State of the Art Packet and
Optical Networking

ribbon®

5G – A Fork in the Road
for Network Transport

5G is Coming – Place Your Bets

Today's network provides one predominant service: all-purpose IP connectivity. End users establish value-added services – Internet searches, Skype meetings, movie downloads, use of AWS or Azure compute resources, etc. – over pre-provisioned and long-duration IP connections, without involving the network operator. These connections are either best-effort Broadband, or VPNs, featuring SLAs for bandwidth and availability.
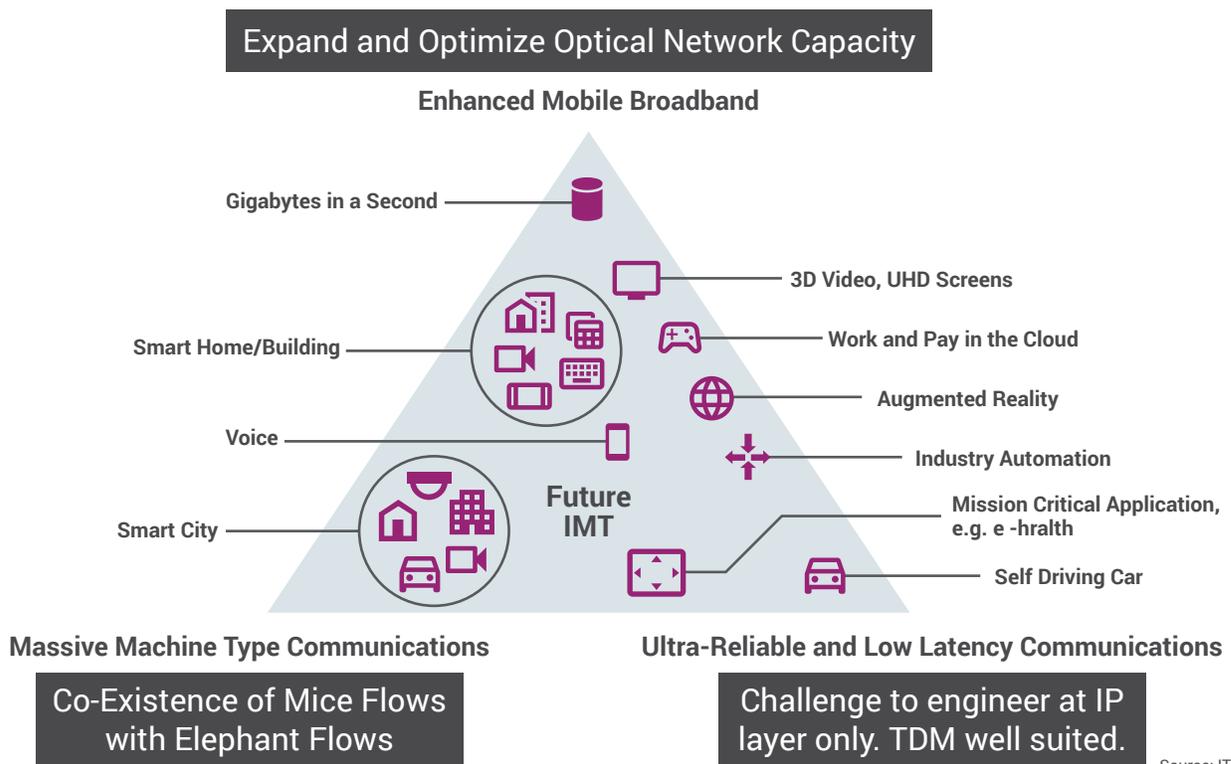
5G has the potential to change this picture dramatically. By definition, 5G technology supports three major service classes: enhanced high-speed mobile and fixed broadband, ultra-reliable and low latency communications, and massive machine-type communications. Combinations of these are possible, as well as the ability to set up short-duration connections dynamically.

The question facing service providers is what strategy they should pursue to evolve their network infrastructure to support 5G services. The answer, it seems, depends on two SP beliefs, how quickly will 5G achieve its full service potential, and whether or not all traffic-engineering challenges are solvable at the IP layer. Based on this, we see three distinct SP strategies emerging for 5G network transport:

1. Business as Usual	**2.** IP-over-DWDM	**3.** Multilayer Optimization

Expand and Optimize Optical Network Capacity

Enhanced Mobile Broadband

Gigabytes in a Second

Smart Home/Building

Voice

Smart City

3D Video, UHD Screens

Work and Pay in the Cloud

Augmented Reality

Industry Automation

Mission Critical Application, e.g. e -hralth

Self Driving Car

Future IMT

Massive Machine Type Communications

Ultra-Reliable and Low Latency Communications

Co-Existence of Mice Flows with Elephant Flows

Challenge to engineer at IP layer only. TDM well suited.

Source: ITU

"When you come to a fork in the road, take it!" is a famous quote from former New York Yankees manager Yogi Bera. Among the many decisions 5G presents to service providers, this white paper examines three prongs of a fork in the road SPs can pursue for network transport.

ribbon

Strategy 1 – Business as Usual

In this scenario, the SP believes that the only 5G service to be deployed in the foreseeable future is enhanced bandwidth. Other 5G services may come at some point, but this class of SP is not basing their network strategy on that occurring anytime soon.

The resulting strategy is simply to extend the network's current IP-connectivity service to support higher bandwidths. This means adding IP routing and switching capacity, as well as underlying optical transport capacity. This is somewhat of a gloomy situation because even in the current pre-5G environment the SP is squeezed to make money, and the increase in fees they can feasibly obtain for 5G enhanced bandwidth may not cover the cost of expanding network capacity.
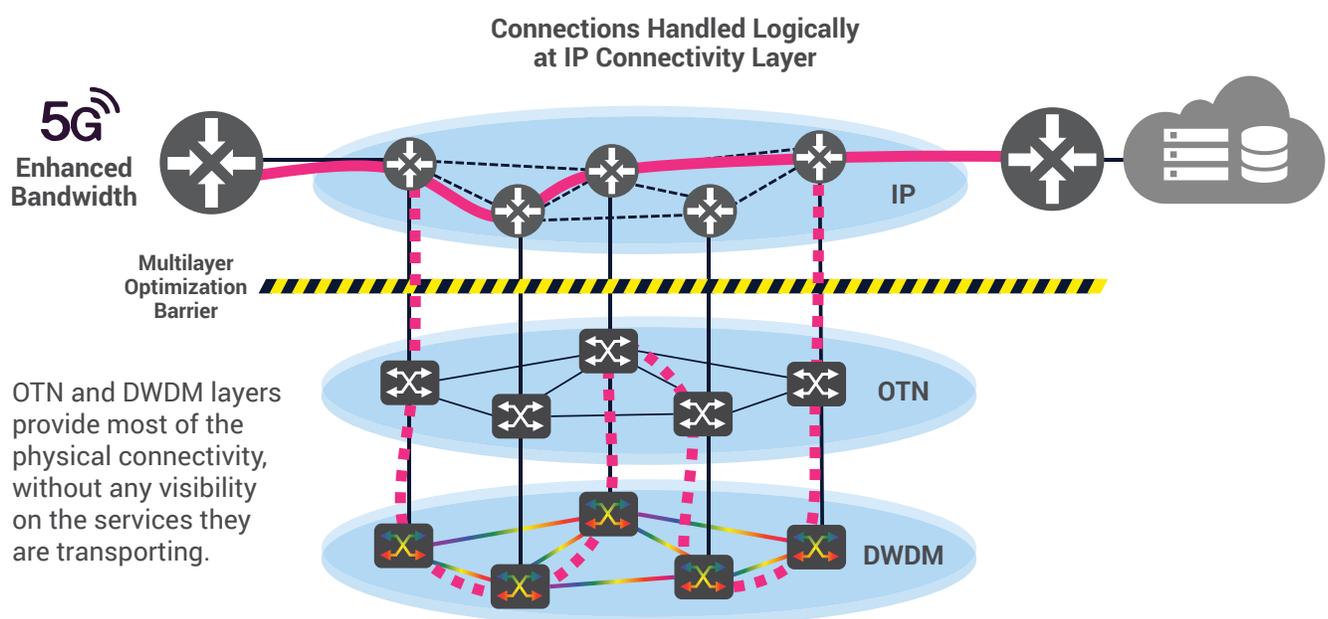
FOMO, or fear of missing out, is a primary reason that SPs in this category pursue this path. They will agree that the treadmill of delivering ever-increasing bandwidth at a more or less constant price is tough, but they have managed until now and no doubt will continue to do so in a higher bandwidth 5G era. They do not see an alternative. If they don't provide the higher bandwidth someone else will do so and they are out of the game. (See the ECI blog, The Great Telco-OTT Symbiosis, which explores this dilemma.) SPs in this category also argue that there may be killer aps waiting, exploiting VR/AR for example, based on which they can raise prices for 5G high bandwidth more significantly.

👍 PROs	👎 CONs
• Maintains the status quo. Besides investing in 5G radio, any additional capital investment for network infrastructure is limited to expanding bandwidth capacity, which is a well-understood process. • Takes a 'wait and see' attitude for other 5G services (beyond high bandwidth) before committing to more significant network modernization investments.	• The service and optical layers are expanded independently, without any multilayer optimization of resources • Loses time against competitors if other 5G services like low latency and IoT take off more quickly than predicted. If this occurs, it may be difficult for the SP to catch up.

Connections Handled Logically at IP Connectivity Layer

5G Enhanced Bandwidth

Multilayer Optimization Barrier

IP

OTN

DWDM

OTN and DWDM layers provide most of the physical connectivity, without any visibility on the services they are transporting.

ribbon

Strategy 2 – IP-over-DWDM

As in the previous scenario, the SP believes that the initial 5G service will be enhanced bandwidth. However, the SP also foresees that other services like low latency and IoT will follow in a few years, and that they need to start making changes to their network infrastructure to prepare for that.

SP's is this category, see this as an ideal opportunity to begin moving to a more streamlined IP-over-DWDM network, with the following general characteristics:
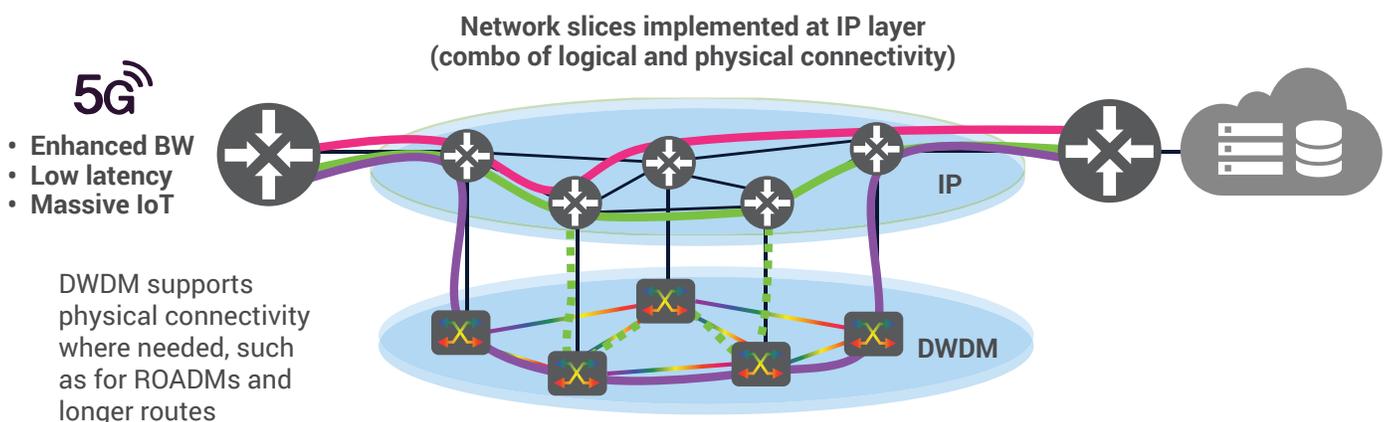
- IP-connectivity is the only network service, as opposed to the predominant one. The IP layer handles all service functionality.

- TDM capabilities like OTN switching are no longer used. This eliminates what is seen as an unnecessary network layer and simplifies network engineering and operations. Circuit emulation handles any legacy TDM services.

- Optical transceivers integrate directly into IP routers and switches for spans up to about 100km. New 400GZR/ZR+ technologies make this particularly feasible.

- Longer spans, and network sections using ROADMs, will continue to rely on a distinct DWDM layer using higher power transponders. Purists may see this as not adhering to an ideal IP-over-DWDM vision, although it is likely necessary.

- Network slicing is introduced to handle different 5G service categories. All slicing mechanisms occur exclusively at the IP layer.

| 👍 PROs	👎 CONs
- Takes a proactive approach to network modernization that recognizes that essentially only a single IP connectivity service exists, today. - Reduces Capex and Opex by creating a streamlined network that eliminates all TDM capabilities, and brings the IP service and DWDM transport layers closer to each other. - In theory, mechanisms at the IP layer handle all 5G service variants, enabling fast deployment and added revenues as these emerge.	- In practice, it is technically very challenging to engineer deterministic services like low latency plus high availability at the IP layer. This gap between theory and practice can lead to ineffective support for deterministic 5G services as they emerge. - Packet is based on a statistical multiplexing paradigm. It is very efficient when there are no QoS guarantees or when they are soft (e.g. average committed BW) yet very inefficient when hard QoS is needed as this negates the benefits of statistical multiplexing

5G
- Enhanced BW
- Low latency
- Massive IoT

DWDM supports physical connectivity where needed, such as for ROADMs and longer routes

Network slices implemented at IP layer (combo of logical and physical connectivity)

IP

DWDM

ribbon

One way to support the different 5G services is to create a separate network for each type. One network for high bandwidth, one for low latency with high availability, one for IoT, and so on. While separate networks clearly carries a prohibitive price tag, this approach points us in the direction the industry is heading; namely, network slicing. In effect, this carves out multiple sub-networks with different performance characteristics, from a common physical infrastructure.
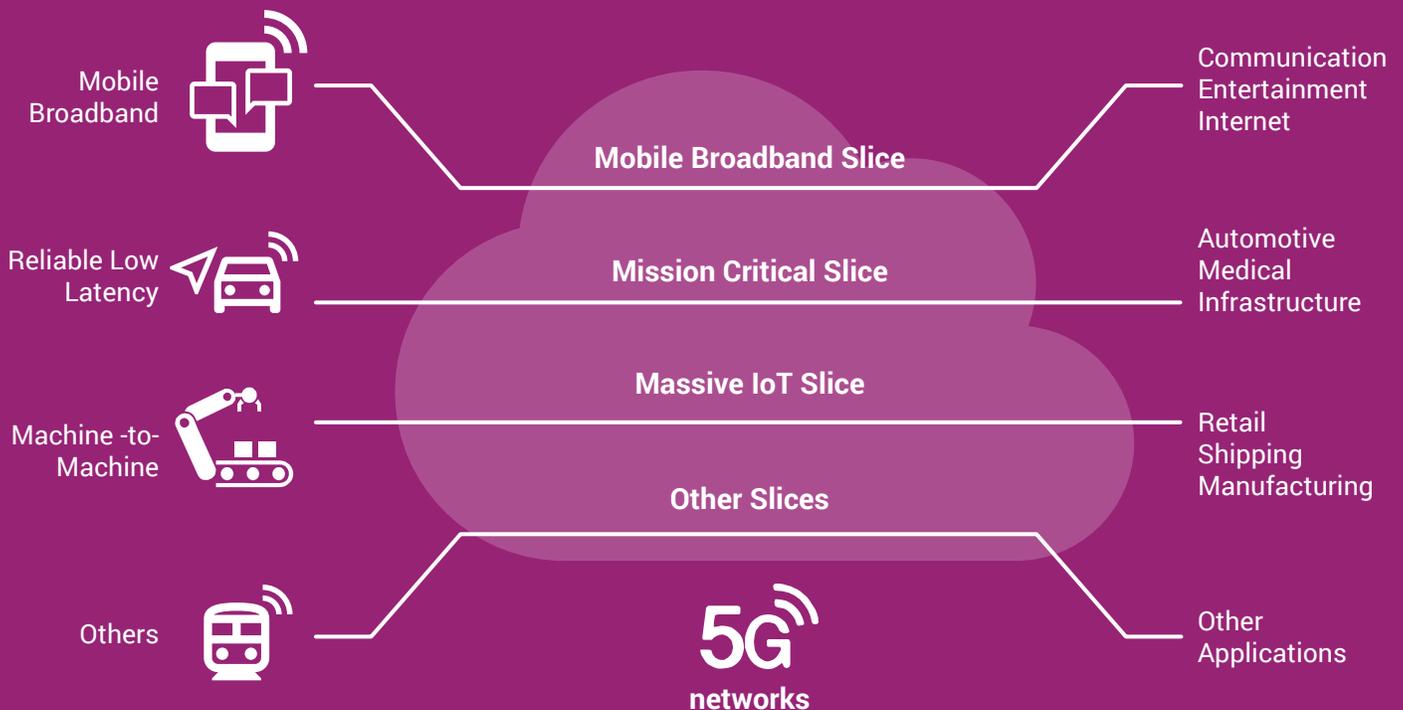
Slices can be composed of 'soft slices' that prioritize resources inherently shared, such as packet processing and virtualized computing, or hard 'hard slices' that dedicate fixed resources, such as router ports and optical channels. As service needs or traffic conditions change, automated software control can shift the resources assigned to the network slices dynamically. At a minimum, 5G services network slices will require at least two main layers to build upon:

- A broadband layer for general-purpose traffic without any specific SLA requirements. This is analogous to the current IP layer.

- A low latency high availability layer, whose implementation approach is the subject of debate. The IP-over-DWDM camp believes this can be engineered adequately solely at the IP layer, while the multilayer optimization camp believes that at least some assistance is required from TDM technologies at the optical transport layer.

Strategy 3 – Multilayer Optimization

As in the previous scenario, the SP envisions that the initial 5G service will be enhanced bandwidth, with low latency and IoT services following in a few years. The major difference is that this class of SP does not have confidence that deterministic 5G services can be supported effectively only at the IP layer.

Their approach is to maintain separate IP-connectivity and optical transport layers, but with one major difference. They will optimize and manage these layers as an integrated whole, rather than independently as is done today. This accomplishes three things. First, it dramatically reduces the capital and operating cost structure of the network by removing redundant resource utilizations, so that the network and thus the SP is more financially competitive. Second, it leaves in place TDM mechanisms such as OTN switching for creating network slices to handle deterministic 5G services.

The third benefit is the furthest reaching. Coordinating and optimizing multilayer resources is undeniably difficult. This will require a powerful set of network flexibility mechanisms and control algorithms. (A subject discussed further in the next section.) An upside, however, is that once these are in place, they can be used to offer a complex services suite featuring different levels of bandwidth, latency, availability, and duration. This enables maximizing revenues, in a similar way that cloud service providers to for compute services, by creating stratified service tiers within each 5G service category, particularly for bandwidth and availability guarantees.

👍 PROs	👎 CONs
• Makes most efficient use of Capex by optimizing overlapping capabilities among network layers.	• Technically complex. Requires new telemetry reporting and real time control algorithms.
• Streamlines network operations over an integrated network, reducing Opex.	• Minimizes redundant backup mechanisms, so can lose network robustness if not done correctly.
• Contains IP and TDM mechanisms to handle all types of 5G services effectively.	• Requires merger of service and optical transport organizations, or close integration under a common master.
• Maximizes revenues via stratified services.	

ribbon
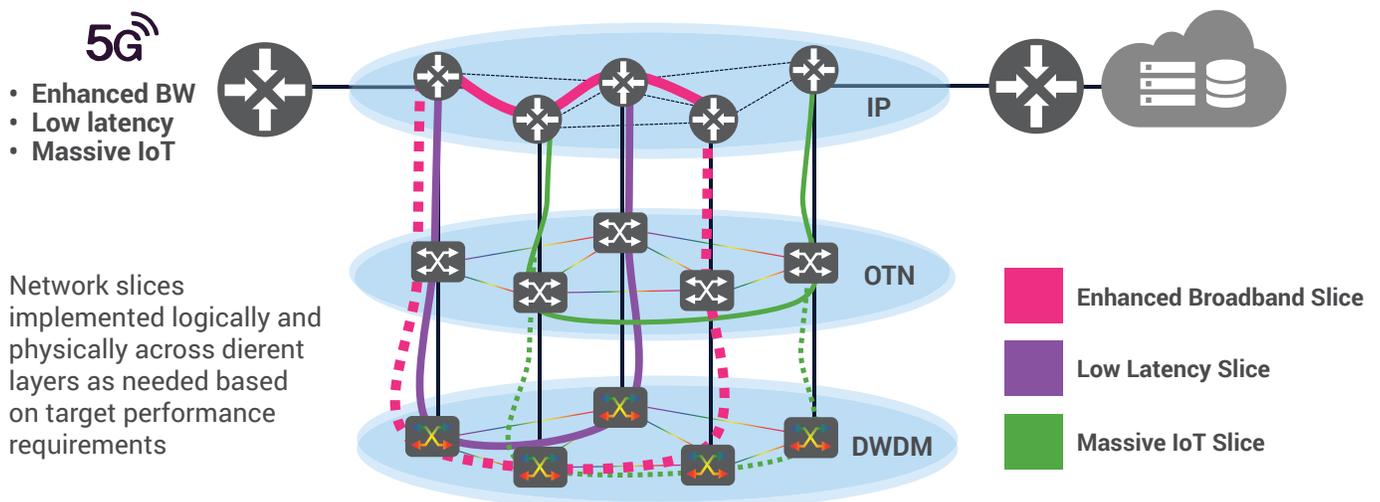
A Deeper Dive into Multilayer Optimization

Of the three strategies listed, multilayer optimization is perhaps the most enigmatic. It is not a new concept, and may indeed be the best approach that combines economic efficiency with multiservice transport effectiveness. Above all, it offers an ability to deliver stratified communications services that emulate the successful cloud services model. To close this paper we dive a little deeper into its details.

Benefits of Multilayer Optimization

MultiLayer Optimization (MLO) and MultiLayer Restoration (MLR) – that we refer to for convenience as MLO – is the most effective way to use network resources. MLO encompasses the IP connectivity layer and the optical transport layer, which includes the OTN electrical and DWDM optical sub-layers. MLO plans, implements, and controls the layers as a whole, rather than independently, as done today. Traffic demand requests are coordinated across layers, avoiding duplicate use of resources, while considering each layer's capabilities and shortcomings. The network structure adapts in real time to traffic demand requests, and strives to meet availability commitments in case of failures.

MLO is not a new concept. It is used occasionally today on a handcrafted basis for router bypass (transit traffic is transported at OTN/DWDM layer, skipping IP routers and saving on router ports and capacity), and for SRLG (Shared Risk Link Groups, making sure that backup paths use diverse fiber paths).
Moreover, paper studies have shown that MLO can provide significant resource savings, from 40-60% in the core and 25% in the metro/aggregation networks, respectively. (Multilayer Capacity Planning for IP-Optical Networks, IEEE Communications, Feb. 2014.)

So Why Hasn't MLO Happened Yet?

If MLO is so beneficial towards overall network efficiency, and can provide the basis for a 5G services network, it is fair to ask - Why don't we have it yet? There are several reasons:

- MLO is a difficult technical problem to solve algorithmically. Solutions depend heavily on service, network, and business models that are often are unavailable or not well understood.

- MLO can be difficult to operate. It not only introduces interdependencies between layers, but also between different telco operating departments.

- MLO removes redundancies between layers, which without proper engineering, can lead to a less resilient network. In addition, as mentioned, proper engineering is a difficult problem.

In short, there hasn't been sufficient justification until now to address MLO challenges. However, the revenue upside from many new services that 5G can introduce may change the game. MLO may be the only economically viable approach to gain the degrees of freedom and flexibility needed to engineer a 5G multiservice network, and unlock the additional revenues this can generate. We need a plan to proceed.

ribbon

Recipe for Implementing MLO

Step 1

Define a Meaningful 5G Services Suite

The fulcrum for justifying MLO is defining a broad and complex 5G services suite. Broadness boosts revenues by supporting multiple service categories for different levels of bandwidth, latency, availability, and duration. Complexity maximizes ROI by supporting tiers within each service category, particularly with respect to availability guarantees.

This actually complements MLO engineering complexity. For example, to perform multilayer restoration, services that are more varied and complex, particularly with respect to availability, can match the cost structure of the network. It becomes possible to create differentiated recovery schemes (committed, slow, eventual, best effort) which is a leap from today's one-class go/no-go approach to recovery.

Step 2

Define the Multiple Network Gears

Define the soft (shared but prioritized) and hard (dedicated) resources that are assignable to network slices. Soft resources can include traffic-engineering constructs like segment routing and dedicated queues, and virtualized routing processing for different classes of service. Hard resources can include router port assignments, flexible Ethernet, flexible ODU, OTN switching, software-definable optical line rates, and lightpath routing.
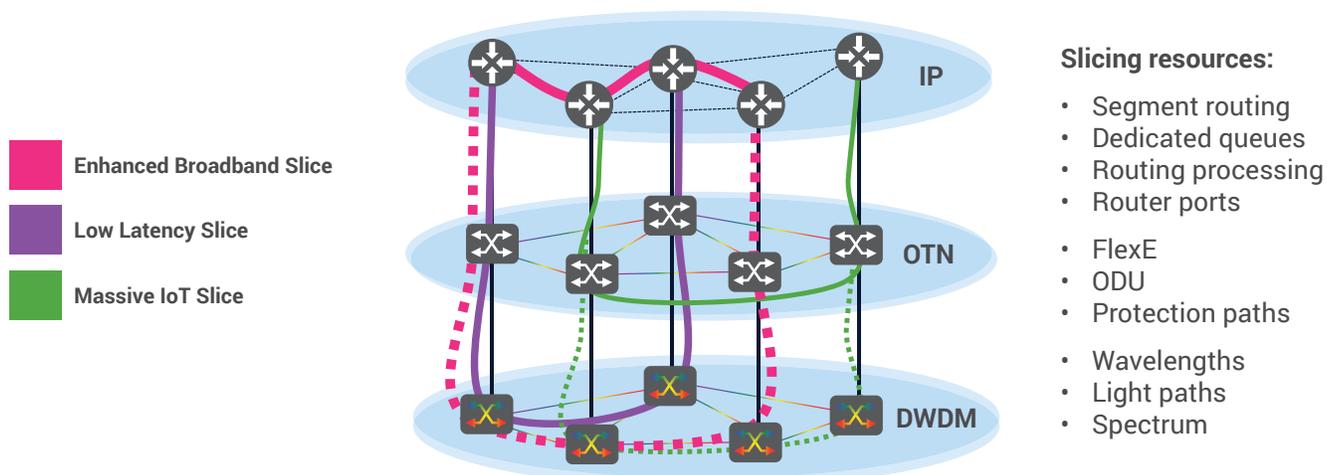
This step also defines the real-time telemetry requirements for each of these "gears", so they can report on their status and performance, and their real-time control interfaces.

Step 3

Define the Network Slice Constructs

It is too complex to map all types of services (step 1) onto a multi-gear network (step 2). We need an intermediary construct between the service types and the limitless network combinations. This is the concept of a network slice. The slices need to be multilayer-optimized, encompassing resources from various layers depending on service needs, with the ability to be organized (orchestrated) in real-time.

It should be possible to nest slices, to create a more complex slice from simpler slices. It should also be possible to offer slices directly to customers, as network-slice-as-a-service.

Enhanced Broadband Slice

Low Latency Slice

Massive IoT Slice

IP

OTN

DWDM

Slicing resources:

- Segment routing
- Dedicated queues
- Routing processing
- Router ports

- FlexE
- ODU
- Protection paths

- Wavelengths
- Light paths
- Spectrum

Step 4

Dynamic Service Creation Tool
Create a tool that aligns most popular service types, from the customer perspective, with the slicing constructs. Customers can choose a range of service availabilities, so long as they are compatible with available network slices that can guarantee the expected SLA. Capabilities that fall out of range are rejected, or accepted with a best-effort commitment.

Step 5

Implement Organic Networks
Steps 1-4 can be implemented using standard engineering approaches and algorithms. The network is planned and initially provisioned. The challenge starts as services and traffic demands appear and disappear, and failures occur. In all likelihood, to keep up with the complexity, machine-learning (ML) control will need to be adopted in stages. ML will also be able suggest service pricing, and based on the usage of the network gears, identify where enhancements would be beneficial to vendors. Above all, ML will continuously optimize the network to make the most efficient use of capital, and even identify where new investments need to occur to keep up with service demands.
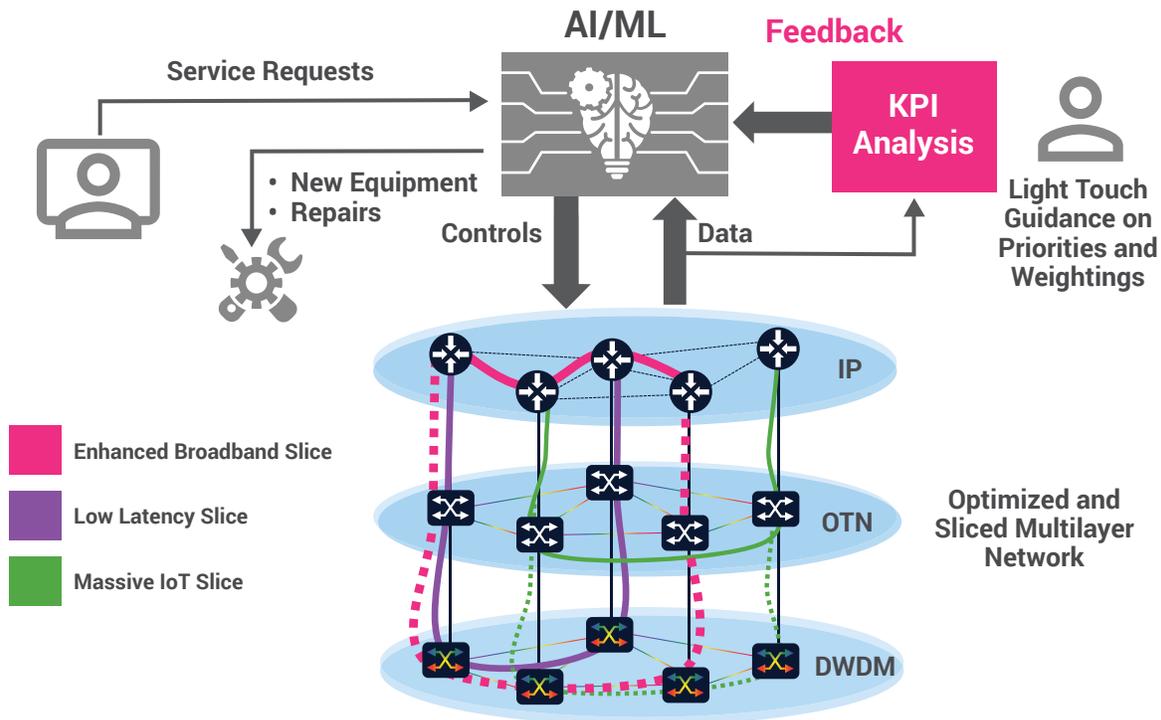
Ribbon Leading the Way

Ribbon is a leader in defining 5G network slicing transport architectures that leverage multilayer and multigear networks, and in creating a new generation of autonomous, automatic service-oriented orchestration tools.

Ribbon delivers unique value to SPs via its cloud-native Muse™ Orchestration platform. Muse customizes and seamlessly integrates its multilayer planning and provisioning tools to each SP's unique network and business needs.

Contact us to find out how Ribbon can help you build your 5G transport network at rbbn.com

ribbon

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.