



State of the Art Packet and
Optical Networking



Demystifying Network Slicing

Background Q&As

Network slicing is undoubtedly one of the hottest telecom buzzwords. Nevertheless, there is confusion over what it is, how it works, and whether it is really needed. We will tackle these questions head-on, using a Case Study approach. We will show how telcos can create differentiated and revenue-maximizing service offerings for both business and consumer segments, by implementing combined **soft and hard slicing**, which can be **dynamically controlled** and is **scalable**. Moreover, all the technologies that telcos need to implement this are available today.

To set the groundwork for the case study, we first go through a quick set of Q&As. ultra-low latency, ultra-high reliability, and highly-accurate timing required by such services.



Q: What Is Network Slicing?

It is the ability to support multiple subnetworks simultaneously, with different performance characteristics, on a common physical infrastructure. A network slice can be composed of both physical and network resources and is implemented at both the data and control planes.

Quoting the IETF directly: “A network slice is programmable and has the ability to expose its capabilities.” This sometimes-overlooked aspect is essential for dynamic control. It is useful to think about a network slice as exposing an API. If there is no API, then it is not a true network slice.

Q: What Problem Does Network Slicing Solve?

It enables efficient delivery of a broad mix of services with different performance characteristics and associated SLAs on a common network. Many foresee network slicing as necessary to support 5G’s multiple services classes, particularly, ultra-reliable low latency communications (URLLC). As such, network slicing is often thought of as purely a 5G-related technology. However, it does not need to be, and in the case study, we show how it can deliver value even for fixed access consumer services.

Q: Don’t We Already Have Network Slicing Today? (e.g. MPLS VPN)

Yes, but...

An MPLS VPN is indeed an example of a network slice and has shown its value in replacing private lines. However, it is very complex and expensive to implement, which is why it is limited to fixed-access business services. In addition, while MPLS VPNs are programmable in theory, their complexity has dictated that almost all implementations are static.

5G mobile services are intermittent or on-demand by definition. If we are to apply network slicing to 5G and consumer services, we will require a much more dynamic, efficient, and scalable approach.

Q: Why Haven't We Implemented Dynamic and Scalable Network Slicing Yet?

It is essentially a chicken and egg problem. Stated another way, it is an “If you build it, they will come” situation. No one has built it yet, even though all of the technologies already exist.

When Apple created the iPhone, they did not need to develop any new touch screen, microprocessor, or any other technology. Their innovation was in integrating existing technologies into a package that they thought customers wanted and were willing to pay for, without customers actually asking for the new iPhone. At the same time, other companies focused on legacy phones with cost in mind, while conducting market research on what customers said they wanted, which of course was just variations of familiar experiences. It can be a similar situation for dynamic and scalable network slicing. The data plane technologies already exist in packet routers and switches, as well as optical switches. So do the SDN and control plane technologies. There must be a more profound reason why telcos are not pursuing network slicing as a means to expand service differentiation. The answer lies in economics and the inertia of existing business models. In order to maximize profitability, service delivery businesses usually develop schemes that create differentiated service offerings – at different price levels – based on differentiated demand for factors, like speed or quality of experience. Examples include package shipping, express lanes on a highway, or seats on an airplane.

Aside from expensive fixed-line business offerings, like private lines or MPLS VPNs, telcos have been reluctant to explore this model further. As a result, they are inevitably losing business to over-the-top innovators, who are not only using the power and ubiquity of Internet-based (best-effort) services, but are also starting to chip away at differentiated connectivity services. For example, AWS (Direct Connect) and Azure (ExpressRoute) already offer direct connectivity to their sites, thus bypassing public Internet access with its unguaranteed performance. Even more, AWS offers a new platform and service called Wavelength that aims at delivering ultra-low latency applications from AWS-based network edges to mobile devices and end-users, positioning themselves as 5G communications suppliers. This is truly an absurd situation. To deliver Wavelength, AWS is creating a network-edge computing overlay on top of existing telco networks, which telcos are in a much better position to deliver in the first place. (For further reading, see the recent LightReading blog, “Telcos Let FAANGs into the Edge at their Peril.”)

5G may very well be the last opportunity for telcos to change their business model, from being just a supplier of raw bandwidth, to being a supplier of **differentiated connectivity services**. To do this, they need to be able to deliver “MPLS VPN-like services” to businesses and consumers, but dynamically, on a pay-per-use basis, and much more cheaply. For such challenges, they must implement a combination of soft and hard network slicing.

Q: What Is the Difference Between Soft and Hard Slicing?

The following brief answer will become clearer during the Case Study discussion.

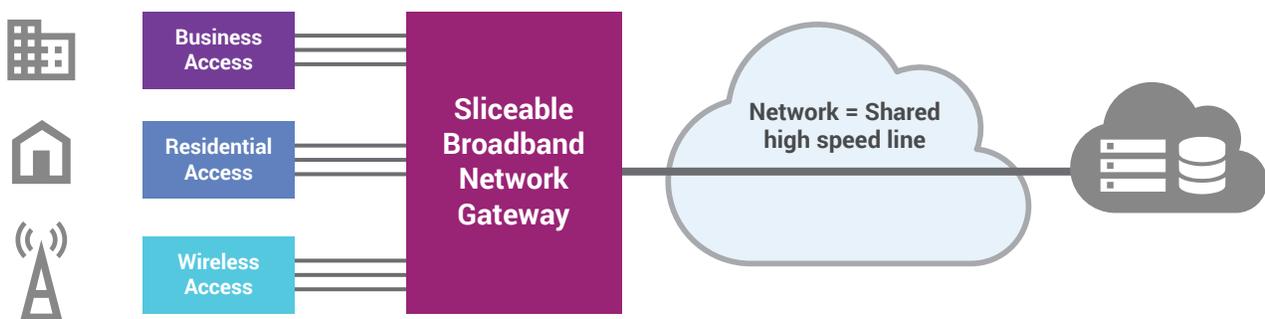
Soft slices provision resources in such a way that, while the services they carry do not, on average, interfere with each other (and one service cannot receive another's packets), they usually compete for resources, such as their position in a buffer queue, or CPU cycles. As a result, services running over soft slices can only be engineered to have an average level of performance. A relatively higher or lower level of performance is achieved by adjusting the degree of "contention" allowed for these resources. However, performance can never be guaranteed absolutely. Just about all forms of packet-processing-based services are examples of soft slicing implementations, including MPLS VPNs.

Hard slices provision resources in such a way that the services that they carry are fully isolated from other services, short of network failures. As a result, services running over hard slices can be engineered to have an absolute or guaranteed level of performance. Examples of resources used to build hard slices are TDM time slots (time isolation) and WDM optical channels (frequency isolation).

Case Study Setup

While 5G is mostly associated with network slicing, to better understand the concept without delving into the complex 5G radio side, we chose to use a non-existent, but not so hypothetical, sliceable broadband network gateway (S-BNG) as an example. The principles discussed apply readily to the more understandable backhaul side of 5G.

The S-BNG must allocate its resources to serve a range of fixed and wireless access services, with different performance expectations and SLA commitments, over a common physical network. The network is simply a high-speed line, which could be 100G or 400G, for example.



While this S-BNG does not exist today, it is fully implementable with off-the-shelf technologies. Initially, we create it here as a thought experiment to describe the possibilities of network slicing. It could, however, be implemented with just a forklift upgrade to the existing network. Doing so is a genuine opportunity for 5G and its wireless access services.

The S-BNG must serve a customer and service mix with target performances, as described below. For our discussion, it is more important to consider the various specifications of Ethernet rates, BW commitment, and latency, in relation to each other. The absolute values shown are only illustrative at this point.

	ID	Service name	Ethernet rate	BW commitment	Latency commitment
Business Access	B1	Business standard	Very fast (e.g. N x 1GE)	<10:1 Contention Ratio, best effort	No spec
	B2	Business prime	Very fast (e.g. N x 1GE)	80% CIR	Low (e.g. 50ms 80% time)
	B3	Dedicated channel	Very fast (e.g. N x 1GE)	Guaranteed	Very low (e.g. 10ms 100% time)
Residential Access	R1	Residential basic	Moderate plus (e.g. 25ME)	<20:1 CR, best effort	No spec
	R2	Residential plus	Moderate plus (e.g. 50ME)	<20:1 CR, best effort	No spec
	R3	Residential gamer	Moderate plus (e.g. 50ME)	80% CIR	Low (e.g. 50ms 80% time)
Wireless Access	W1	Wireless basic	Slow (e.g. 10ME)	<30:1 CR, best effort	No spec
	W2	Wireless plus (5G)	Moderate plus (e.g. 50ME)	<30:1 CR, best effort	No spec
	W3	Wireless gamer	Moderate plus (e.g. 50ME)	80% CIR	Low (e.g. 50ms 80% time)

Assessing Network Slicing Combinations

In the remainder of this white paper, we examine the pros and cons and four setups of the BNG that use different combinations of soft and hard slicing. The assessments cover:

- Effectiveness** – the ability to support tiered services and to guarantee SLAs, so that total revenue-per-bit is maximized.
- Complexity** – the extent to which queuing and scheduling mechanisms are employed.
- Utilization** – the relative use of the BNG’s theoretical maximum frame processing capacity.

While this is not an exhaustive analysis, it illustrates how using various packet queuing and TDM scheduling mechanisms produces different results. It indicates using a combined soft and hard slicing as the only way to support a mix of best-effort and committed services, with a range of statistical and deterministic performance guarantees. The table below summarizes the setups.

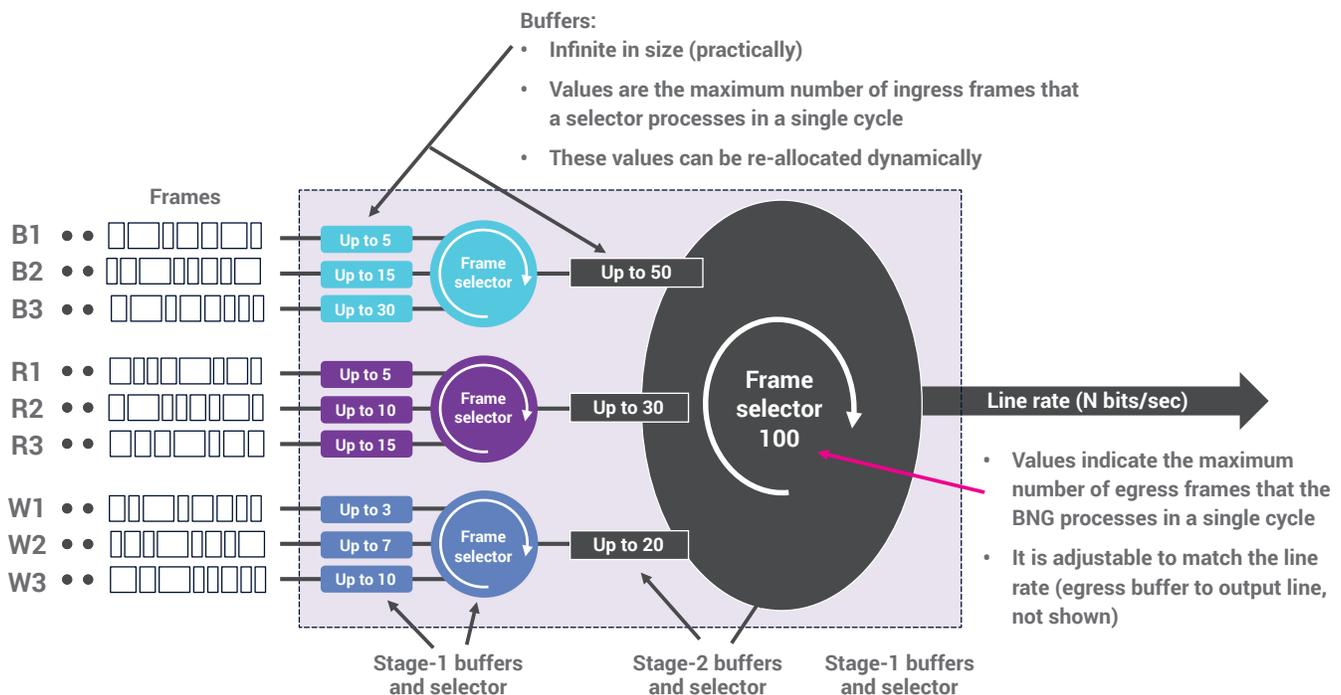
Setup	Description	Relative Comparison of BNG Setups		
		Effectiveness at meeting service commitments	Implementation complexity	Frame processing capacity utilization
Soft slicing, best effort only	Uses WFQ exclusively for all services. By playing with weightings, tries to meet all expectations as best as possible.	LOW Cannot deliver any bandwidth or latency guarantees	LOW Uniform implementation throughout	HIGH Maximum use
Soft slicing, with queue prioritization	Supplements WFQ above with AF and EF queue prioritization to improve performance for selected services.	MEDIUM Supports statistical bandwidth guarantees, but cannot guarantee specific latency or jitter	MEDIUM Uses a mix of queue prioritization mechanisms	MEDIUM Some unused slots due to prioritization schemes
Combined soft and hard slicing	Adds TDM hard slicing capabilities to the previous setup for services requiring hard bandwidth and latency guarantees.	HIGH Supports a mix of best effort and committed services, with a range of statistical and deterministic performance guarantees	HIGH Combines queue prioritization with TDM techniques	MEDIUM Some unused slots due to prioritization and TDM schemes
Hard slicing only	Uses TDM hard slicing exclusively for all services.	HIGH Supports deterministic performance guarantees for all services. While overkill for non-committed traffic, makes sense when there are many QoS and SLA levels requiring deterministic latency and jitter, perhaps in a future AR/VR dominated services world.	LOW Uniform implementation throughout	LOW Most unused slots due to use of TDM throughout

Acronyms: WFQ = Weighted Fair Queuing; BE = Best Effort; AF = Assured Forwarding; EF = Expedited Forwarding; TDM = Time Division Multiplexing

Note that actual queue management includes more mechanisms, such as policers, shapers, buffer managers (WRED, RED), etc. In addition, some implementations use token-based assigned BW-per-queue and not simple frame counters. Here, for the sake of simplicity and clarity, we will discuss the principles as they are reflected in the main queuing mechanism.

Setup #1 – Soft Slicing - Best Effort Only

In all setups, the BNG uses two stages. In the first stage, each of the services (business, residential, and wireless) accesses its own individual selector. Then, the BNG combines the traffic at a second stage selector. The first basic setup of the BNG uses a weighted fair queuing (WFQ) scheduling algorithm at both stages.



Starting at the bottom of the second stage, the selector processes 'up to' 20 frames coming from the wireless access services. When 20 frames are processed, or when there are no more frames to process, the stage-2 selector processes 'up to' 30 frames in the residential services buffer, followed by 'up to' 50 frames in the business services buffer, and then the cycle repeats.

The stage-1 selectors follow the same process for the different streams of services. Each service has its own buffer and the stage-1 selectors draw from those buffers to feed the stage-2 buffers.

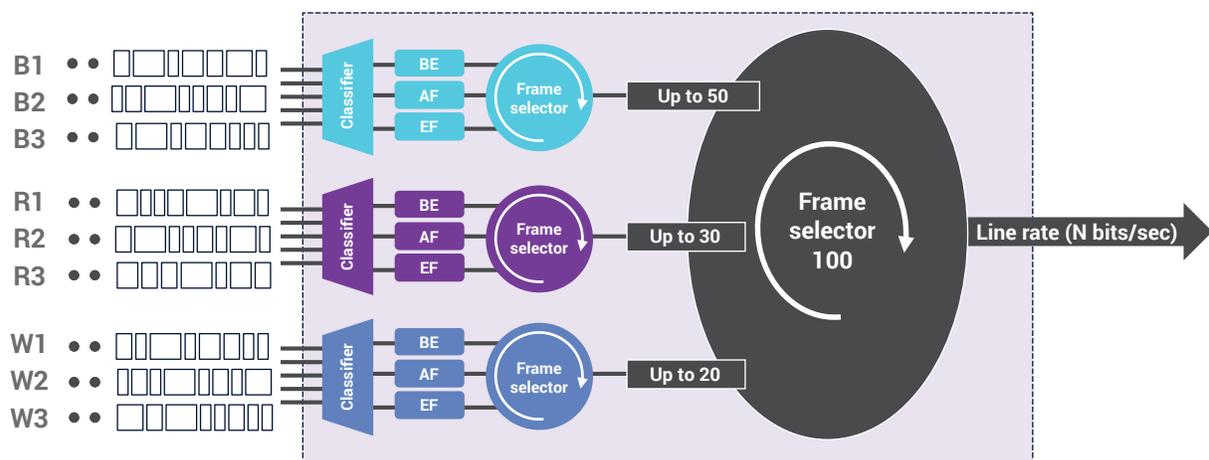
By adjusting the size of the stage-1 and stage-2 weights, it is possible to change the relative performance for each service category. As shown, the business services receive - on average - 50% of the available bandwidth, residential services 30%, and wireless access services 20%. Individual services are prioritized similarly within each category.

This setup has low complexity and makes the most efficient use of the BNG's processing capacity, in that it can handle the highest level of traffic demand compared to the other setups. Its big limitation is that, other than applying "weighting", it processes all services equally. In addition, it cannot guarantee the bandwidth or latency performance of any particular service, such as B2, B3, R3, and W3 services.

Setup #2 – Soft Slicing - Best Effort with Queue Prioritization

We can improve on the above model by introducing strict priority queuing at the first stage. Specific queues have precedence over others and get their frames processed before other queues, regardless of the round-robin schedule. (It is possible to extend priority queuing to the second stage of the BNG, but for simplicity, we follow the WFQ model above.) We show three queue types:

- **Expedited Forwarding (EF)** – Supports services with low delay, low loss, and low jitter, such as voice and interactive video. In the defined services, these include B3, and probably R3 and W3. Frames in this buffer queue are almost always selected with strict priority over other traffic classes.
- **Assured Forwarding (AF)** – Supports services with guaranteed bandwidth, so long as the traffic does not exceed a subscribed rate. This is achieved by engineering the AF traffic, such that in the worst case, the drop probability of AF frames in any queue in the network is very low. Frames in this buffer queue are next in priority to EF, and are processed, only if there are no pending EF frames. In the defined services, AF priority would be applied to B2, and possibly to R2 and W2.
- **Best Effort (BE)** – Supports best-effort services. The selector only processes frames in this buffer queue when the EF queues are empty. This class is, by definition, overprovisioned to a predetermined contention level, so that frame drops and fluctuations of available BW are expected. This is not problematic, so long as the bulk of traffic transported is at best effort.



Some may consider the addition of queue prioritization as good enough for high-level QoS guarantees, but this approach has several problems:

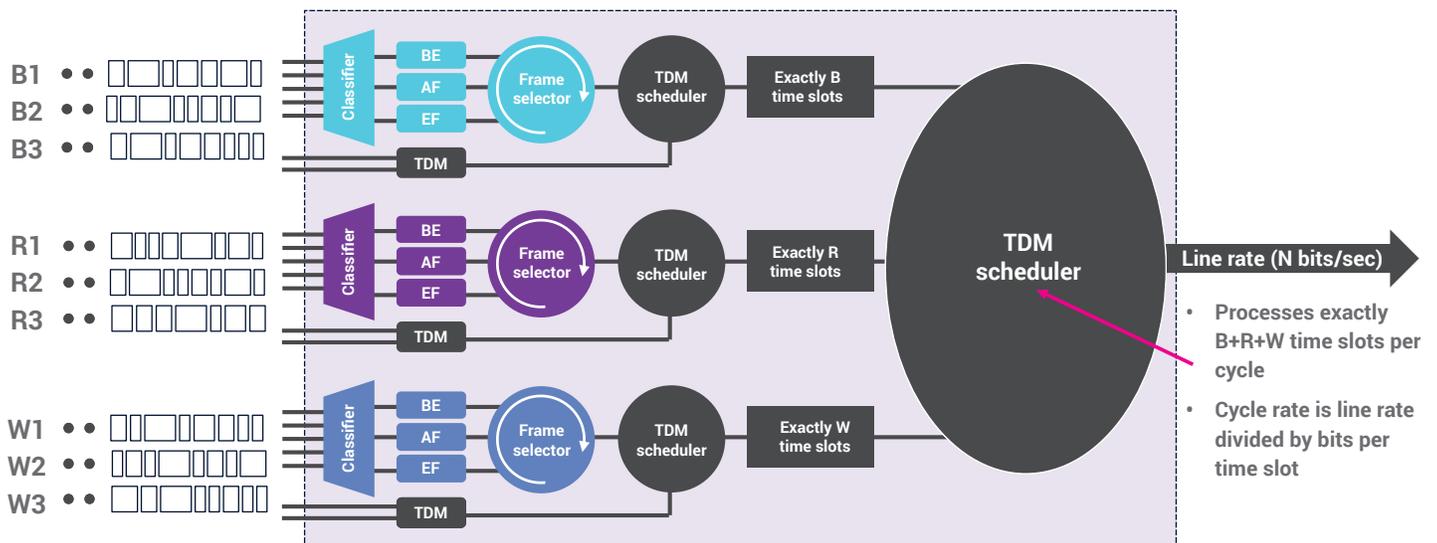
- There is still no full isolation between flows. EF traffic would have to wait an unspecified time until the current frame (which can be anywhere from 64B to 1500B, or even more) is delivered, which leads to jitter, even on the EF traffic.
- Strict priority behavior is very aggressive and can easily lead to starvation for the non-EF traffic, even if the total traffic is less than 100% of link capacity. This limits the portion of EF traffic to 5-10% (or less) of the total traffic.

Setup #3 – Combined Soft and Hard Slicing

To deliver both high utilization and hard QoS guarantees, we must add hard slicing. This affects both stages of the BNG. At stage-2, the frame selector becomes a TDM scheduler, and the algorithm for the stage-2 buffers goes from processing “up to” a given number of frames, to processing “exactly” a given number of time slots for each service category.

For example, if the BNG is processing the wireless access buffer, and after processing (let’s say) 12 frames, there are no more frames to process, it still waits for the remainder of the allocated time slots, and will process any new frames appearing within that time. (The total cycle rate for the stage-2 scheduler is determined by dividing the egress port line rate by the slot width. For example, if the line rate is 400Gbps and the slot width is 50 bits, then the cycle rate is 8G slots per second.)

The effect of this change is that it dedicates a percentage of the BNG’s overall egress bandwidth to each service category. While this reduces the overall efficiency of the BNG (because in some cases, the BNG is not transmitting anything while waiting for frames on an empty queue), it does provide a method to better achieve the offered performance for each service category.

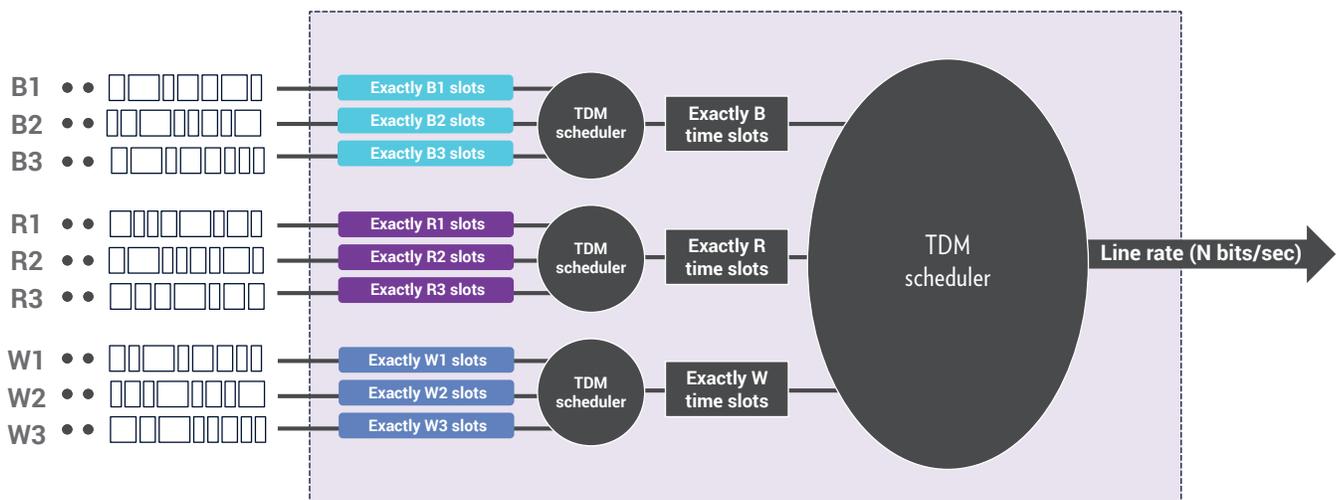


To take this a step further and introduce the means to guarantee the performance for any particular service, we also add optional hard slicing to the first stage of the BNG. We accomplish this by adding a TDM scheduler in series with the soft slicing frame selector. This dictates that all traffic coming from stage-1 is now mapped into fixed time slots. Some time slots are reserved for soft-slicing traffic coming from the EF/AF/BE queues, and some slots are reserved for specific services. For example, we would likely assign dedicated time slots to service B3 that requires both hard bandwidth and latency guarantees. Adding TDM hard slicing to stage-1 similarly reduces the overall efficiency of the BNG, because in some cases, it now processes empty slots.

Setup 3 is the only setup we have analyzed so far that can process both statistical traffic as well as traffic with hard QoS guarantees. While this does come at some loss of efficiency, the result can be compared to airlines not moving passengers from coach to business class when the latter is not full and they can still sell more coach seats. In the case of the BNG, the situation is even better, because it can allocate resources and assign priorities dynamically between the hard and soft slices, and within the soft slices. Through proper market research and service differentiation, BNG setup 3 provides a powerful tool to maximize revenues, profits, and customer satisfaction. Moreover, because it can be constructed using a building block approach, it is inherently highly scalable.

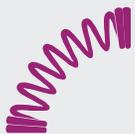
Setup #4 – Hard Slicing Only

This setup is shown primarily for completeness. It is analogous to setup 1, but instead of each queue being best effort, it is now exact. In effect, the BNG becomes a two-stage OTN switch. While it has the capability to provide every service with performance guarantees, this would only make sense when there are many services with SLA levels requiring deterministic latency and jitter. Perhaps this may be the case in the future, if we move into a world where the majority of services deliver real-time VR/AR experiences. However, for today's world, it is an overkill approach, when it is perfectly acceptable to offer a majority of services with statistical multiplexing that combines best effort with committed traffic.



TDM Re-Emerging

For the last few decades, TDM technologies, such as voice-centric SDH/SONET, have steadily been crushed by the weight of a data-centric statistical packet network. They have been assigned to the bin of legacy technologies. Yet, like buds appearing after a desert storm, we are now seeing a new generation of TDM technologies emerging to work in concert with packet networks. These will facilitate creating hard-plus-soft slicing networking solutions for dealing with committed and low-latency services traffic.



Flexible Ethernet (FlexE) is an OIF standard that enables multiplexing Ethernet connections to an optical network in ways that are more flexible than just one-to-one mapping of speeds. At its heart is a timeslot calendar mechanism that acts as a scheduler to multiplex different Ethernet clients, at Physical Coding Sublayer (PCS) symbol level, over any number of optical channels. By allocating predefined timeslots to each Ethernet client flow, it achieves hard slicing and isolation between the different flows.



Time Sensitive Networking (TSN) is a set of IEEE Ethernet standards to enable deterministic real-time communication over Ethernet. TSN achieves determinism by using frame level time synchronization and a scheduler that is shared between network components. By defining queues based on time slots, TSN ensures a bounded maximum latency for scheduled traffic through switched networks, guaranteeing the latency of scheduled communications.

Discussion and Conclusion

Over the past fifty years, the telecommunications network pendulum has swung from TDM-centric voice to packet-centric data. These incorporate special packet processing mechanisms, such as prioritized queues, to improve performance statistically for voice, interactive video, and other services that are sensitive to latency and jitter. However, these mechanisms can only create soft network slices and are unable to deliver hard latency or jitter SLA guarantees. The up-and-coming 5G world, where URLLC services are a defined service mode, will exacerbate the situation of having “not quite good enough” solutions for latency- and jitter-sensitive services.

In anticipation of finding solutions for this growing challenge, new “hard slicing” technologies, like FlexE and TSN are emerging that can deliver deterministic performance. This paper shows how we can combine soft and hard slicing mechanisms to deliver a range of statistical and deterministic performance guarantees for a broad variety of services with different bandwidth and latency characteristics.

All the technological elements to implement this approach are available today. Nevertheless, telcos have not pursued it, mostly because they have not seen the additional revenue benefits as sufficient to make up for the added implementation complexity.

However, the situation is worth reconsidering, based on the additional revenue potential of 5G services. A dynamically configurable soft-plus-hard sliced network gives telcos a “holy grail” for offering customers a suite of tiered services with different price points and performance expectations. The immediate benefits (even before 5G) are higher revenues, higher profitability, and greater customer choice and satisfaction. Even more, it gives telcos an ability to claw back market share from OTTs, who have already started playing this tiered communication services game at the expense of telco-provided raw bandwidth.

All it takes is the resolve to get started.

Contact us for more information about how Ribbon can help you meet and overcome the 5G network connectivity challenge

About Ribbon

Ribbon Communications (Nasdaq: RBBN), which recently merged with ECI Telecom Group, delivers global communications software and network solutions to service providers, enterprises and critical infrastructure sectors. We engage deeply with our customers, helping them modernize their networks for improved competitive positioning and business outcomes in today's smart, always-on and data-hungry world. Our innovative, end-to-end solutions portfolio delivers unparalleled scale, performance, and agility, including core to edge IP solutions, UCaaS/CPaaS cloud offers, leading-edge software security and analytics tools, as well as packet and optical networking leveraging ECI's Elastic Network technology.