Convergence of Services Using Network Slicing

A Practical Implementation

A Technical Paper prepared for SCTE by

Fernando X. Villarruel
Chief Architect, MSO Practice
Ciena Corporation
1185 Sanctuary Pkwy, Alpharretta
fvillarr@ciena.com

Geoff Eaton
Senior Solutions Architect
Ciena Corporation
2351 Alfred-Nobel Blvd, Saint Laurent, Quebec
geaton@ciena.com

Marco Naveda
Senior Director, Architecture
Ciena Corporation
3500 Carling Ave, Ottawa, Ontario
mnaveda@ciena.com

Table of Contents

List of Figures

1. Introduction

By now the convergence of service signals over one digital MSO network needs no further motivation. Over the last some years there have been resources to describe the financial and engineering benefits of this transition (Chamberlain, 2017), (BAUMGARTNER, 2019), (Chamberlain, 2017).

In this paper we turn to the mechanism for creating the right form of infrastructure to make convergence an implementable reality. According to (Chamberlain, 2017) there are three requirements to make convergence successful: powering, backhaul and site availability. We find this a good generalization where power refers to the active power availability in the MSO outside plant, and site availability implies the fiber, coax, and pole access real estate to work with. In this paper we focus particularly on the backhaul section for successful convergence, which includes the connectivity of hubs and customer endpoints with the recent differentiator of digital fiber, see Figure 1 below.

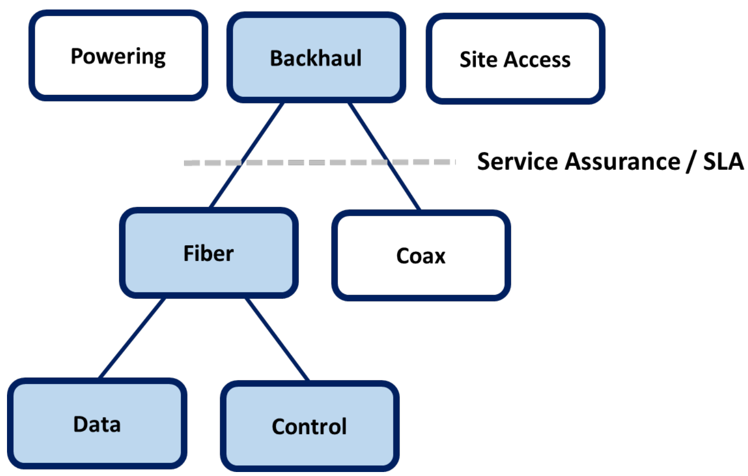

Figure 1, Infrastructure Needs for Convergence

This paper more specifically focuses on the transmission and management of multiple services over the fiber infrastructure, otherwise known as the converged interconnect network (CIN), spanning the fiber access endpoints or aggregation points to the signal packet cores at hubs or headends (Villarruel, 2020). While our focus is on the fiber networking, we go beyond to point out general implications and expectations for service assurance—code word for service layer agreements (SLA). The combination of transmission, management and sensitivities to service assurance is what drives the need for network slicing, thus the need for a paper that covers this topic.

Network slicing as a principle to streamline convergence of services for the cable industry has been described in useful detail in the reference paper: "Framework For Convergence Of Services On The MSO Network Using the Principles of Network Slicing" (Villarruel, 2020). This paper uses this reference paper as a launching point and further shows a practical implementation of network slicing on multiple services implemented in a laboratory setting.

We point out two disclaimers. First, we refer to network slicing as a general term, but the term network slicing also has a particular meaning within the 5G mobile space, as standard bodies like 3GPP have specific definitions for qualities of network slices for specific subservices. We use the generalized definition of network slicing because we can apply similar principles to a new space, in this case the MSO industry. Secondly, the work we show here is descriptive and not prescriptive in that there can be other ways to organize and share network resources and enforce service assurance.

2. Network Overview

The purpose of this paper is to position an end-to-end digital network that has the right sensitivities for convergence. At a high level we partition the converged network in two sections, access aggregation and core network transmission. See Figure 2 below. We position Flexible Ethernet (Flexed) technology for access aggregation and segment routing (SR) for the core network transmission. We also describe the principle of a Transition Map between FlexE and SR. This paper aims to describe the methodology for network slicing in the fiber backhaul and finally present a proof of concept from a lab buildout of the described principles.

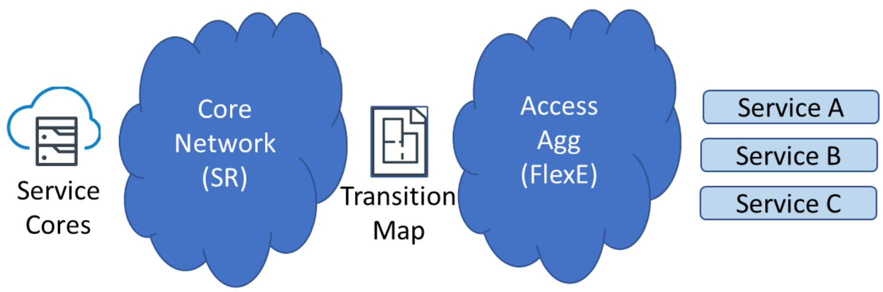

Figure 2, Network Overview

3. Convergence and Network Slicing

It is useful first review what we mean by service convergence and its relation to network slicing. Network convergence has been shown to include several access technologies with several services and subservices embedded, and span outside plant access and core networks, as seen in Figure 3 below (Villarruel, 2020).

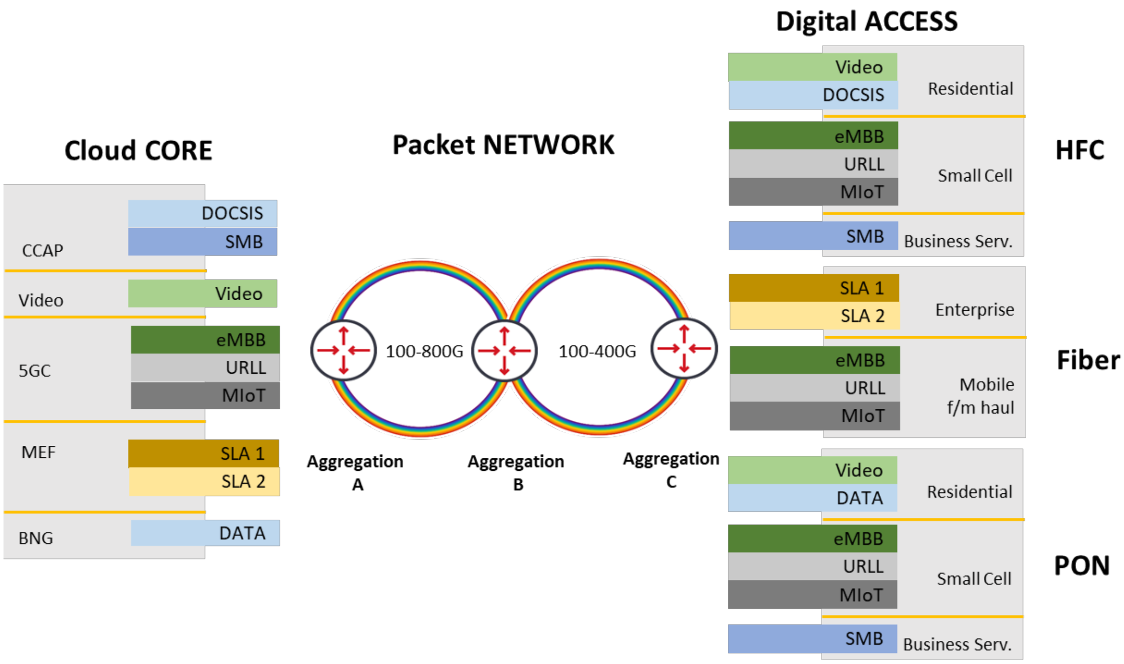

Figure 3, Convergence and Network Slicing

Notice that while Figure 3 is the view of an end-to-end system that works together, there is a delineation of domains, as each would have its own functions and expectations. See Figure 4. In the context of network slicing we call these resources the access would have its own resources, the network domain would have its resources and the core would have its resources. Connecting them might be a global orchestrator, but it is worthwhile to evaluate principles for each section in its own light. As alluded to in the introduction, the focus of this paper is on the network domain.

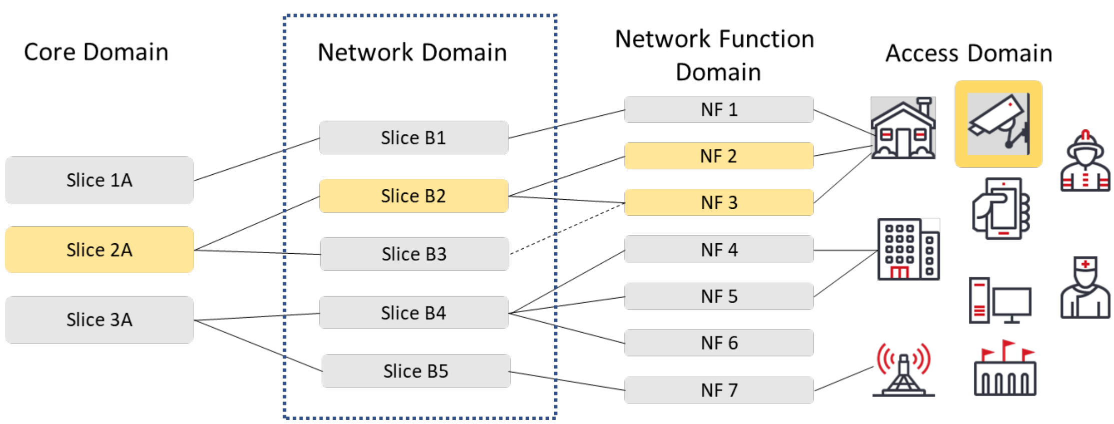

Figure 4, Slice Domains and Functions

Formally network slicing is: "a method to run multiple end-to-end logical networks on a common set of resources." But within and end to end framework there are domains which also have their own slices, so end-to-end network slicing is really a patch of slices from each domain to make a particular service. Using the example shown in (Villarruel, 2020), Figure 4, where each domain has a partition of resources that create domain slices, a subservice would use a unique collection of slices from each domain for end-to-end transmission of a subservice, otherwise known as a service blueprint. Consider the example shown of a residential camera subservice, which might need a network function of face recognition and encryption, a tunnel from the network and a unique profile from a service core. The emphasis of this paper, if we take Figure 4 as reference, are the slice types available and created in the network domain. In this case the network domain is specifically the Ethernet and IP resources available to MSO systems that have evolved to a digital fiber plant, with the advent of distributed access architectures.

For a converged system it takes some effort to create the slice profiles and instantiate services as a collection of slices and manage their evolution through a lifecycle, but the primordial challenge is how to share sources in a way that allows legacy networks the same service assurance that they've had as independent networks. This is the fundamental challenge and so it must be addressed first, before any high-level orchestration or management is determined. The question is then, how does one converge services in a digital network such that they remain autonomous? The methodology for network slicing answers this question and is thus the kernel of this paper. We particularly revolve our work around utilizing concepts of soft and hard slicing, where hard slicing describes mechanisms that are rigid in how resources are dedicated, and soft slicing refers to mechanism that are made for flexibility in the sharing of resources. For the application of these hard and soft slicing concepts we leverage Flex Ethernet (Flex-E) and managed segment routing (SR) respectively. We now take a focused view of Flex Ethernet and Segment Routing.

4. Flexible Ethernet

Flexible Ethernet, otherwise known as Flex-E, is a mechanism that was created to provide transmission rate flexibility between client and line signals. This capability addresses the variety of client rates new and installed, along with line side transmissions and their ability to work together in a network. Flex-E dissociates the client rate from the physical interface. It does this by adding a shim layer to the 802.3, 100GBASE-R protocol stack and the 802.3bs 200/400GBASE-R protocol stack, separating the MAC from the PCS and PMA layers. See Figure 5.

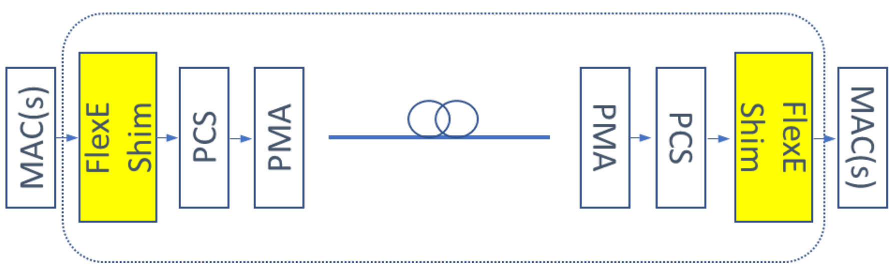

Figure 5, Flexible Ethernet Protocol

More generally, the Flex-E shim layer allows for not just one client MAC, but various, such that a collection of various clients at different rates can be transmitted over one or more Flex-E channels, see Figure 6. Specifically, FlexE clients have their own separate MAC, reconciliation layer and MII above the FlexE Shim layer, while layers below the PCS keep the intact rates as specified for their own Ethernet implementation (Optical Internet Working Forum, 22).

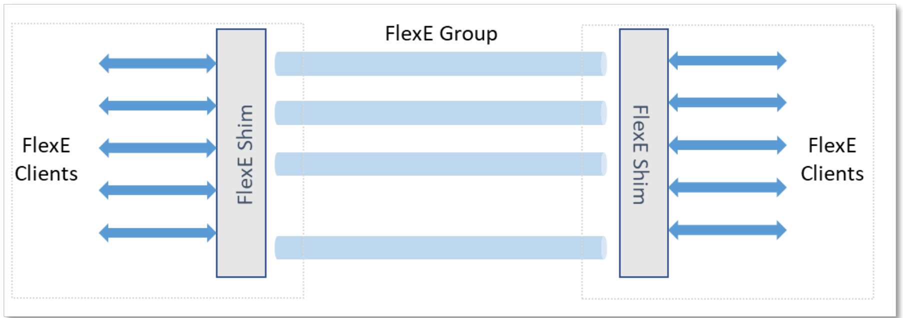

Figure 6, Clients and FlexE Group

The shim layer is shown in more detail in Figure 7; it shows the muxing function of a shim layer, where a collection of clients present themselves in a serial 66 bit stream, created from an ethernet MAC, operating at 10, 40 or mx25 Gb/s.   Entering the calendar function, the client signals are adapted to the FlexE channel rate by a clock transition and an insertion or deletion of idle blocks. The calendar consist of 66 bit based time slots, and the control function controls the sequential insertion of client blocks to the calendar, along with insertion of necessary overhead. The subcalendars are logical groupings of calendar slots that prepare for interleaving that puts together signaling for the expected rate of the FlexE channel. Show in the Figure 7 is the description of the muxing function. There is also a demuxing function which is effectively the inverse of the muxing function.

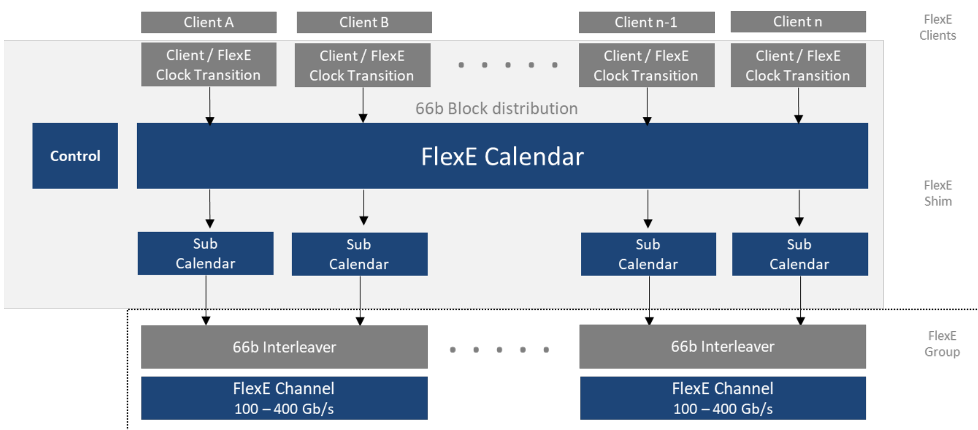

Figure 7, FlexE Shim Layer

The functionality of the FlexE calendar is interesting and a full treatment is beyond the focus of this paper, but as an example Figure 8 shows the calendar function in more detail for a 100Gb/s FlexE channel. Here the calendar has a granularity of 5G slots and every 20n slots create the sufficient 100Gb/s signal for a subcalendar. This calendar shows the serially inserted bitstream of mixed PHY clients at 10Bb/s and 25Gb/s. Similarly, the calendar can have up to 25G slots, a quality more natural for 25 G clients and greater speed FlexE channels at 200 or 400 Gb/s.

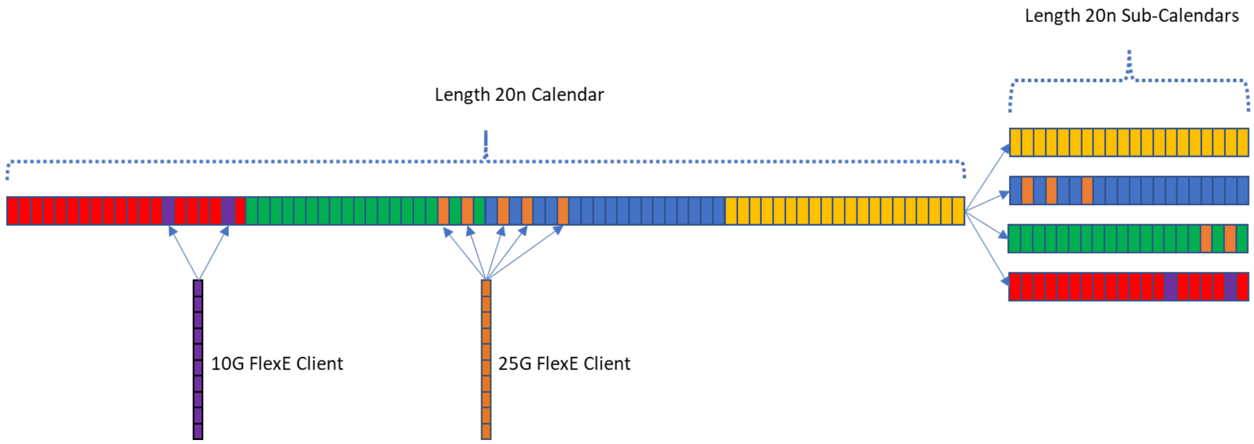

Figure 8, FlexE Calendar

With the presented characteristics FlexE has the capability of three basic functions, as shown in the Figure 9. Bonding allows the user to create a higher rate service, like 400G, with a collection of lower rate interfaces. The example shows 2, 100G Ethernet PHY signals carrying a 200Gb/s service. This is sometimes applied as an alternative to LAG without the extra layer of hashing algorithms. Subrating, service rates that don't match a greater service to be carried over on interface of a larger channel. More specifically the transmission rate is downsized in order to match the necessary client rate. In the example we show a 75G service utilizing a 100GB/s ethernet interface. And finally, there is the capability of channelization. which aggregates lower rate channels onto a higher rate interface. In the example we have clients of various native rates that manage to use one 100G FlexE channel. It is the case that these capacities can be mixed and matched depending on the circumstance.

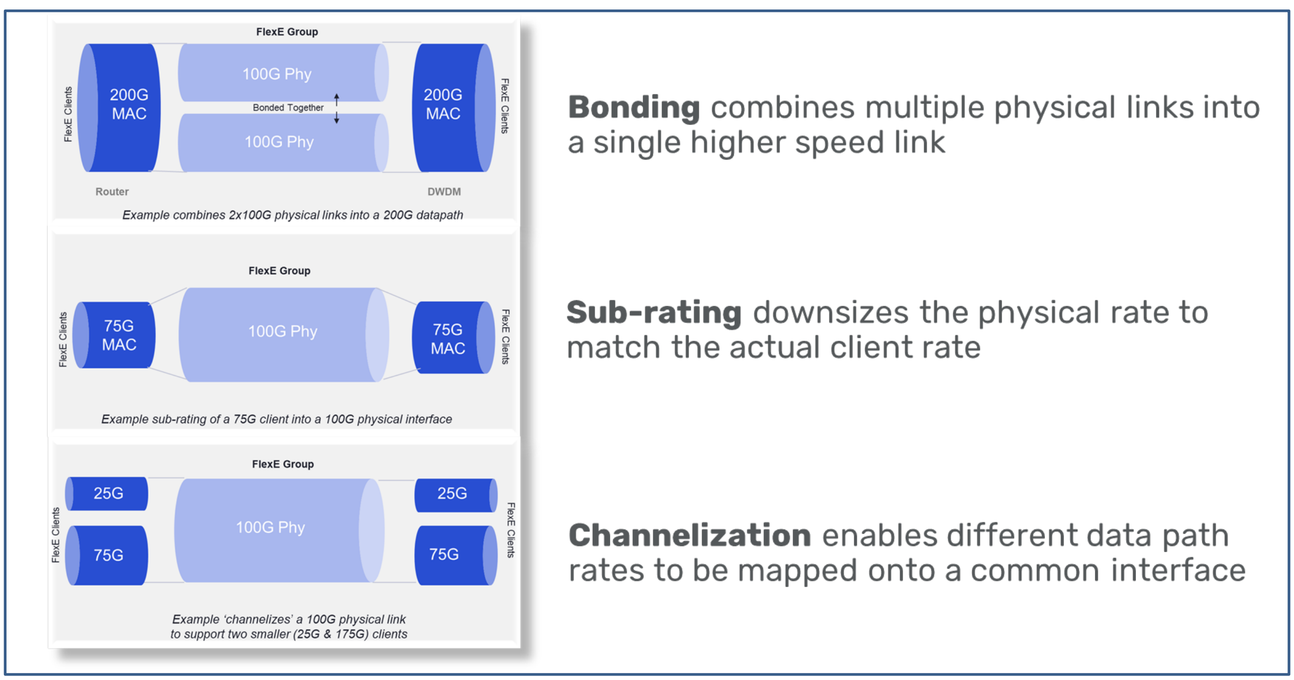

Figure 9, FlexE Functions

For our MSO convergence discussion channelization allows for the transmission of 10 and 25Gbs signals to leverage naturally higher rate backhauls, without having to go through an extensive forwarding plane, which is particularly useful in low footprint low energy platforms, such as outdoor shelves and nodes. Thus, FlexE is a good resource for the access outside plant, creating a robust combining mechanism that resembles straight forwarding muxing functions, like OTN, but with a more streamlined overhead.

But if FlexE is a good option for the outside plant access, what about the core network? This brings us to our next portion of end-to-end network slicing—segment routing. SR is an IP function which allows us to present a solution for soft network slicing.

5. Segment Routing – Data Plane

Segment routing is a topic that has been well described in other resources so here is a brief treatment reviewing basic principles (Bonica, 2017). SR has also started to show merit specifically in MSO networks as a way to enhance operational efficiencies for business services delivery (Yeo, 2019). SR per the name focuses on breaking down network topologies into segments, where a segment is a path between two routers. As seen in Figure 10, the router segments can be between two nearest neighbors like (AB) and (BC), and only need local awareness between nearest neighbors. Or segments can also describe routes that span across routers like (AC), and these segments need global awareness beyond just neighboring routers—these segments identifiers (SID)are described as adjacent SIDs or binding SID's respectively. Routing protocols can then be applied to segments as is done with regular routing that relies on unique addressing of source and destination. The routers that participate with capability to recognize SIDs and routes accordingly is called the SR domain.

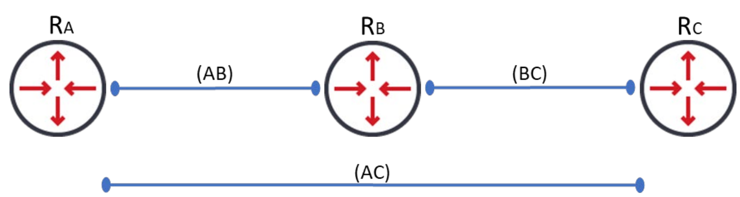

Figure 10, Segment Routing IDs

There are two methods for encapsulating SID's, one is to use labeling space in the MPLS header, and the other is for to use an IPv6 header with an SR extension. In either case the principles for inclusion of labels in headers throughout the life of a packet is the same. SR aims to keep the route of a packet known to the packet itself and reduce the calculation of routes as is done in legacy routing. This occurs by encoding the path of a packet with a stack of labels that describe the route to take and for the most part avoids large route calculations at every hop, as is done in regular routing by analyzing IP addresses and deriving paths from large look up tables.

In Figure 11 we show a very simple example of a six routers SR domain with a collection of local SID's, advertised and known only to individual nearest neighbors, and binding (global) SID's advertised and known to all (which allows packets to skip treatment at certain hops.) In Figure 12 and Figure 13 we show SR label stacks and accompanying payload progress through a route. In Figure 12 an instance using only local SID's shows a path from router Ra to router Rf. Note how the ingress packet from Ra has the whole path encoded on its label stack. At Rb, the router pops off the incoming label and forwards the packet onto its next location, at Rc the router pops off the incoming label and forwards the packet onto its next location, and so on until it reaches Rf, where only the payload is left as was desired. In Figure 13 we note that the labels used include global SID's, where label processing skips over Rb and Re, with a local label between Rc and Rd. Here Rb and Re pass the packet onto the destination known globally without any manipulation of labels. At Rc and Rd however the label stack is depleted on route to Rf where it terminates as desired. Because of this simple approach, SR relies heavily on Interior Gateway Routing protocols, which simplify the necessity to know routes beyond local networks. There are other nuances and complications to label stack building and recognition, but this example highlights the general interworking of SR.

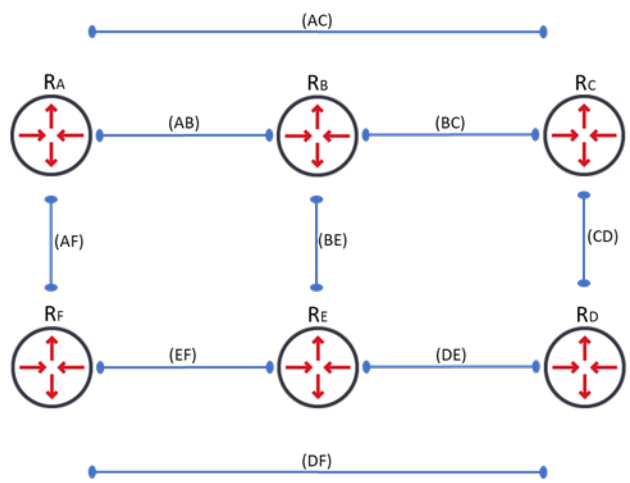

Figure 11, Routing Mesh and SID assignments

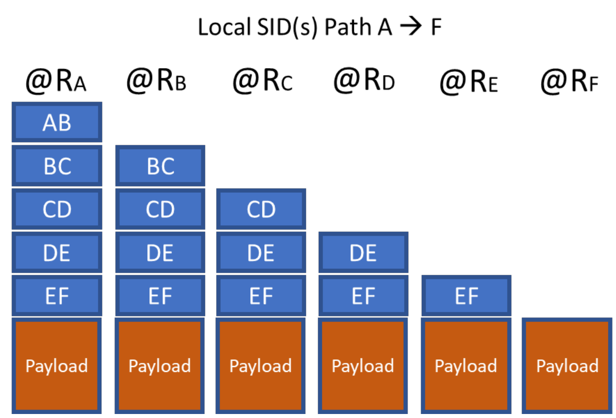

Figure 12, Local SID Path Progression

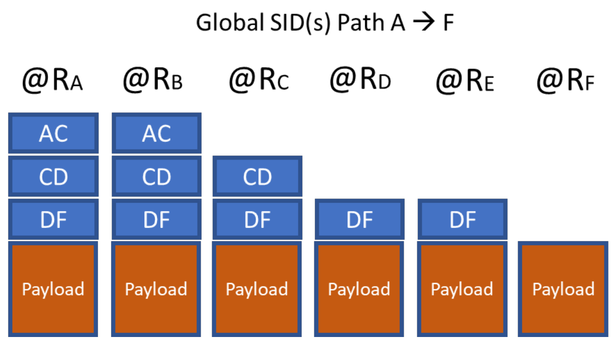

Figure 13, Global SID Path Progression

6. Segment Routing – Control Plane

In parallel to the data plane in SR there is also a control plane that calculates paths and decides what segments to visit, then communicates to routers instructions on how to build label stacks. The methodology used to control SR could allow us the ability to create an umbrella of reliable service assurance, which manifests itself as sellable SLA's, thus its importance. The decisioning of paths is done in path compute engines (PCE) where the cost of segments is evaluated according to variable state qualities such as latency or congestion, or other metrics. As shown in Figure 14, there are two schools of thought for location of path computation engines. One is to have path computation local, where each router participates in computing route characteristics and contributes cost of segments to a database that is available to all IGP participants. Each router then has the ability to influence SR labeling accordingly. The second school of thought is to have a centralized path compute engine which peers into networks and collates a global view of segment constraints based on telemetry and external influence and thus can create a wholistic cost of routes. The centralized PCE then has responsibility, not the local routers, for how label stacks are built and so directly controls the flow of information throughout the network.

For network slicing a centralized PCE, or "off-board" PCE is preferred because it can execute policy structures that could be missed if left to generic SR support mechanisms like Topology Independent Loop-Free Alternate (TI-LFA). A centralized PCE then takes control of creating SR stacks and distributing to origination points without burdening any path responsibility to middle hop routers. This is a particularly efficient way to meet service SLAs, particularly as SLA contracts can include certain non-typical networking constraints. For example, consider an example that has been considered recently, where a sensitive government service cannot cross country borders. If left to on board PCE it would be difficult to maintain data from crossing borders as pure networking mechanisms would not naturally digest that policy. With a centralized view of routes and costs expectations outside of typical network activity this important policy can be met.

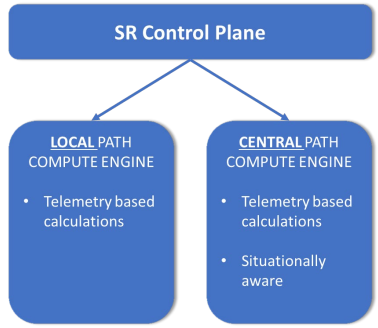

Figure 14, SR Control Plane Options

7. Transition Mapping

We have shown the operating principles of FlexE and SR, and while these are powerful independently in the case of a converged end-to-end network there needs to be an elegant and efficient, service aware, transition between the FlexE and SR domains, as shown in Figure 2. Figure 15 shows in more detail and how the Transition Map works in the context of an end-to-end network as it converges services. For Figure 15 this example we borrow some of the typical services and subservices listed in Figure 3. The transition function is a mapping between established FlexE channels and segment routed flows, where channels and flows have specific meaning and relationship to services and subservices carried by them from customer endpoints and handed off to service packet cores, then vice-versa.

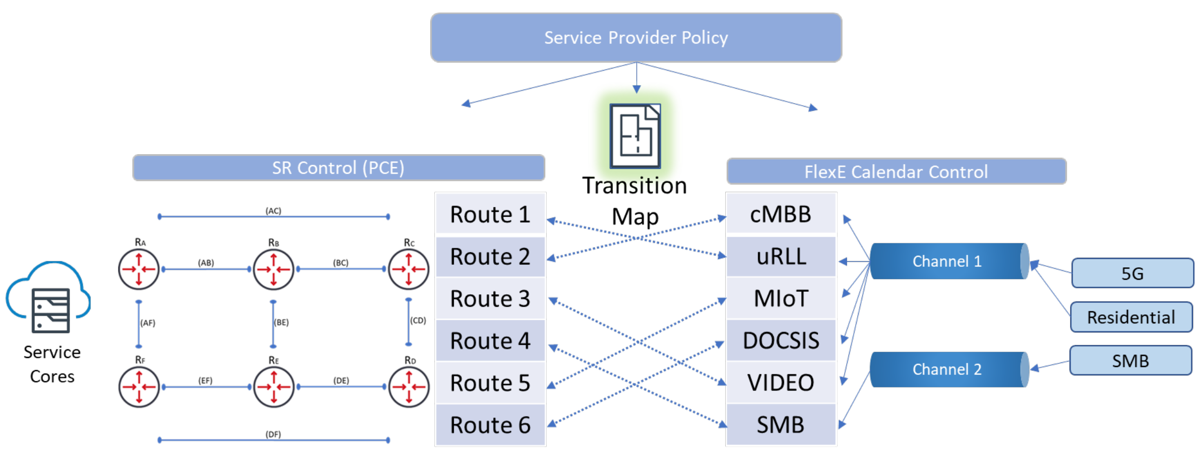

Figure 15, Transmission Mapping and Context

The transition function map executes a policy determined by the service provider which describes the relationship of subservices to SR labels and services to FlexE channels, in accordance to known SLAs. Note that the function map is intimately connected to the PCE and Calendar control functions, as they shed light to policy on how the map should be built.

In practice the transition function can be statically provisioned or dynamic. In a static set up very specific FlexE channels map directly to particular IP flows with variations over time executed by hand. On the other hand, in a dynamic state the mapping between Flex channels and SR labels could change depending on situations that allow services to best meet their SLA—again this points to the relationship of PCE and Calendar control with the policy mechanism and the transition map.

Note that SR and FlexE are well defined processes and mechanisms while the transition map is not. This implies that vendor differentiation is created around management of policy via PCE and Calendar control, and importantly the robustness of the Transition Map, how it is built and how it evolves as the system changes.

As an example, in the world of 5G, where the original work for formally defining network slicing, there is a standard list of attributes that define a slice, and that list is populated according to the particular needs of a service thus creating particular network slices, which then get assigned to subservices (GSMA, 2016). The service provider policy shown in Figure 15 would have an understanding of these slice definitions and their particular attributes and express this to the centralized PCE, the FlexE Calendar control and the transition Map where the policy is ultimately executed. In the MSO world, this is ripe definition work for CableLabs, where MSOs could benefit from Network Slicing attribute tables which could readily be assigned to subservice types.

8. Laboratory Set Up

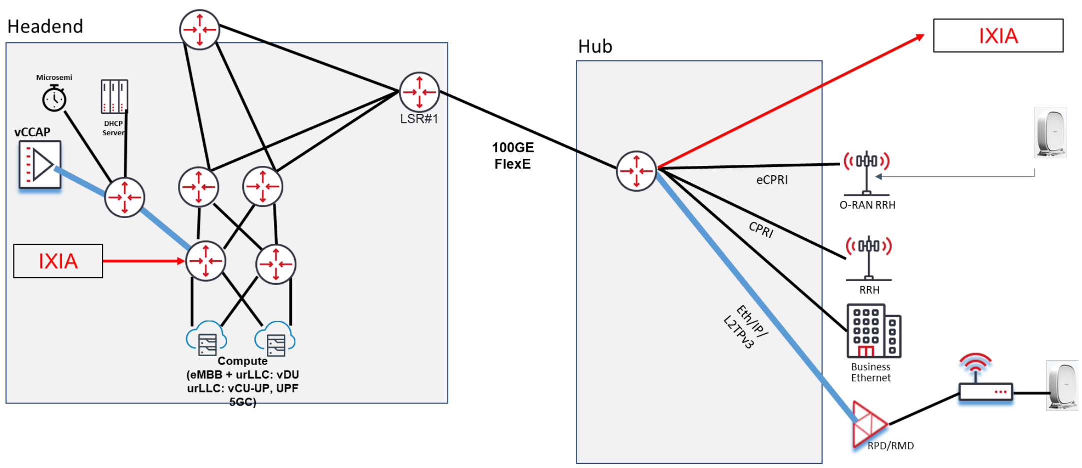

Figure 16, Planned Laboratory Setup

Figure 16 shows the planned laboratory setup with a convergence of mobile, residential and business services. Note that while this set up is built for this paper, it is also part of a more general laboratory set up for convergence in general, so some of the portions go beyond what would be needed for our testing, but useful in to verify other architecture questions.

The setup represents and-end-to end network with equivalents of outside plant, Hub and Headend. At the Headend there are the packet cores and supporting applications. There is cloud implementation of a 5G core and accompanying fabric, not in focus here. There is a virtual CCAP from a leading provider and supporting IEEE 1588 Master clock running ITU-T G.8275.2, and a DHCP server. There is also two entry routers to the Headend, an edge router located at the top, and the LSR#1 which is of our focus.

At the outside plant equivalent there is mobile signaling, from legacy backhaul in CPRI and next generation fronthaul in eCPRI supporting macro and microcells respectively. There is a point-to-point ethernet connection representative of business services. There is also a cable modem feeding into a Remote PHY Device, which ultimatately terminates at the CCAP. There is also an Ixia signal generator present to add rogue traffic which assists in testing the robustness of the FlexE and SR tunneling mechanisms.

We show that the service signals coincide on an aggregation switch, at the hub, which handles the execution of FlexE channels for all services. The aggregated signals are backhauled from the Hub to the Headend on a 100G signal, which is spent according to the FlexE channelization. The router at the headend LSR#1 terminated the FlexE channels and using a transition map executes service or subservice signals onto specific SR labels and flows, which have their own quality of service COS assignments. The flow prioritization is kept throughout the rest of the SR domain. Note that the signals are taking a non-direct path after LSR#1, as would be determined by typical routing protocols. This is to show the deliberate nature of the SR policy.

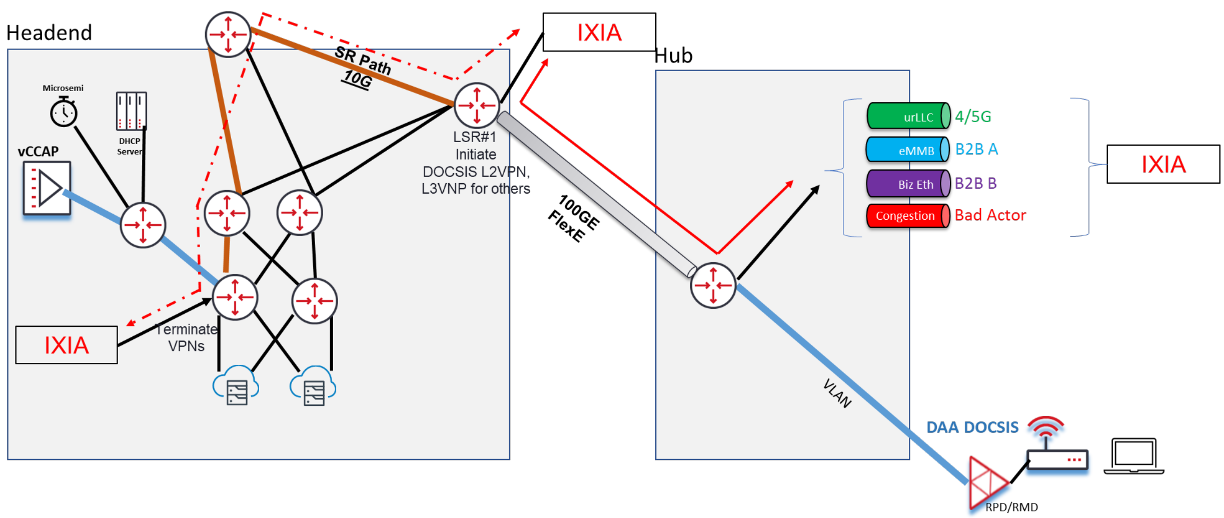

Figure 17, Actual Laboratory Setup

In Figure 17 we show the actual laboratory set up, which is a functionally networking equivalent to what is shown in Figure 16. Figure 17 shows the residential cable signaling is the only service that is running specifically on its native endpoints and cores. The system on the other hand mimics mobile and Business Ethernet flows via an Ixia traffic generator at the Hub. There is a 4/5G flow to mimic a signal from a wireless network source. The B2B-A and B2B-B signals represent enterprise grade circuits and the Ixia also provides a congestion flow to stress the FlexE channel—shown in the solid red line. In the Headend, at LSR#1 there is a breakout to another Ixia. This Ixia temporarily evaluates the FlexE signals and is the permanent stop for the congestion signal added to test the FlexE group. It is also the source for another “bad actor” signal into the SR domain in order to stress its class of service (CoS) mechanism—shown in the dot and slash red line. The LSR#1 is also where VPNs are originated, L2VNP for the vCCAP and RPD circuit, along with an L3VPN for the other signals. This is done to mimic typical business grade scenario in the case of L3VPNS, and in the case of DOCSIS signaling because we did not want to disturb same network settings for the vCCAP and the RPD. Note, that the FlexE segment relies on a 100G coherent direct connect signal, while the SR segment first path uses a 10G signal, for reasons that will be apparent in the next section. The VPNs are terminated ate the router before the last Ixia, through the representative cloud fabric.

The reason for differences between Figure 16 and Figure 17 was shortened set up and test time, a result of product transit, availability, and laboratory access in our current state with a pandemic.

9. Signal Settings and Test Results

	FlexE Settings	
Service Type	FlexE Group (Gbps)	FlexE Channel (Gbps)
B2B - A	100	5
B2B - B		5
4G/5G		5
DAA - DOCSIS		5
Bad Actor – FlexE*		75

Figure 18, FlexE Channel Settings

Figure 18 describes the FlexE bandwidth settings for each service type: Each service is assigned 5Gbps hard slices, and the "bad actor" congestion circuit is assigned 75Gbps, for a total composite of 95Gbps of client traffic and one backhaul signal and interface of 100Gbps.

Traffic Type	Tx (Mbps)	Rx (Mbps)
B2B - A	2G	2G
B2B - B	3G	3G
4G/5G	4G	4G
DAA - DOCSIS	25Mbps	25Mbps
Bad Actor	**90G**	**75G**

Figure 19, FlexE Throughput Results

Figure 19 shows the transmit and receive results for the Flex Ethernet group signal. We note three items. First when a bad actor that has a hard slice limit of 75G, as shown in Figure 18, tries to use more than its assigned SLA its signaling is curtailed back down to 75G, as determined by the FlexE channel setting, and second in this exercise of limiting the bad actor, none of the other service signals are affected. This is as expected. Third we note that the DAA DOCSIS signal is very low. This is per the one modem setting we had on the vCCAP. With more time we would have built a more extensive system. As it stand, with regards to bandwidth scale, the DDA DOCSIS signal is noise.

Service Type	DSCP	MPLS EXP	SR CIR	SR EIR (10G)
B2B - A	0 (Routine)	0	2G	8G
B2B - B	24 (Flash)	3	3G	7G
4G/5G	46 (Critical)	5	4G	0G (no burst)
DAA - DOCSIS	32 (Flash Override)	4	200Mbps	9.8G
Bad Actor – SR	1 (Priority)	1	200Mbps	9.8G

Figure 20, CoS Settings in SR Domain

Figure 20 shows the CoS settings for the service signals in the SR domain. The priority settings are expressed in the IP domain following the Differentiated Services Code Point (DSCP) structure. The MPSL priority settings are expressed in the MPLS Experimental bits mechanism (EXP), from lowest to highest. Basically, the MPLS priority structure can reflect the IP priorities to expedite any packet processing without digesting any IP headers. Note, the DCSP uses six bits and 64 possible values. The MPLS EXP on the other hand uses only three bits, and generally reflects the first three bits setting of the DCSP setting. Note, this allows for a further DSCP structure of priorities for subservices at the customer domain, to use as necessary.

Figure 20 also shows the bandwidth assignments for the different signals. The Committed Information Rate (CIR) shows the minimum bandwidth that the given signal will receive, this is typically a contractual SLA value. The Excess Information rate is the burstable rate possible, and it is the difference between the capable signal interface and the CIR. Note that the SR segment in egress from SLR#1 is 10G, thus the EIR values shows are calculated at 10G minus CIR. This is the nature of a flexible network slice as discussed earlier in the paper.

Traffic Type	Tx (Mbps)	Rx (Mbps)
B2B - A	2G	2G
B2B - B	6G	3G
4G/5G	4G	4G
DAA - DOCSIS	25Mbps	25Mbps
Bad Actor	10G	800Mbps

Figure 21, Transmission Results with SR CoS and FlexE

Figure 21 shows the end-to-end transmission results for the SR and FlexE segments. The focus here however is the priority structure set of the SR domain. Note the “Bad Actor” signal, which was assigned 200 Mbps, is trying to push 10G through the system which would take up all the line rate value. The output however shows the limiting nature of the soft slice as the system determined only 800 Mbps worth of capability. Also note the B2B-B signal, which has a limitation of 5G in the FlexE domain, and a limitation of 3G in the SR domain. When it tries to put through its 6G signal, it gets throttled to 3G. Also note, that the DOCSIS signal is unaffected. Ultimately the 10G SR segment is all used up based on the decision structure set by CoS.

Figure 22 below shows a snapshot of the raw data presented above. The one item to note in the tests was latency nature of the DOCSIS signal. We conjecture that this loss value and the high latency is due to the very small nature of the DOCIS signal in this case. We describe it as being in the noise and thus prone to calculation error, which is not the case of all the other signals which show miniscule latency.

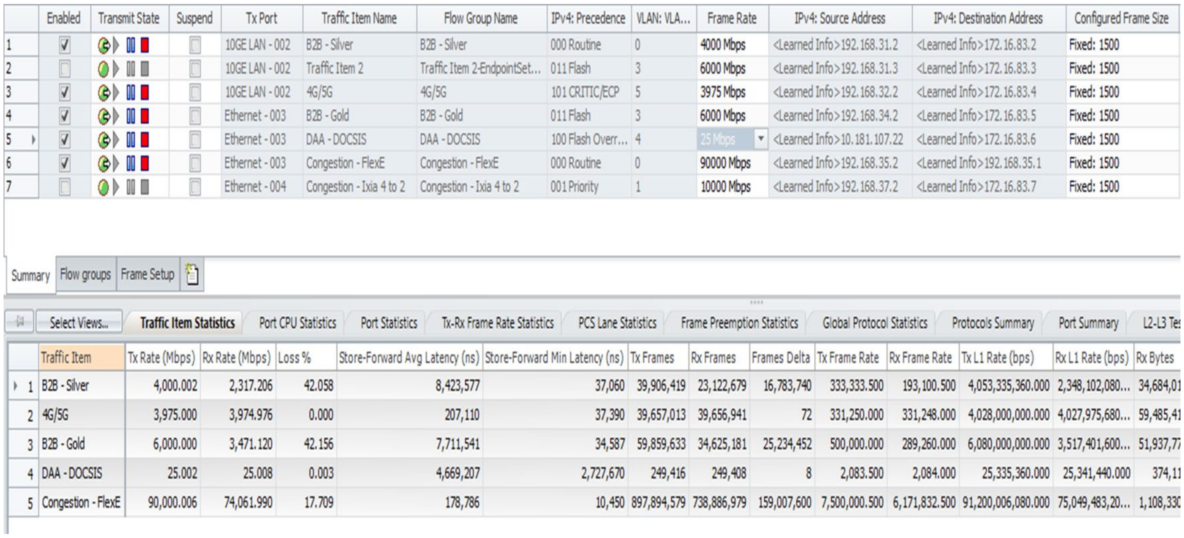

	Enabled	Transmit State	Suspend	Tx Port	Traffic Item Name	Flow Group Name	IPv4: Precedence	VLAN: VLA...	Frame Rate	IPv4: Source Address	IPv4: Destination Address	Configured Frame Size
1	✓			10GE LAN - 002	B2B - Silver	B2B - Silver	000 Routine	0	4000 Mbps	<Learned Info>192.168.31.2	<Learned Info>172.16.83.2	Fixed: 1500
2				10GE LAN - 002	Traffic Item 2	Traffic Item 2-EndpointSet...	011 Flash	3	6000 Mbps	<Learned Info>192.168.31.3	<Learned Info>172.16.83.3	Fixed: 1500
3	✓			10GE LAN - 002	4G/5G	4G/5G	101 CRITIC/ECP	5	3975 Mbps	<Learned Info>192.168.32.2	<Learned Info>172.16.83.4	Fixed: 1500
4	✓			Ethernet - 003	B2B - Gold	B2B - Gold	011 Flash	3	6000 Mbps	<Learned Info>192.168.34.2	<Learned Info>172.16.83.5	Fixed: 1500
5	✓			Ethernet - 003	DAA - DOCSIS	DAA - DOCSIS	100 Flash Overr...	4	25 Mbps	<Learned Info>10.181.107.22	<Learned Info>172.16.83.6	Fixed: 1500
6	✓			Ethernet - 003	Congestion - FlexE	Congestion - FlexE	000 Routine	0	90000 Mbps	<Learned Info>192.168.35.2	<Learned Info>192.168.35.1	Fixed: 1500
7				Ethernet - 004	Congestion - Ixia 4 to 2	Congestion - Ixia 4 to 2	001 Priority	1	10000 Mbps	<Learned Info>192.168.37.2	<Learned Info>172.16.83.7	Fixed: 1500

Summary | Flow groups | Frame Setup

Select Views... | Traffic Item Statistics | Port CPU Statistics | Port Statistics | Tx-Rx Frame Rate Statistics | PCS Lane Statistics | Frame Preemption Statistics | Global Protocol Statistics | Protocols Summary | Port Summary | L2-L3 Tes

	Traffic Item	Tx Rate (Mbps)	Rx Rate (Mbps)	Loss %	Store-Forward Avg Latency (ns)	Store-Forward Min Latency (ns)	Tx Frames	Rx Frames	Frames Delta	Tx Frame Rate	Rx Frame Rate	Tx L1 Rate (bps)	Rx L1 Rate (bps)	Rx Bytes
1	B2B - Silver	4,000.002	2,317.206	42.058	8,423,577	37,060	39,906,419	23,122,679	16,783,740	333,333.500	193,100.500	4,053,335,360.000	2,348,102,080...	34,684,01
2	4G/5G	3,975.000	3,974.976	0.000	207,110	37,390	39,657,013	39,656,941	72	331,250.000	331,248.000	4,028,000,000.000	4,027,975,680...	59,485,41
3	B2B - Gold	6,000.000	3,471.120	42.156	7,711,541	34,587	59,859,633	34,625,181	25,234,452	500,000.000	289,260.000	6,080,000,000.000	3,517,401,600...	51,937,77
4	DAA - DOCSIS	25.002	25.008	0.003	4,669,207	2,727,670	249,416	249,408	8	2,083.500	2,084.000	25,335,360.000	25,341,440.000	374,11
5	Congestion - FlexE	90,000.006	74,061.990	17.709	178,786	10,450	897,894,579	738,886,979	159,007,600	7,500,000.500	6,171,832.500	91,200,006,080.000	75,049,483,20...	1,108,330

Figure 22, Raw Data Sample From Ixia.

10. Conclusion

This paper motivates the principles for hard slicing in an Ethernet domain via Flexible Ethernet, and soft slicing in a Segment Routing domain via the CoS structure. The paper then successfully shows these principles exercised on a laboratory system with multiple types of services. The systems shown is highly representative of what MSOs are trying to achieve with one digitized access and metro network. Please contact the authors for more information and context.

The authors would like to thank Regis Mounkala, Kaylan Kanchumarthy and Tina Wu, Solution Engineers at Ciena for their lab support and dedication to excellence.

Abbreviations

3GPP	Service Layer Agreement
CIN	Converged Interconnect Network
CIR	Committed Information Rate
CoS	Class of Service
DSCP	Differentiated Services Code Point
EIR	Excess Information Rate

FlexE	Flexible Ethernet
GSMA	Groupe Speciale Mobile Association
IGP	Interior Gateway Protocol
IP	Internet Protocol
MAC	Media Access Controller
MII	Media Indipendent Interface
MPLS	Multiprotocol Label Switching
MPLS EXP	MPLS Experimental Bits
MSO	Multiple Service Operator
OIF	Optical Internetworking Forum
OTN	Optical Transport Network
PCE	Path Computation Engine
PCS	Physical Coding Sublayer
PHY	Physical Layer Implementation
PMA	Physical Medium Attachment (sublayer)
SLA	Service Layer Agreement
SR	Segment Routing
TI-LFA	Topology Independent Loop Free Alternate

Bibliography & References

BAUMGARTNER, J. (2019, April 16). *Comcast, Charter MVNO Deals Are Bad for Everyone – Analyst.* Retrieved from LightReading.com: https://www.lightreading.com/mobile/comcast-charter-mvno-deals-are-bad-for-everyone---analyst-/a/d-id/750837

Bonica, R. (2017, June 6). *A Segment Routing (SR)*. Retrieved from NANOG Archives: https://archive.nanog.org/sites/default/files/1_Bonica_Tutorial_Segment_Routing.pdf

Chamberlain, J. (2017). Competitive Advantages Of HFC Networks for. *SCTE Tech Expo 2017.* Denver: Society Of Cable Television Engineers.

Chamberlain, J. (2017, April 11). *Reality Check: Competitive advantages of HFC networks for wireless convergence.* Retrieved from RCR Wireless News: https://www.rcrwireless.com/20170411/opinion/reality-check-competitive-advantages-hfc-networks-wireless-convergence

GSMA. (2016). *Generic Network Slice Template Version 2.0.* GSM Association.

Optical Internet Working Forum. (22, June 2018). *Flex Ethernet 2.0.* Retrieved from oiffourm.com: https://www.oiforum.com/wp-content/uploads/2019/01/OIF-FLEXE-02.0-1.pdf

Villarruel, F. (2020). Framework For Convergence Of Services On The MSO Network Using The Principles of Network Slicing. *Cable Tech Expo.* Virtual: Society Of Cable Television Engineers.

Yeo, E. (2019). Segment Routing Proof of Concept for Business. *SCTE Cable Tech Expo.* New Orleans: Society of Cable Television Engineers.