-to-End Network Slicing

3.1. ITU-T PON Slicing Overview

What is a PON Slicing? Allocating a portion of PON capacity to a group of users with each group having its own DBA. All DBAs are managed by a hierarchical traffic scheduler. The benefits of PON slicing are that each slice and members in a slice can have configurable bandwidth and latency properties. Any bandwidth unused above guaranteed or committed information rate may be shared with users within each slice and even among all slices, to not waste unused capacity.

The Broadband Forum (BBF) is examining a standards approach to manage end-to-end (E2E) slicing, from cloud orchestration and software domain controllers, network access layer, and the customer premises devices including the optical network termination (ONT) and residential gateway (RG). Orchestration and domain controllers will handle end-to-end service orchestration and flexible programmability of end-to-end network slicing, networking automation, policy enforcement, usage-based billing, and proactive network monitoring.

In Figure 5, sourced from the ITU-T Supplement 74 (12/2021), this illustrates network slicing use cases and network slicing types [ITU-Sup74]. There are four slice types defined in PON slicing: 1) enhanced fixed broadband (eFBB), 2) guaranteed reliable experience (GRE), 3) Internet of things (IoT), and 4) High Bandwidth Low Latency (HBLL). In the 3GPP 5G slicing section below there will be similarities to the slice types defined in ITU-T Sup74. This illustration has many use cases for PON slicing across many market segments like residential, wholesale, smart city, industrial, hospital or any other enterprise-oriented segments [ITU-Sup74]. Network slice instances and slice types will meet the corresponding QoS requirements. There are many other groupings of service or customer types that could be placed into PON slices not shown in this figure.

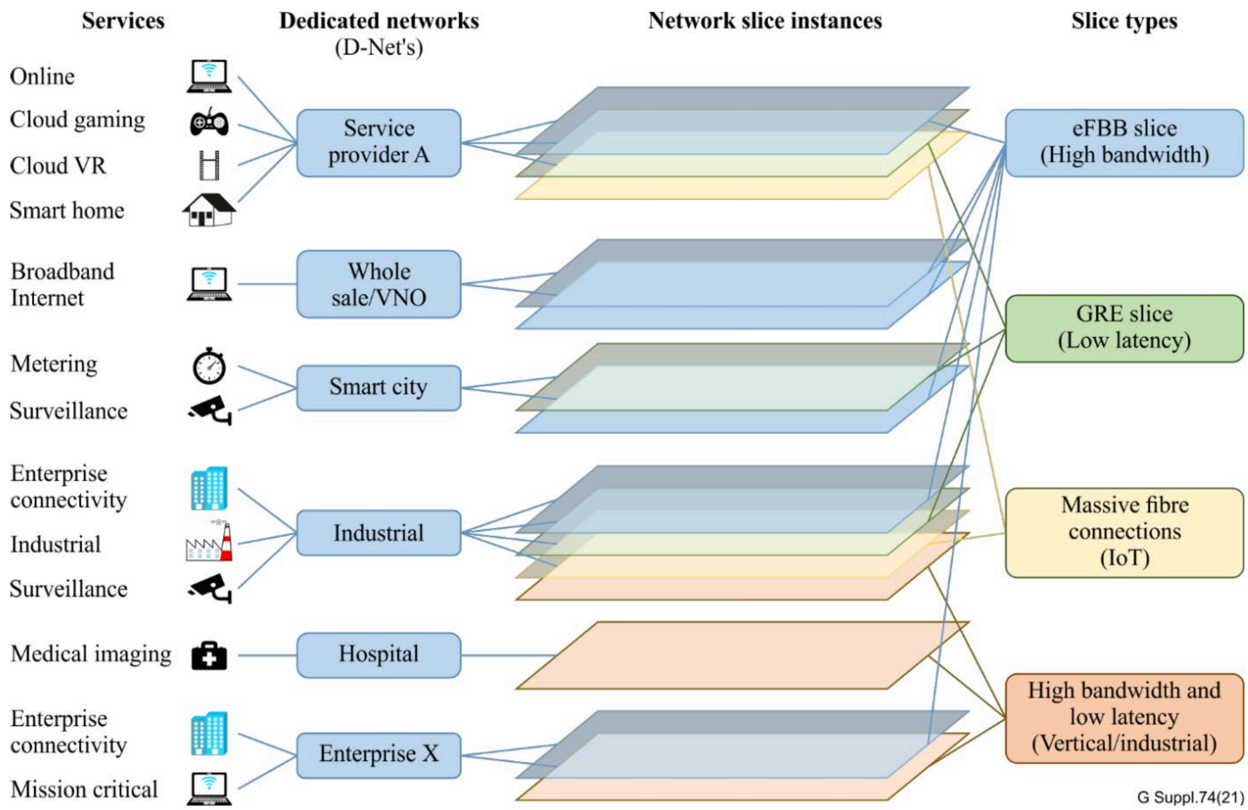

Figure 5: ITU-T G. Supp 74 PON Slicing for Services Requiring Different Slice Types

Figure 6 illustrates a PON system and a high-level end-to-end network slice through several network segments. Slices are a virtual network and Figure 6 illustrates an ITU-T end-to-end network slicing example [ITU-Sup74]. The PON system defines a system network interfaces (SNI) to interconnect with the IETF network slicing network. The user network interface (UNI) connects to the end user customer network. The PON-based access network includes PON OLT equipment, which may support multiple OLT Channel Terminations (CT), and each OLT CT supporting an ODN with multiple subtending ONUs. In Figure 6 the segments are identified as:

- (2A) represents the OLT slicing;
- (2B) represents PON slicing, and each OLT CT can carry multiple PON slices to/from the PMD+TC (Physical Medium Dependent and Transmission Convergence functions of the associated ONUs;
- (2C) ONU slicing represents slicing in the part of the ONU behind its PMD+TC function [ITU-Sup74].

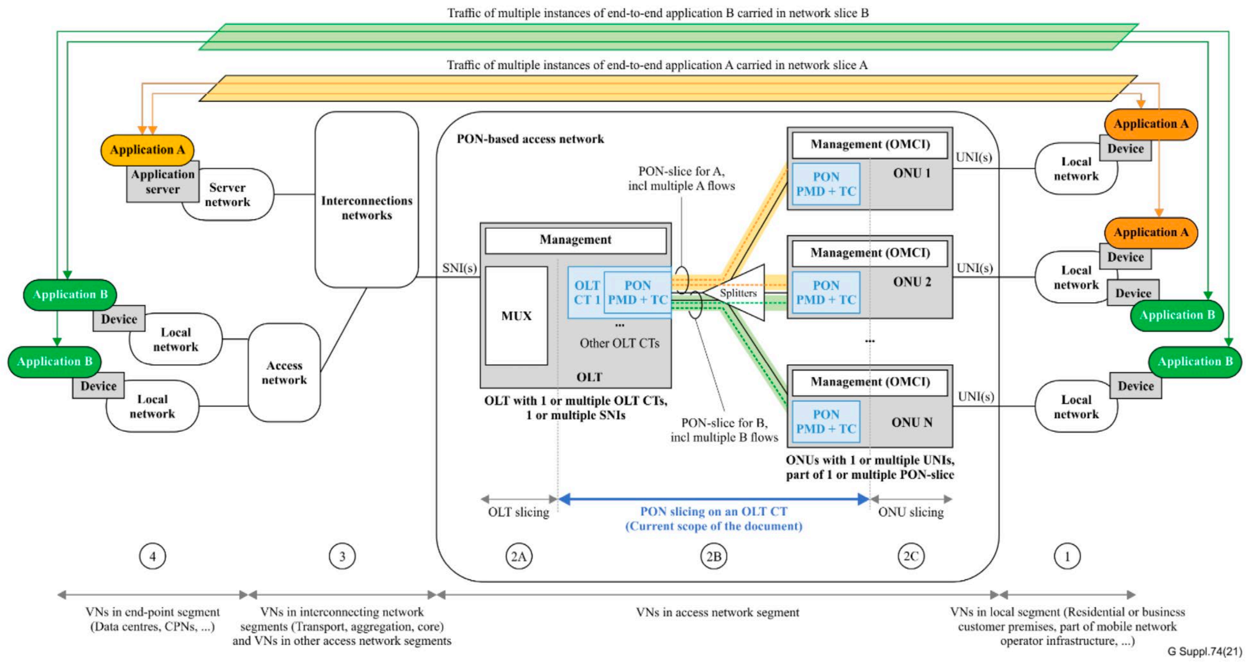

Figure 6: ITU-T G. Suppl.74 PON System and PON Slicing

3.1.1. PON Slicing Functional Architecture

A PON system is composed of an OLT and ONUs used to transport network slices. The OLT has SNIs to interconnect with devices that may enable IETF network slicing. The OLT CT has the PON PMD+TC (Physical Medium Dependent and Transmission Convergence) functions that connect to multiple ONUs. The OLT performs dynamic bandwidth assignments (DBA) in the upstream and hierarchical scheduling downstream. Prior to supporting PON slicing concepts, an OLT had a single DBA that allocated bandwidth using a fairness algorithm at a flow level. In a PON slicing architecture the OLT has an upstream DBA function at the slice level and a hierarchical scheduler managing multiple PON slices each with a DBA function. In the downstream direction, the OLT scheduling hierarchy is extended a level for PON slicing. The ONU has PON PMD+TC connecting with the OLT and UNIs connecting to end user customer devices. ITU-T PON slicing definition are listed below sourced from the ITU-Sup74:

OLT slice: An optical line termination (OLT) slice is one partition of the traffic management functions of the OLT intended to facilitate network slicing over a PON-based access network. An OLT slice might have its own management of traffic and other parameters and appear as an independent logical OLT. A sliced OLT would appear to north-bound network management systems as several logically independent OLTs.

ONU slice: An optical networking unit (ONU) slice is one partition of the traffic management and buffering functions of the ONU intended to facilitate network slicing over PON-based access network. An ONU slice might have its own management of traffic and other parameters. ONU slicing involves the partition of the traffic management and buffering functions of the ONU. A conventional ONU has corresponding capabilities (e.g., traffic management) that are controlled via OMCI. A sliced ONU could have multiple similar functional instances, each controlled independently. Clearly, an indeterminate coordination or unification of all these different functional instances is assumed. Clause 6.5 describes the two possible scenarios from an ONU perspective, namely slice-aware ONU, and slice-unaware ONU. Further considerations for ONU slicing are for future study.

PON slice: A PON slice is a group of one or more flows associated with one or more ONUs that are treated as a single entity by a hierarchical traffic scheduler.

Slice-aware and Slice-unaware: In a slice-aware optical access network segment, the optical line termination (OLT) is slice-aware, while the optical network unit (ONU) can be either slice-unaware or slice-aware:

- In the case that all services associated with the ONU belong to one slice, the ONU can be slice-unaware. In the case where the ONU is slice-unaware, the mapping/de-mapping between PON slices and the service content of network slices happens at the OLT.
- In the case that services associated with the ONU belong to different slices, depending on QoS requirements per slice, the ONU may or may not need slice-awareness in order to support the QoS requirement for the user data associated with each slice. In the case where the ONU is slice-aware, the mapping/de-mapping between PON slices and service content of network slices happens at both the OLT and at the ONU.

3.1.2. PON Slicing Use Cases

Slicing scenarios include an ONU participating in a single slice as shown in Figure 7. An ONU participating in multiple slices is shown in Figure 8. A PON slicing architecture can support a mix of ONUs each with different slicing configurations. A PON port can enable multiple PON slices and support ONUs that may connect to a single slice while other ONUs may have flows that connect to multiple slices on the PON as shown in Figure 9. These figures are sourced from the ITU-T G. Suppl.74.

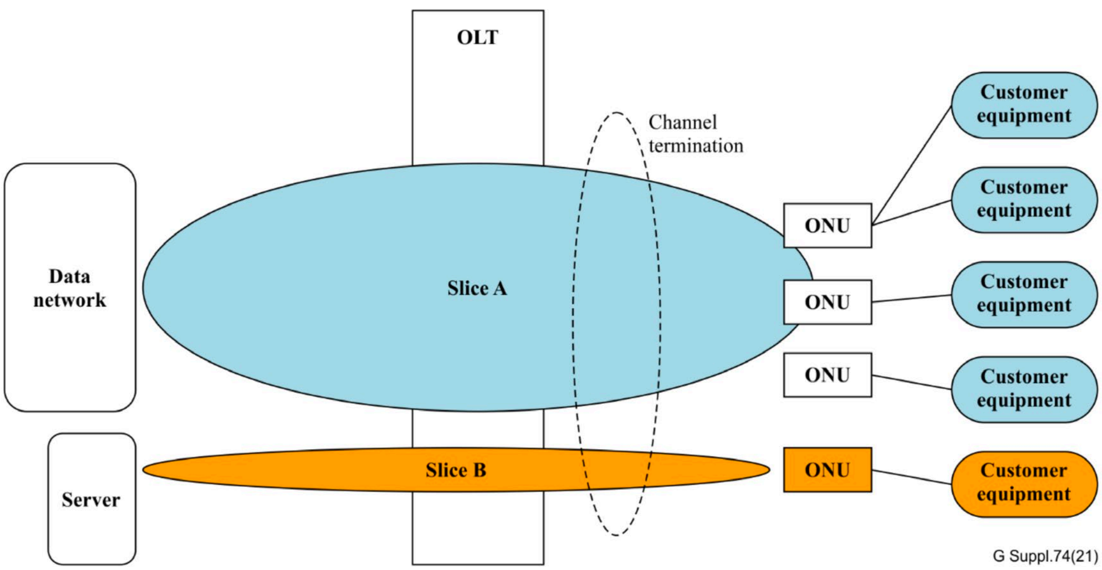

Figure 7: ONU Dedicated to a Single Slice

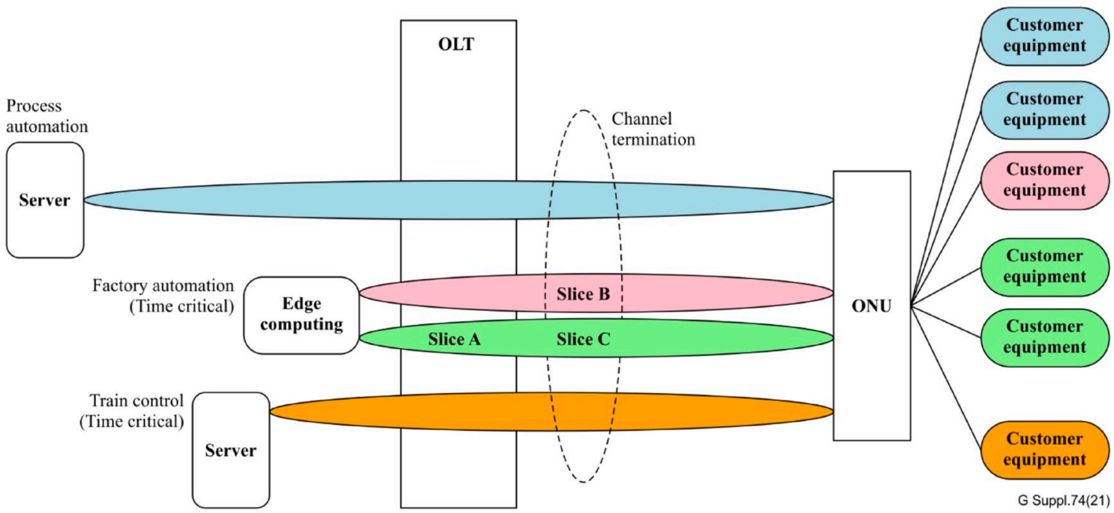

Figure 8: One ONU Carrying Multiple Slices Each Slice with Different Requirements for Rate, Delay, Isolation, and Routing

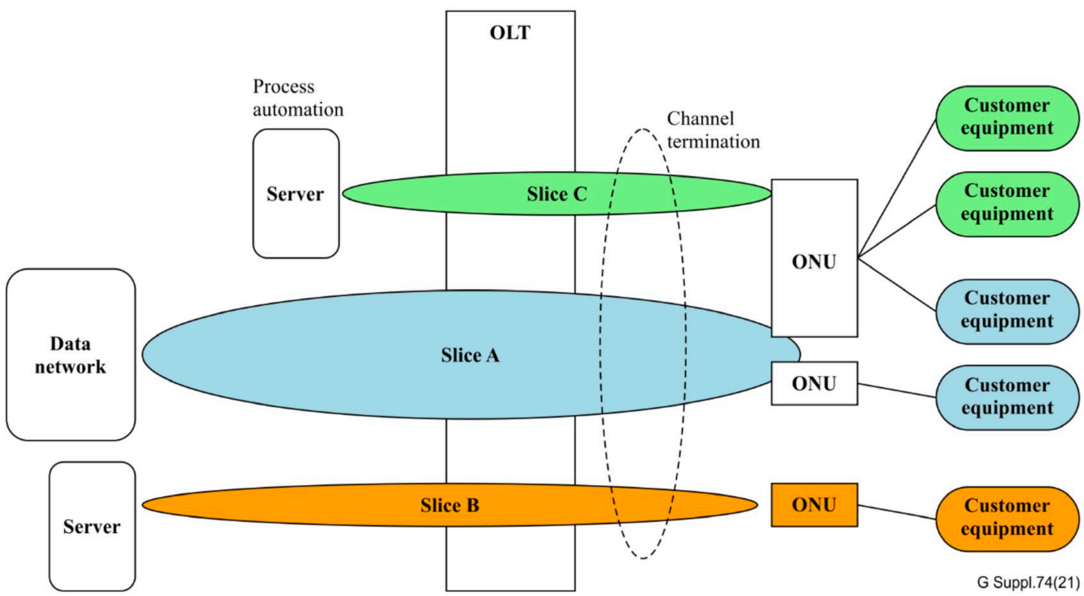

Figure 9: Mix of ONUs carrying single or multiple slices

The use of PON slicing could be part of the E2E network slicing architecture that a service provider enables to support various uses cases. There are other standards organizations developing network slicing technology and standards, for example the 3GPP for 5G mobile slicing, IEEE for Wi-Fi slicing, and at the network core, distribution layers, and provider edge network segments would use network slicing defined by the IETF. The following sections provide an overview of these network slicing segments.

3.2. 3GPP 5G Network Slicing Overview

The 3GPP defined 5G network slicing to enable mobile operators to partition their networks for specific customer use cases that provide different amounts of network resources for different types of traffic. A 5G network slice allocates resources to support service level agreements (SLA) connectivity bandwidth (speed) and latency parameters. "Network slicing allows multiple logical networks to be created on top of a common shared physical network", according to Verizon [VZ]. Verizon also identified use cases for 5G network slicing to include Internet of Things (IoT) in a manufacturing environment, operating autonomous vehicles, separating customer traffic into different slices like AI-driven video analytics and point-of-sale information, and autonomous forklifts in a factory [VZ]. The 5G use cases enhanced mobile broadband (eMBB) targeted at high data rates across wide coverage areas [SANOG36]. The ultra reliable low latency communication (URLLC) targets 1 millisecond of latency, security, and reliability of 99.999% targeted at autonomous driving and mission critical applications [SANOG36]. Massive Machine Type Communication (mMTC) serves large number of devices that transmit small amounts of data targeted at low-cost endpoints [SANOG36].

3.3. IEEE Wi-Fi Network Slicing Overview

Many of the concepts of network slicing have been foundational in IEEE Wi-Fi for many years. For example, consumers and enterprises use Wi-Fi to support private and guest networks on a Wi-Fi access point, creating a virtual network. Service providers have enabled carrier Wi-Fi on public Wi-Fi access points (APs) and residential Wi-Fi APs. IEEE Wi-Fi can logically separate Wi-Fi networks with different policies and security on the same Wi-Fi physical infrastructure. Wi-Fi network slicing architecture can be implemented via service set identifier (SSID). According to a Wireless Broadband Alliance (WBA) paper published in 2018, Wi-Fi network slicing can be implemented using several techniques. In the WBA paper, a controller-based architecture can dynamically allocate VLANs with different groups of users as illustrated in Figure 10. The use of dynamic VLAN assignment enables the slice selection to be based on network policy, rather than handset configuration. The WBA stated that the Wi-Fi industry is widely using concepts of network slicing within enterprise deployments to isolate corporate traffic and users from guest users via VLAN [WBA2018]. Service providers are offering carrier Wi-Fi users the same capabilities of partitioning resources to support private and public devices [WBA]. The use of Basic Service Set Identifier (BSSID) can support Wi-Fi slices with a single BSSID or multiple BSSIDs as shown in Figure 10 and Figure 11 respectively.

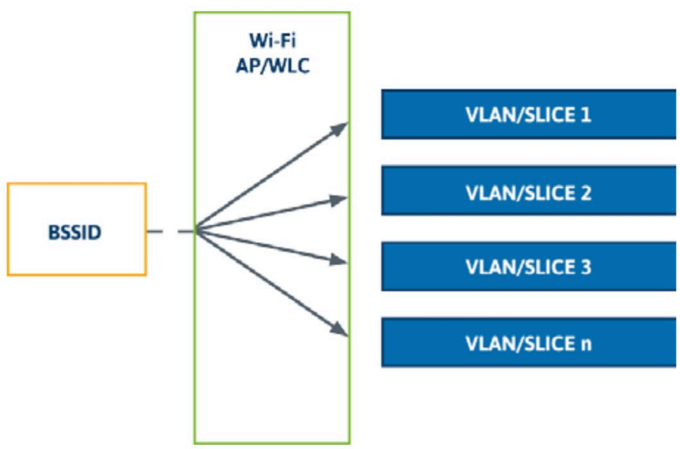

Figure 10: Slice support using single BSSID [WBA2018]

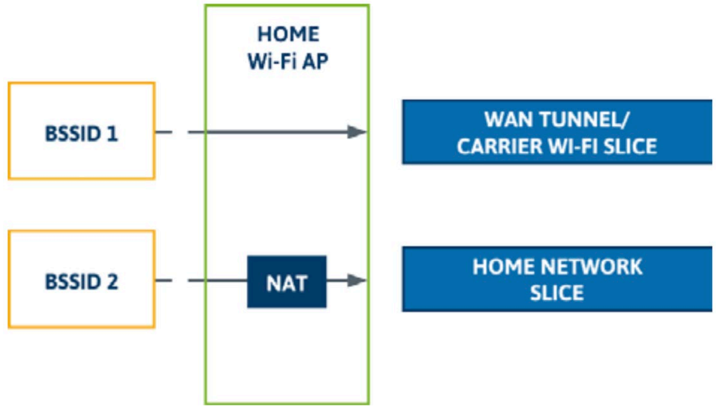

Figure 11: Slice support using multiple BSSID [WBA2018]

IEEE Wi-Fi has defined many network slicing capabilities, and the 2018 WBA paper defined other examples listed below:

- The ability to move a Wi-Fi device from one network slice to another, and to remove a UE from a network slice
- Ability to isolate traffic between different network slices in the same network
- Ability to define resources for a network slice
- Ability to define prioritization between slices, in case network resources become over-subscribed
- Ability to enable a Wi-Fi device to be simultaneously connected to more than one network slice
- Slicing of Wi-Fi Core Networks and Transport Networks
- Management and Orchestration of Sliced Wi-Fi Networks

The use of Wi-Fi and PON slicing could be a compelling solution for service providers to offer 5G mobile offload to several or all the mobile operators. A community Wi-Fi service offering could be offered by the service provider. As described in the PON slicing section there are several methods that could be used to steer Wi-Fi traffic to PON slices.

3.4. IETF Network Slicing Overview

Service providers may want to enhance the end-user's experience and their service offerings with the use of network slicing across the packet network (IP/MPLS). The physical network, in this case, is the provider edge, distribution and core network layers, that can be partitioned into multiple logical networks or network slices for specific services or customers. In the packet network domain virtualizing the network logically has been done for over two decades, with the use of layer 2 or layer 3 virtual private networks (VPNs) and with the use of traffic engineering (TE). Network slicing builds on VPNs and TE, with service guarantees such as throughput, latency, and jitter, that will not impact other slices [SANOG36]. Additionally, network slicing can reserve resources such as bandwidth and perform network isolation for performance, traffic separation, security, and privacy [SANOG36]. Finally, end-to-end multi-domain service and network orchestration and can manage domain controllers and network slices [SANOG36].

The IETF has defined network slicing with the use of segment routing to include deterministic capacity, latency and reliability [ECI]. There are two segment routing forwarding or data plane instantiations choices these are 1) SR-MPLS (MPLS data plane) or 2) SRv6 (Segment Routing over IPv6 data plane). There are some shortcomings with SRv6 segment identifier headers and segment routing mapped to IPv6 referred to as SRm6 and these are considered to have been solved according to Juniper Networks

[Juniper]. The use of SR-MPLS data plane may be a preferred path by some service providers as this leverages the mature MPLS hardware with likely software upgrades [Filsfils]. The use of SRv6 does not use the MPLS data plane and may have 66 percent less data plane entries and counters and does not use of RSVP-TE for TE/FRR [Filsfils].

Beyond data plane the IETF further defines SR policy with service level agreement (SLA) for bandwidth and latency parameters. SR uses source-based routing and path computation element (PCE) for precise, deterministic paths to be created across the network [ECI].

Segment Routing uses Traffic Engineering (SR-TE) and Flexible Algorithm (Flex-Algo). The use of Flex-Algo enhances SR-TE on-demand next hop (ODN) and Automated Steering traffic for intent-based instantiation of traffic engineered paths [Filsfils]. The determination of delay uses a probe measurement at both ends of the network, with PM query and PM response packets. The network slices using SR will have a three-tiered delay service capability to include [Filsfils]:

- Minimizing Routing Cost Metric (Low Cost Network Slice)
- Minimizing Delay (Low Delay Network Slice)
- Minimizing Cost with Maximum Delay Bound Slice

The OLT is placed at the edge and may serve as a customer edge (CE) or a provider edge (PE) connecting with either the aggregation or distribution layers that then connects to the core network. Service providers that use a layer 2 OLT will use VLAN hand offs and L3 OLTs can perform provider edge (PE) functions and participate in the IETF network slicing architecture using SR as shown in Figure 12.

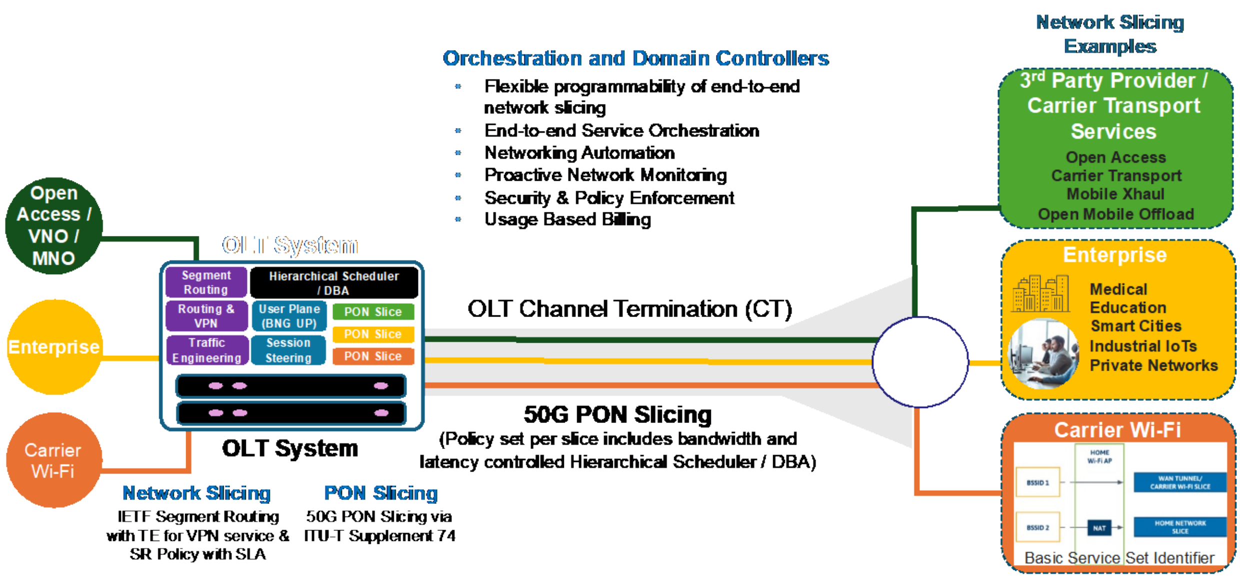

Figure 12: End-to-End Network Slicing with Integrated OLT System Architecture

In Figure 12, there are several network slicing service examples that could be supported by PON slicing, as shown on the far right of this figure. This shows an end-to-end network slicing architecture that combines Wi-Fi, PON slicing, and IETF network slicing using segment routing. The consolidation of network functions to the OLT system, as shown in the figure above, enables the operator to configure and managed both PON slicing and IETF network slicing on the same network device. This enables the slicing from the customer premises to the OLT, and the OLT then instantiates the IETF network slice. This greatly simplifies the network operationally and reduces the total cost of ownership.

4. Service and Network Drivers

We categorized the market into three (3) segments to include business services, aggregation services, and residential services, as seen in Table 1. These segments were defined into use cases with associated data rates and then measured them against each PON technology type.

In the Wi-Fi 7 use case a range of 30 Gbps to 46 Gbps is shown. This considers the theoretical maximum of Wi-Fi 7 which is 46 Gbps, although actual results are expected to be much lower. According to the Wi-Fi 7 standard IEEE P802.11be amendment document called, IEEE Standard 802.11-2020. This defines a "standardized modifications to both the IEEE Std 802.11 physical layers (PHY) and the Medium Access Control Layer (MAC) that enable at least one mode of operation capable of supporting a maximum throughput of at least 30 Gbps, as measured at the MAC data service access point (SAP)" [IEEE P802.11be]. This is why both 10G (XGS-PON) that has a maximum capacity of approximately 8.5 Gbps and 25G (25GS-PON) has a maximum of 21 Gbps are both shown in the table to not support the Wi-Fi 7 use case. However, 50G (ITU-T 50G PON) that has 42 Gbps of usable capacity is shown to support the Wi-Fi 7 use case. It is likely that initial Wi-Fi 7 products may support 20+ Gbps, however future versions may see higher data rates.

The 5G midhaul and backhaul use cases are sourced using data rates found in the O-RAN Open Xhaul Transport Working Group 9 specification. This technical specification defined the requirement for midhaul and backhaul to share the same capacity projections for the site types defined in Table 1 [O-RAN WG9]. The O-RAN working group defined 5G Xhaul with a conservative provisioning bound for both medium and large sites that use 3 sectors [O-RAN WG9]. The small sites use a single sector and peak rates were used [O-RAN WG9].

The source materials for IEEE Wi-Fi 7 and O-RAN mid/backhaul are fully sourced in the bibliography & references section of this document. Table 1 illustrates the use cases and the capabilities of XGS-PON (10G), 25GS-PON (25G), and ITU-T 50G PON (50G) technologies to support each use case. The use of 50G PON will support all the use cases.

Table 1: Service Use Cases and PON Technology Assessment

Segment	Service / Aggregation Site	Gbps	10G	25G	50G
Business Services	Business Max Service Tier and Peak Traffic	<8.5	✓	✓	✓
	True 10G Services	10	✗	✓	✓
	True 25G Services	25	✗	✗	✓
Aggregation Services	Wi-Fi 6 / Wi-Fi 6e Access Point Transport	9.6	✗	✓	✓
	Wi-Fi 7 Access Point Transport	30 – 46	✗	✗	✓
	5G Mid/Backhaul - Small Site	2.0 – 5.7	✓	✓	✓
	5G Mid/Backhaul - Medium Site	15.2	✗	✓	✓
	5G Mid/Backhaul - Large Site	36.8	✗	✗	✓
	Transport To/From MDU/Remote Cabinet/Node	10 – 40	✗	✗	✓
Residential Services	Max Service Tier and Peak Traffic	<8.5	✓	✓	✓
	Max Service Tier and Peak Traffic	>8.5	✗	✓	✓

5. Traffic Engineering and Capacity Planning

5.1. Residential Service Tier Growth

Nielsen's Law of Internet Bandwidth states that the highest service tier or speed tier offered to consumers grows at a 50% compound annual growth rate (CAGR). In Figure 13, we use Nielsen's y-axis logarithmic scale to show the exponential growth of 50% annualized growth of high-end user's connection speed [NielsenLaw]. According to Nielsen, the data analysis beginning with the acoustic 300 bit per second (bps) modem in 1984 [NielsenLaw]. We have extended the logarithmic scale to the year 2044 to illustrate the highest service tier or speed offered to consumers would be over these next 20 years, as seen in Figure 13. As Nielsen started in 1984 with a 300-bps modem Figure 13 illustrates the values on the logarithmic scale every 5 years beginning in 1984 and running through 2044.

Nielsen's Law has been fairly accurate for much of the last 40 years, with actual service provider high-end user's connection speeds growing near or above Nielsen's 50% CAGR. In recent years, service providers have accelerated service tier or connection speed growth far exceeding Nielsen's Law of 50% CAGR logarithmic scale. In 2022 and 2023 several service providers launched a 5 and 8 Gbps service tiers, exceeding Nielsen's Law, this leap in service tier growth was a result of the adoption of XGS-PON. In early 2024, Google Fiber launched a 20 Gbps residential service tier, far exceeding Nielsen's Law growth rate projection of 3.3 Gbps. Though not labeled in the figure below, Nielsen's Law logarithmic scale projects by the year 2030 a 38 Gbps service tier, so if a service provider deployed 50G PON this could be supported.

Our prediction is that Nielsen's Law of 50% CAGR will not continue for another 40 years, moreover we predict a significant decline in the CAGR of service tiers offered in the 2030's. Our prediction of Nielsen's Law obsolescence is based on many factors such as, a continued 50% CAGR of the highest service tier will not be noticeable or needed by consumers. Additionally, the cost to service providers to enable the access network to sustain the Nielsen 50% CAGR for service tier (speed) through the 2030's will not be economically sustainable. For example, extending Nielsen's Law forecast to the year 2044 would see a service tier of a Terabit per second (Tbps) between 2038 and 2039 and 11 Tbps by 2044.

Service providers offering a top service, or speed tier, to consumers will not continue to grow at 50% CAGR forever, our prediction is that Nielsen's Law will break likely in the 2030s.

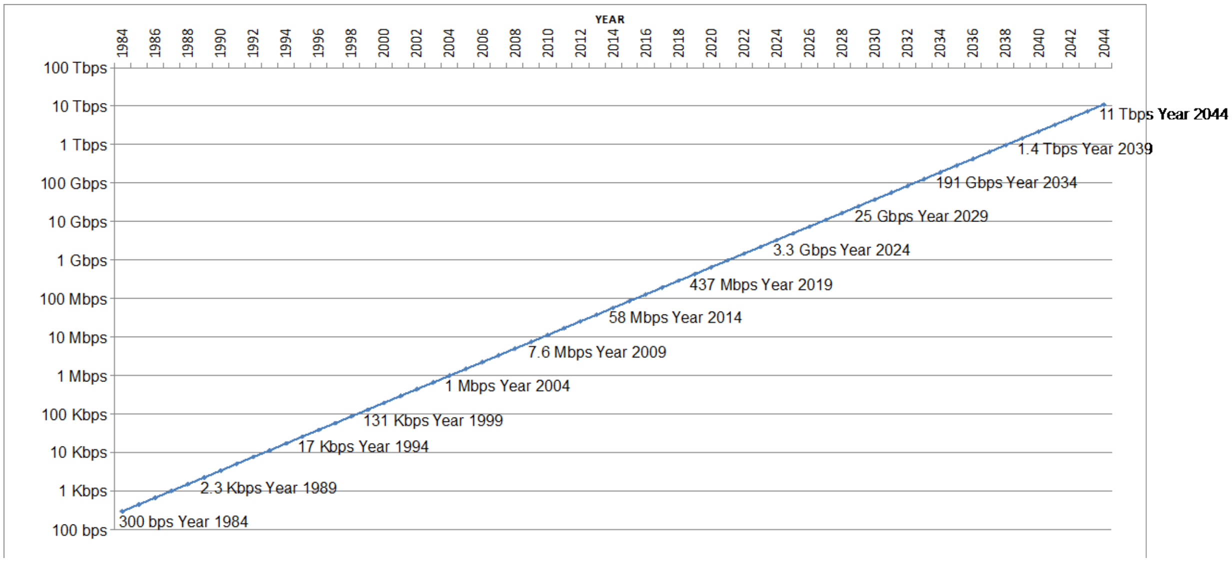

Figure 13: Nielsen's Law of High-end User's Connection Speed at 50% CAGR

5.2. Residential Per Subscriber Traffic Usage Growth

Nielsen's Law covers speed tier growth rates, but there is another critical factor needed for capacity planning of networks, systems, and technologies, this is the traffic or usage growth rates during peak periods. This measures the traffic through the system during peak periods as well as over time, to calculate a per subscriber traffic or bandwidth that is then used to determine a compound annual growth rate (CAGR) of traffic, which is critical for future planning. We researched the busy hour busy day (BHBD) traffic or bandwidth per subscriber data rates back to the year 2000 [BHBD Sources]. These BHBD traffic calculation are referred to as kilobits (Kbps) per subscriber or today known as megabits (Mbps) per subscriber as measured during peak traffic periods. Over several decades of collecting peak period traffic calculations from several public sources this data can calculate the traffic per subscriber and over time to determine a CAGR, see Figure 14. We use traffic per subscriber data and the number of subscribers sharing a network / port to calculate the capacity during peak traffic periods. We also use this data for network and technology planning. The figure shows in the year 2000 the Kbps per subscriber during peak periods was 6.176 Kbps and by the year 2022 it was 3,500 Kbps per subscriber or 3.5 Mbps per subscriber, this period had a 33.4% CAGR for user traffic [BHBD Sources]. The figure also plots different time intervals for measuring the CAGR. The figure shows a prediction of a traffic CAGR of 20% and when applied from 2022 to 2030 this estimates 15 Mbps per subscriber BHBD and by 2040 estimates 93 Mbps per subscriber BHBD.

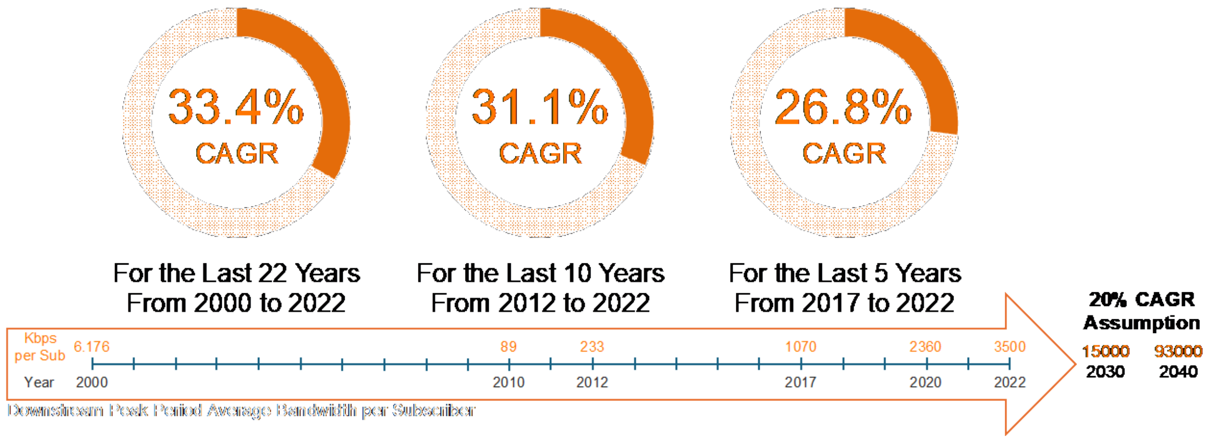

Figure 14: Historical Traffic Growth Rates over 22 Years [BHBD Sources]

5.3. Residential Service and Usage Growth Predictions

There are many applications that may continue to drive traffic growth rates. These could include the transition of over-the-top video services, or Internet Protocol television, that could move from 1K to 4K streams and even 8K streams in the future. This would mean a significant increase in network utilization because of the increase in the bit rate of the streams as seen in Table 2 and the streaming bit rate values are sourced from several references [Streaming] [Netflix] [8K Streaming]. There is a substantial increase in bit rates per stream when moving from 1080p or 1K to 4K and then from 4K to 8K. The ITU-T supplement 74 uses video streaming bandwidth after encoding for 4K at a bit rate of 54 Mbps, and 8K programs ranging from 80 Mbps to 140 Mbps per stream, both higher than this table illustrates [ITU-Sup74]. The table below was compiled from various sources and has a much lower 4K and 8K streaming bit rate than the ITU-T Sup74 forecasts. The transition of video streams from 1K to 4K and then 4K to 8K will influence traffic growth rates in the future. Network planners will need to make sure that the network technologies selected will have the capacity to support future service tier and traffic growth rates.

Table 2: Streaming Platform Data Rate Projections

Streaming Platform	SD (480) or HD (720p)	High Definition (HD) 1080p	Ultra High Definition (UHD/4K)
Netflix	3 Mbps	5 Mbps	15 Mbps
YouTube	3 Mbps	7 Mbps	15 Mbps
Hulu	1.5 Mbps	3 Mbps	8 Mbps
Amazon Prime Video	0.9 Mbps	3.5 Mbps	25 Mbps
Disney+	5 Mbps	10 Mbps	25 Mbps
HBO Max / MAX	5 Mbps	10 Mbps	25 Mbps
Apple TV+	1 Mbps	6 Mbps	25 Mbps
Paramount+	1.5 Mbps	3 Mbps	25 Mbps

8K Ultra HD 30 - 50 Mbps per Stream

Table 3 illustrates the virtual reality (VR) throughput and latency targets according to the Wireless Broadband Association (WBA) Annual Industry Report 2023 published in October 2022 [WBA2022]. The paper cited Gartner that predicted by the year 2026, that 25% of people will spend at least one hour per day in a virtual shared space, thus driving enormous pressure on home Wi-Fi networks and access networks [Gartner]. The introduction and use of VR represents a change in consumer behavior that will influence traffic growth rates in the future.

Table 3: Virtual Reality (VR) Throughput and Latency [WBA2022]

ESTIMATE THROUGHPUT AND LATENCY FOR VR/AR TECHNOLOGIES

	VR Resolution	FPS	Equivalent Resolution	Maximum Throughput (Mbsp)	Maximum Streaming Latency (ms)	Maximum Interactive Latency
Early VR	1K X 1K	30	240p	25	40	10
Entry VR	2K X 2K	30	SD	100	30	10
Advanced VR	4K X 4K	60	HD	400	20	10
Extreme VR	8K X 8K	120	4K	1000-2350	10	10

As shown in Figure 14, the traffic growth rates seem to be declining as we look at the recent 5 years of traffic growth, from 2017 to 2022 showing a 26.8% CAGR. Considering a longer duration analysis such as the last 22 years this has a higher CAGR of 33.4%. If the migration of streaming service video moves from 1K to 4K and then 4Kto 8K in the period from 2022 to 2040, then video streaming could cause a per user traffic increase, however the consumer may not be watching more streams. This means that through no change in consumer behavior the capacity increases are caused be machines using more data, in this case 4K and 8K streams instead of 720p or 1K streams. The use of VR that uses massive streaming bandwidth as well as long duration sessions could be yet another driver for continued traffic growth. Considering the last 22 years had a 33.4% CAGR and the last 10 years had a 31% we needed to find a value that was reasonable for a forecast from 2022 to 2040, an 18-year span, and our estimates below use a 20% CAGR for this time duration. Our traffic engineering models are highly flexible, and we have modeled many other traffic growth rates.

As shown in Figure 15, the traffic CAGR of 20% is applied to several subscriber group counts including 32, 64, 128, 160, and 192 sharing a capacity channel. These traffic projections are not bound to a particular access technology, though the subscriber count per shared link of 128, 160 and 192 may be found in DOCSIS® networks and not PON. The traffic projections in this analysis grow at 20% CAGR through 2040. Figure 15 shows a downstream usable GPON capacity limit of 2.3 Gbps and XGS-PON link capacity limit of 8.5 Gbps. When the traffic lines are below the capacity limits this shows the available capacity for service tier to pass a speed test during busy hour busy day (BHBD). The service tier of 1 Gbps is used to show the project time period when GPON may reach the limit to support that service tier, assuming the 20% traffic CAGR and 32 subscribers sharing a link. The analysis shows that GPON may support a 1 Gbps service tier under these assumptions until 2035 or 2036 as shown in Figure 15. A similar analysis is performed for XGS-PON considering an 8 Gbps service tier and a 5 Gbps service tiers. Note from Figure 15 that a 5 Gbps service tier is supported a decade longer than 8 Gbps. A service provider could extend the life of GPON and XGS-PON by moving the higher service tiers and top traffic users up to the higher capacity PON technology, leaving lower service tiers customer on GPON or XGS-PON. For example, if a service provider launches an 8 Gbps service tier on XGS-PON and at some point, the BHBD traffic prevents passing a speed test, the operator could move those 8 Gbps subscribers to 50G PON to extend the life of XGS-PON, at some point 5 Gbps runs out and those subs could be moved to 50G as well. Next generation PON technology needs to have enough of a capacity increase to support new services as well as keep up with service and traffic growth rates and have a long useful life.

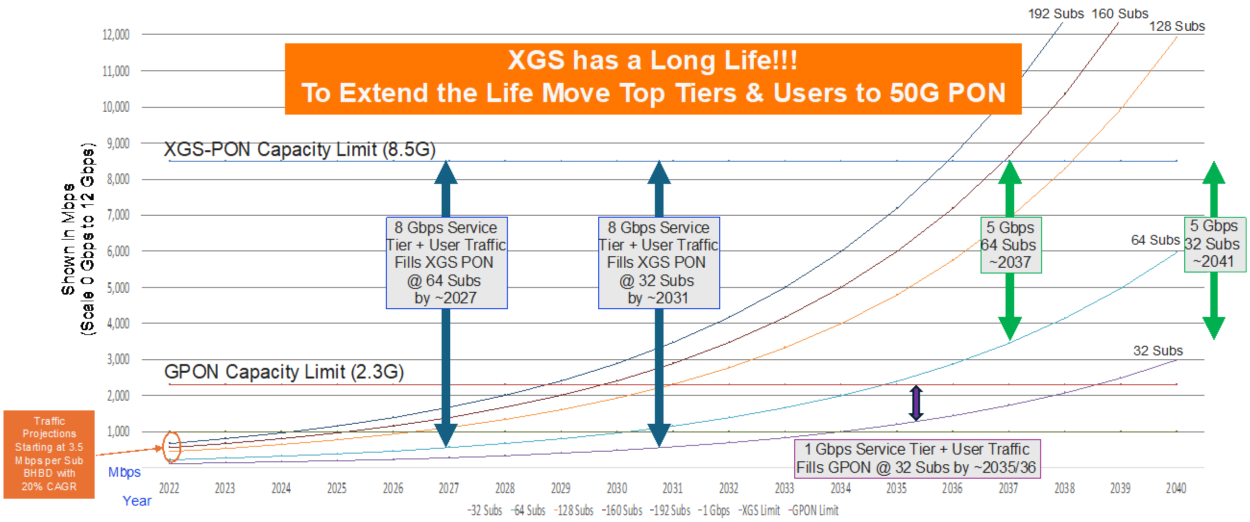

Figure 15: Traffic CAGR of 20% with GPON and XGS Available Capacity Projections

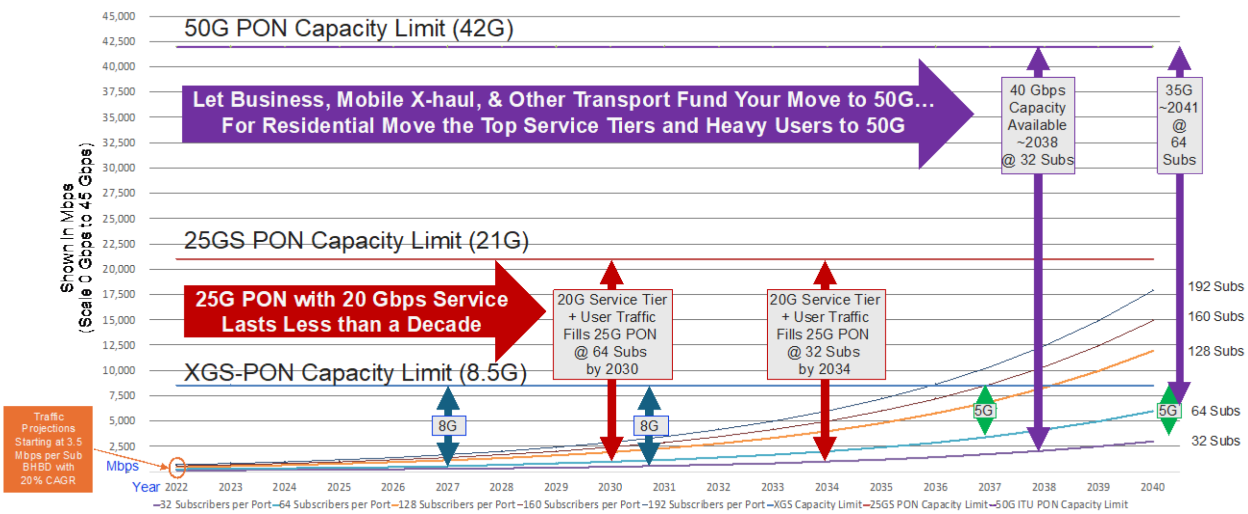

Figure 16: Traffic CAGR of 20% with 25GS PON and 50G PON Available Capacity Projections

In Figure 16, a downstream usable capacity limit of 21 Gbps for 25GS-PON and 42 Gbps capacity limit for 50G PON are considered. If a service provider launches 25GS-PON in new markets instead of XGS-PON, then launches a 20 Gbps service tier, at some point the BHBD traffic prevents passing a speed test. The forecast based on 64 and 32 subscribers sharing a 25GS-PON port are shown in Figure 16, and this forecasts the year a speed test may not be passed using the assumptions described above. This figure also calculates the available capacity at different time periods and subscriber count for 50G PON as well.

6. Conclusion

Service providers are deploying XGS-PON in large scale as of the date of this publication. A few service providers have launched 25GS-PON while others are waiting for ITU-T 50G PON for many reasons described in this paper.

ITU-T 50G PON has the Capacity! The use of ITU-T 50G PON meets the service provider's current and future use cases for business services such as 10 Gbps and 25 Gbps, mobile Xhaul, IEEE Wi-Fi 7 access point transport, aggregation layer functions, and future residential service tier increases.

ITU-T 50G PON has better technology and economic flexibility! The use of ITU-T 50G PON specifies a cost-effective single channel 50G downstream. 50G also specifies that 50G OLTs can have a dual-rate receiver to support both 50x50 ONUs and 50x25 ONUs on the same OLT interface using the same wavelengths. This compelling feature enables economic flexibility for service provider to use one OLT interface port and have a choice of symmetric or asymmetric ONUs with likely different price points.

ITU-T 50G PON supports ITU-T PON slicing! The ITU-T supplement 74 PON slicing is an optional specification and when implemented in the 50G PON OLTs and ONUs this will enable programable PON slices of capacity, QoS, and latency for groups of subscribers. 50G PON slicing technology efficiently uses capacity above guaranteed to be shared by others member of the slice and across the entire PON interface. The use of PON slicing is cost effective compared to optical Ethernet that dedicates wavelengths, ports, space, and power, per customer and even if just a little capacity is used. The use of 50G PON and PON slicing can unlock new revenue streams, while reducing capital and operational costs. The consolidation of network functions to the OLT system, such as BNG and provider edge functions means that the OLT system can support both the PON slicing and IETF network slicing domains.

Abbreviations

50G PON	ITU-T 50G PON
AE	Active Ethernet
AP	Access Point
APON	Asynchronous Transfer Mode (ATM) PON ITU-T G.983
ASIC	application-specific integrated circuit
BBF	Broadband Fourm
BHBD	Busy hour busy day
BPON	Broadband PON ITU-T G.983
Bps	bits per second
BSSID	Basic Service Set Identifier
CAGR	compound annual growth rate
CPE	customer premises equipment
CT	Channel Terminations
DBA	Dynamic Bandwidth Assignment (G.984, G.9807, 9804, Suppl 74)
DOCSIS	Data over cable system interface specification
E2E	end-to-end
FEC	Forward Error Correction
FTTH	Fiber To The Home
FWA	fixed wireless access
Gbps	Gigabits per second
GPON	Gigabit-capable Passive Optical Networks ITU-T G.984
HSP	Higher Speed PON
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
ITU-T	International Telecommunication Union - Telecommunication Standardization Sector
LDPC	Low-density Parity Check
Mbps	Megabits per second
mmWave	Millimeter wave
MPLS	Multiprotocol Label Switching
MSA	Multi-Source Agreement
nm	nanometer
ODN	Optical distribution network
OLT	Optical Line Termination
ONT	optical network termination
ONU	optical network unit
PE	provider edge
PMD	Physical Medium Dependent
PON	Passive Optical Network
QoS	quality of service
RAN	radio access network
RG	residential gateway
SNI	system network interfaces
SR	segment routing
SSID	service set identifier
Tbps	Terabit per second
TC	transmission convergence

TE	traffic engineering
UE	user equipment
UNI	user network interface
VPN	virtual private networks
VR	virtual reality
WAN	wide area network
WBA	Wireless Broadband Alliance
XGS-PON	10 Gigabit PON

Bibliography & References

[BHBD Sources] Figure 14: Historical Traffic Growth Rates over 22 Years [BHBD Sources]

- Source for 6.176 Kbps per subscriber BHBD in the year 2000, "Bandwidth Monitoring Parameters for Capacity Management", page 3, "200 or 300 customers per DS-1", (used the average in this model), Dennis Cleary, NCTA 2000).
- Source for 89 Kbps per subscriber BHBD in the year 2010, 233 Kbps in the year 2012, and 1070 Kbps in the year 2017, "Traffic Engineering in a Fiber Deep Gigabit World", Ulm, et al., Cable-Tec Expo 2017.
- Source for 2.36 Mbps per subscriber BHBD in the year 2020 and 3.5 Mbps in the year 2022, "Broadband Capacity Growth Models", Ulm, et al., Cable-Tec Expo 2022

[ECI] Network Slicing, Cut a long story short, ECI, www.ecitele.com

[Filsfils] Clarence Filsfils, "SRv6 Standardization Deployed at Scale", November 18, 2021, Web https://www.segment-routing.net/tutorials/2021-11-18-CKN-SRv6/

[G.9804.1] Rec. ITU-T G.9804.1 (2019)/Amd.1 (08/2021)

[Gartner] Gartner, Gartner Predicts 25% of People Will Spend At Least One Hour Per Day in the Metaverse by 2026, February 7, 2022, Web Source https://tinyurl.com/yb6g5lxr

[IEEE P802.11be] P802.11be Amendment to IEEE Standard 802.11-2020, Wi-Fi 7, Web, https://mypr-nodejs.standards.ieee.org/mypr-file/par/6886/mypr

[ITU-Sup74] Supplement 74 to ITU-T G-series Recommendations Network slicing in a passive optical network context

[Juniper] Juniper Networks, What is segment routing? https://www.juniper.net/us/en/research-topics/what-is-segment-routing.html

[NielsenLaw] Jakob Nielsen, April 4, 1998 · Updated Jan. 23, 2023, "Nielsen's Law of Internet Bandwidth", A high-end user's connection speed grows by 50% per year, Nielsen Norman Group Nielsen Norman Group, Web Source https://www.nngroup.com/articles/law-of-bandwidth/

[O-RAN WG 9] O-RAN Open Xhaul Transport WG 9 - Xhaul Requirements, O-RAN.WG9.XTRP-REQ-v01.00 Technical Specification https://orandownloadsweb.azurewebsites.net/specifications. Frequency Range (FR1) refers to frequencies below 7.225 GHz and (FR2) refers to frequency bands from 24.250 GHz to 52.6 GHz spectrum (also referred to as "millimeter wave range"). Refer to Table 13: Last mile provisioning for 5G Backhaul and Conservative provisioning bound for medium and large sites due to 3 sectors (peak used for small sites due to single sector)

[SANOG36] Dhruv Dhody, "Network Slicing & related work in IETF", Web, https://www.sanog.org/resources/sanog36/SANOG36-Conference-ietfnetworkslicing_Dhruv.pdf

[Streaming] What is Good Internet Speed Needed for Streaming? By Gank Content Team, June 16, 2023, https://ganknow.com/blog/internet-speed-needed-for-streaming/

[Netflix] Internet connection speed recommendations, https://help.netflix.com/en/node/306

[8K Streaming] How Much Bandwidth Will You Need to Deliver 8K? By Streaming Media Editorial Staff Short Cuts, May 22, 2019, https://www.streamingmedia.com/Articles/ReadArticle.aspx?ArticleID=131687

[VZ] Verizon, What is 5G network slicing?, https://www.verizon.com/business/resources/articles/s/5g-network-slicing-do-you-have-the-team-you-need/

[WBA2018] WBA, *"Network Slicing, Understanding Wi-Fi Capabilities"* Web, https://www.wballiance.com/wp-content/uploads/2018/03/Network-Slicing-Understanding-Wi-Fi-Capabilities.pdf

[WBA2022] WBA Annual Industry Report 2023 Industry Reports; October 2022 via Source Mangiante. https://wballiance.com/resource/wba-annual-industry-report-2023/