State of the Art Packet and
Optical Networking

FlexE Applications in Transport Networks

White Paper

Table of Contents

FlexE Applications in Transport Networks

FlexEthernet (FlexE) is defined in Optical Internetworking Forum's (OIF) FlexE 2.0 as a manipulation of Ethernet rates by adding a layer of abstraction above the physical layer to provide greater flexibility in offering network services through bonding, subrating, and/or channelization of the Ethernet rates. FlexE offers advantages over traditional Ethernet, making it suitable for networks that require low latency and deterministic packet delay variation. These qualities are considered suitable for 5G transport that must deliver services with stringent network requirements. FlexE can also isolate traffic within the same physical interface through channelization, making it a viable option for transport network slicing. This versatility is applicable for 5G networks and is also expected to be adapted by fixed and mobile converged networks, critical infrastructures, wholesale networks, utilities, and other verticals.

Following OIF's introduction, FlexE is now being adopted by ITU-T recommendation for Metro Transport Networks G.mtn as an integral part of ITU-T Study Group 15 (SG15). ITU-T SG15 is scheduled to be finalized in 2020. As an enhancement, G.mtn will add protection and OAM functions to OIF's FlexE, making it perfect for ultra-lowlatency and high-availability networks. Ribbon is leveraging these enhanced network capabilities in a transport solution that enables operators to utilize FlexE as an underlay, with the option to combine it with statistical multiplexing on the upper layers. G.mtn's FlexE is part of Ribbon's network slicing options, providing a transport ecosystem that achieves the desired balance between traffic separation and cost efficiency. This white paper describes the network and service possibilities using FlexE.

FlexE Overview

FlexE adds a layer between Layer 2 (MAC layer) and Layer 1 physical interface rates (PHY layer). The main difference between standard Ethernet and FlexE is how the MAC layer is mapped onto the PHY layer, using a 'FlexE Shim'. Essentially, FlexE uses this shim to dissociate the Ethernet rate on the client side from the actual physical interface rate. This makes it possible to map Ethernet MAC rates, which can be greater than or less than the Ethernet PHY rates. The terminologies on the next page, as defined in OIF FlexE 2.0, best describe how this is achieved.

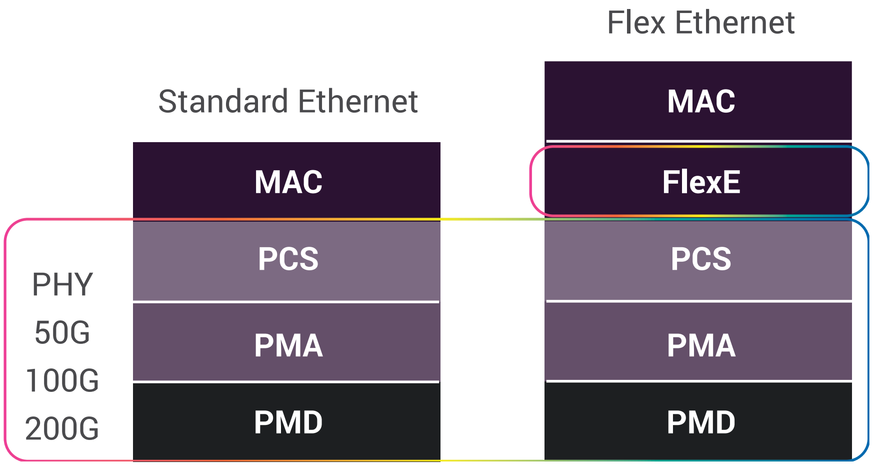

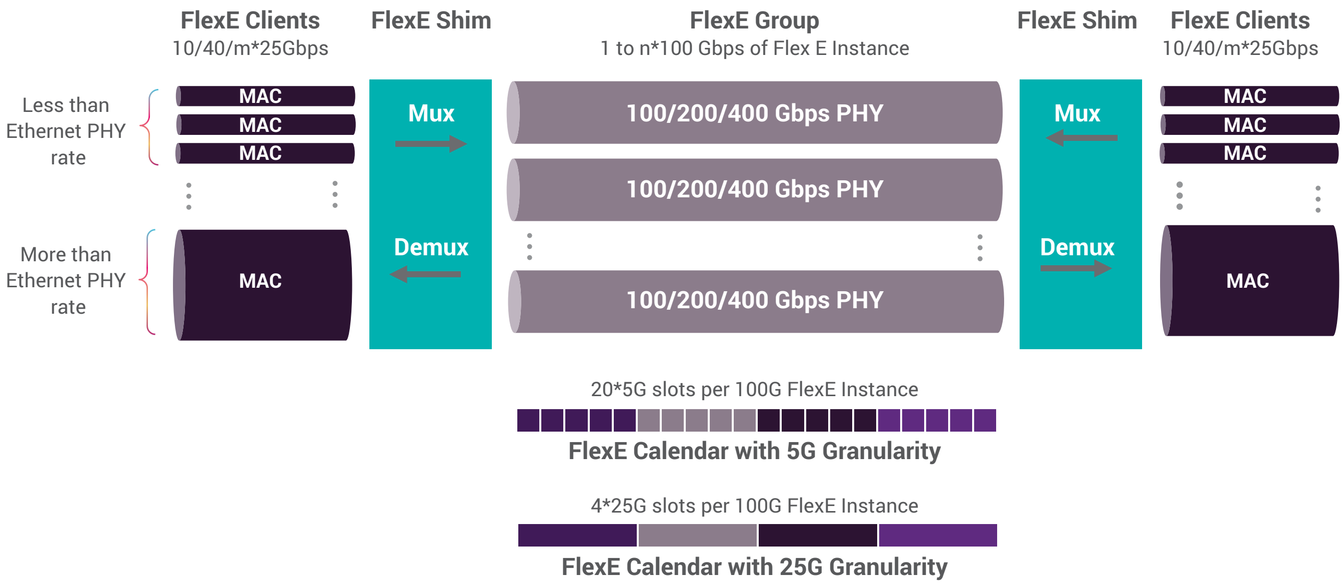

FlexE client is an Ethernet flow, based on a MAC data rate that may or may not correspond to any Ethernet PHY rate. This can be 10, 40, or m*25Gbps.

A **100G FlexE instance** is a unit of information consisting of 100G of capacity, able to carry FlexE client data, together with its associated overhead.

FlexE group refers to a group of one or many 100 Gbps FlexE instances, which may consist of 100, 200, or 400Gbps PHY/s. A 50Gbps PHY will be added within the OIF FlexE 2.1 Project.

FlexE calendar is used to assign TDM-like slots called calendar slots on each of the 100G FlexE instances. The calendar can have either 5Gbps or 25Gbps slot granularity. For example, there are 20*5Gbps calendar slots per 100Gbps FlexE instance and there are 4*25Gbps calendar slots per 100Gbps FlexE instance. The calendar slots, which carry the FlexE client traffic, are logically interleaved with other frames responsible for mapping, control, synchronization, error correction, and faults.

FlexE mux maps the FlexE clients to the FlexE group, and **FlexE demux** de-maps the FlexE clients from the FlexE group.

Below is an illustration of how a 10G client is muxed onto a 100G FlexE instance via 2 separate 5G calendar slots. The FlexE group comprises 4x100GbE PHYs. The 20x5G calendar slots that correspond to each 100G FlexE instance are then mapped onto the corresponding 100GbE PHYs. The overhead and other control frames are not shown here, for simplicity.

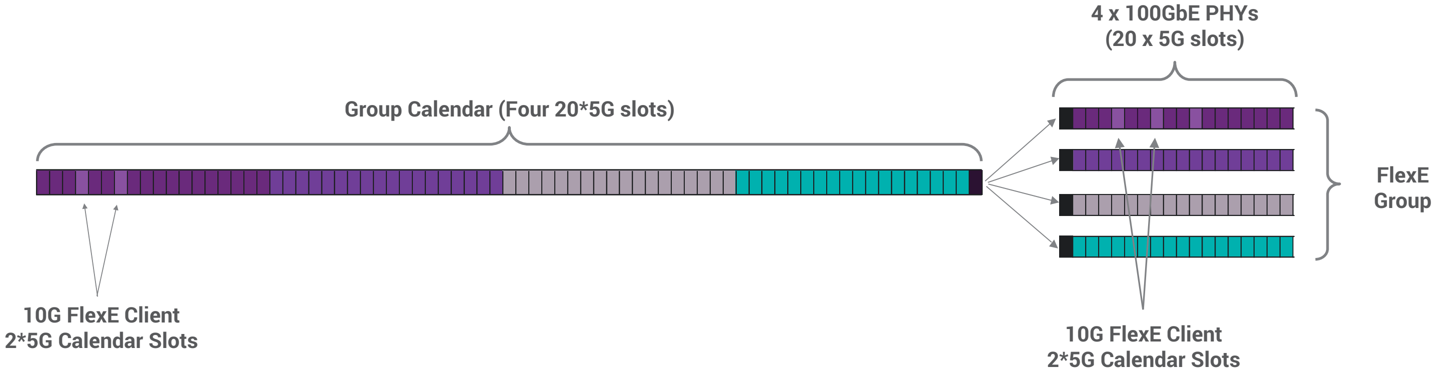

FlexE Mechanisms

The flexibility in mapping the MAC to PHY, together with the TDM-frame structure, achieves many possibilities in physical traffic separation, traffic multiplexing, and traffic cross-connection.

FlexE's three main mechanisms include:

1

Bonding of Ethernet PHYs

Bonding allows aggregating lower rates to form a higher rate. It is an alternative to LAG (link aggregation), in applications requiring deterministic performance and higher efficiency. Below is an example of bonding 4x100GEs to achieve a 400GE MAC rate between two routers. If a 100GE PHY fails, the link will continue to operate via the remaining 3x100GEs.

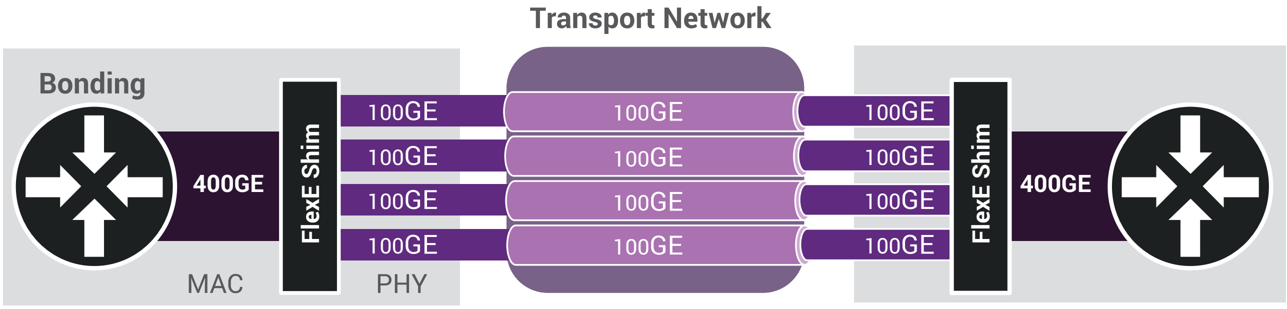

e.g. 400G MAC over 4 bonded 100Gbps PHYs

2

Sub-rating of the Ethernet PHY

Sub-rating allows a link to use a portion of a PHY or a group of bonded PHYs. In the example below, two routers are connected by 150G MAC rate over 2x100GEs. Each of the 100GEs are carrying only 75G of traffic. The unused extra capacity may be reserved for a planned capacity upgrade or it can be further used by other links, as shown in the case of the channelization example on the next page.

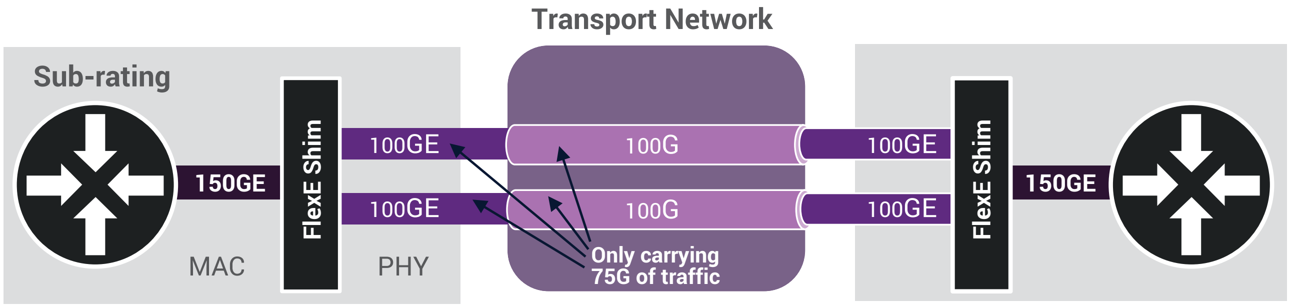

e.g. 150G MAC over 2x100Gbps PHYs

3 **Channelization within a PHY or a group of bonded PHYs**

Channelization allows one link to carry several sub-rates. There is total separation of traffic between the sub-rates, which is ideal for wholesale applications and network slicing. In this example, two routers are connected via 6x10GEs, 2x25GEs, 1x40GE, and 1x50GE over two bonded 100GEs.

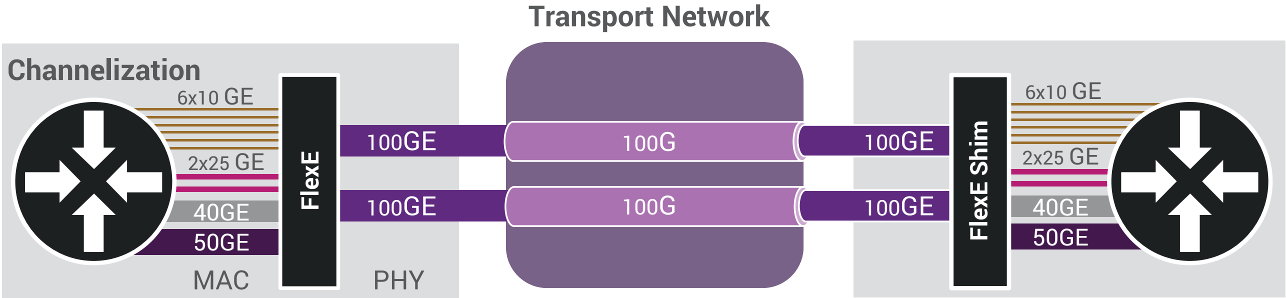

e.g. 40G MAC, 50G MAC, 2x 25G MACs and 6x10G MACs over 2 bonded 100Gbps PHYs

Underlying Transport Options

FlexE-unaware transport

This is the term used when the underlying optical transport does not process FlexE. This is useful when the network operator has an older optical layer network or is leasing standard Ethernet rates. In such cases, the PCS code word must be carried across transparently, which could be a challenge for some OTN networks. Also, all PHYs of a FlexE group must take the same fiber route, due to skew tolerance requirements. In the illustration below, the transport network uses Standard 100GE interfaces and carries the FlexE traffic across the transport network transparently.

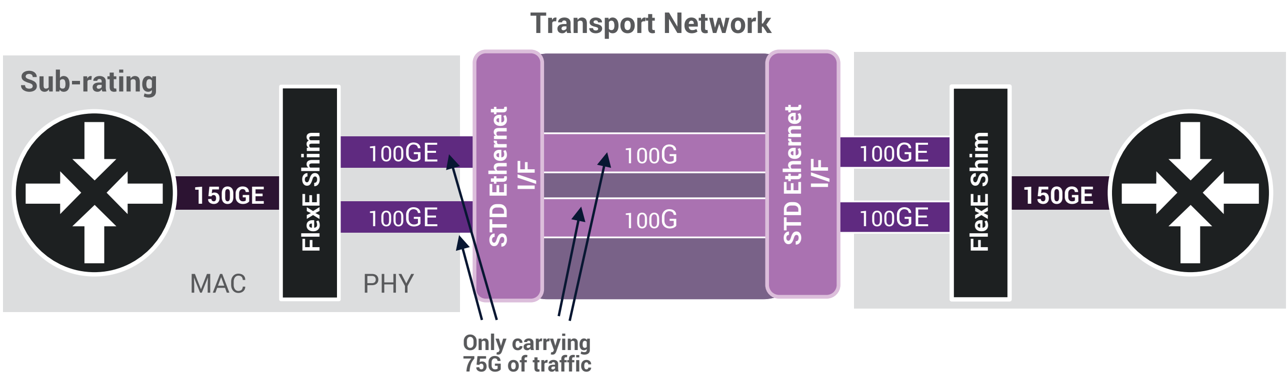

FlexE-aware transport

The transport network can process FlexE, but does not terminate it. It removes empty calendar slots before the FlexE traffic is forwarded across the transport. Here, the capacity is used more efficiently. As in the case of FlexE unaware transport, all PHYs of a FlexE group take the same fiber route. In the example below, 2x75Gs FlexE traffic is packed into 1x150G wavelength, using a FlexE-aware DSP (digital signal processor).

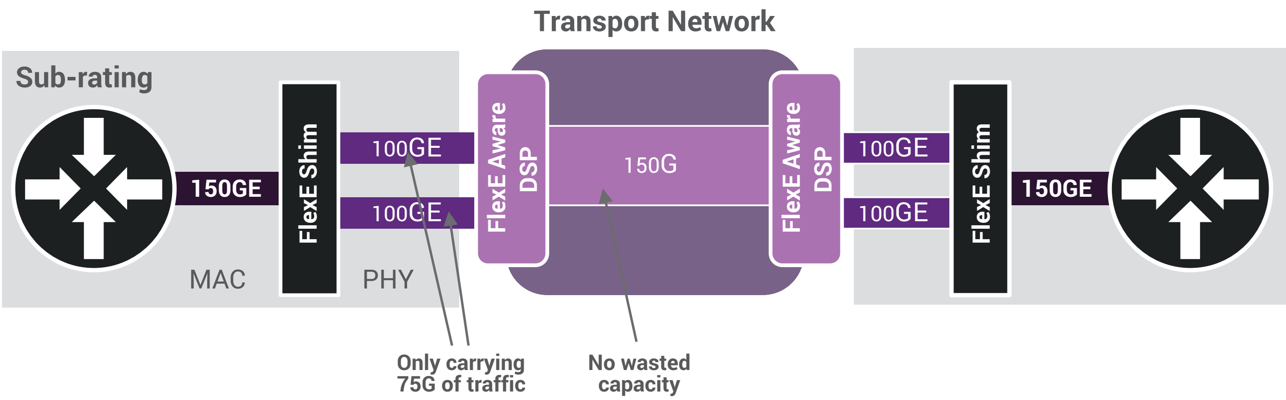

FlexE-terminated transport

FlexE is terminated before it crosses the transport network. This transport type has less skew issues because the distance between FlexE Shims is limited to a maximum of about 40km (Ethernet link distance). Instead of carrying PHYs, FlexE clients are carried across the transport network as shown below. The FlexE clients can be transmitted over wavelength or sub-wavelength services over OTN.

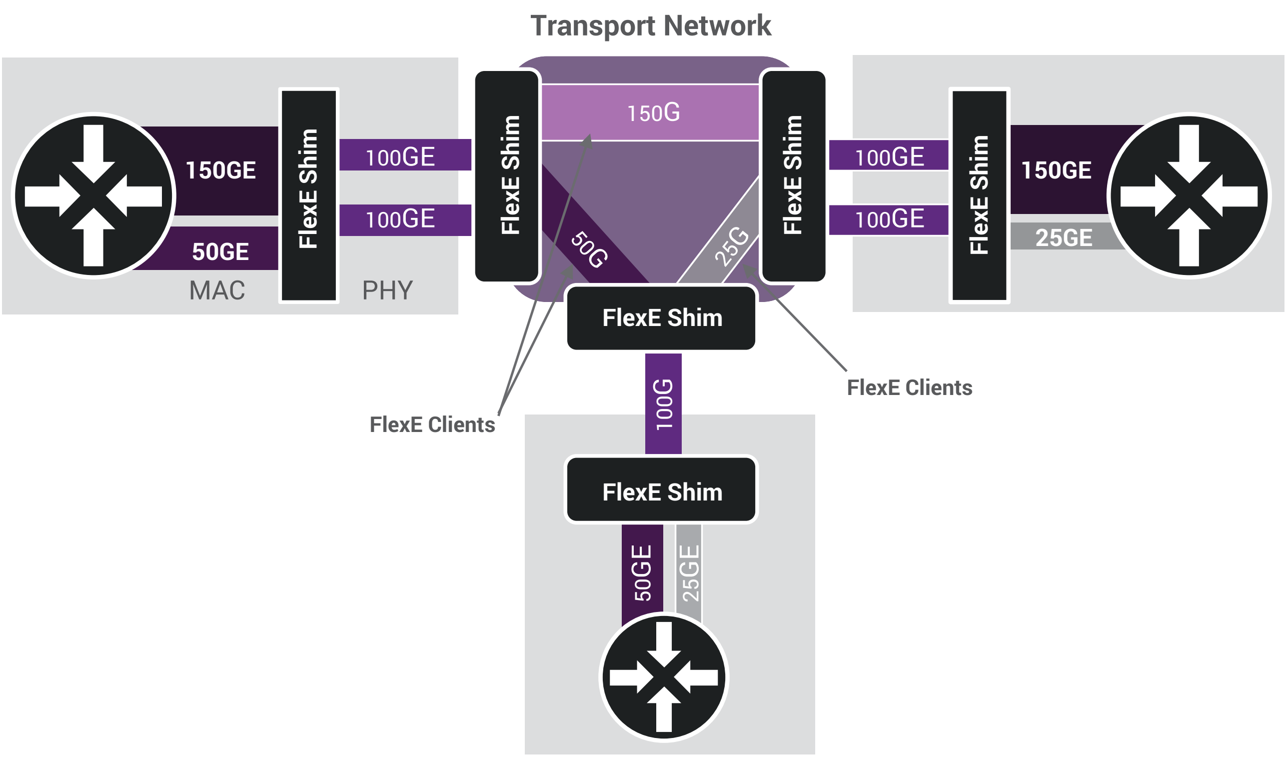

Deterministic Low Latency and Deterministic Packet Delay Variation

While packet networks employing statistical multiplexing are widely used for their cost effectiveness, the challenge remains, how to protect critical and high-priority traffic from traffic surges. With new low-latency services slowly becoming a trend, transport networks need to be more deterministic, while carrying the usual best-effort traffic. As an Ethernet-based transport, FlexE enables networks to carry packet services, while providing deterministic low latency and deterministic packet delay variation. Here are ways to achieve this:

Router bypass mechanisms

With FlexE, the switching process does not happen at the MAC Layer, compared to the traditional Layer 2 switches. The switching happens almost entirely at the PHY Layer. This reduces latency on paths, which pass through multiple hops, by eliminating the associated delays caused by packet processing at the intermediate nodes. In the example below, VPN AC with pass-thru traffic at Node B bypasses the packet processor almost instantaneously, while VPN ABC's payload is handled by the packet processor. In effect, traffic from A to C will experience lower latency when traversing VPN AC, compared to traversing VPN ABC. The router bypass mechanism, as illustrated on Node B, will have less than 2µs added latency. Thus, the cumulative latency of the end-to-end services can be kept to a minimal deterministic range.

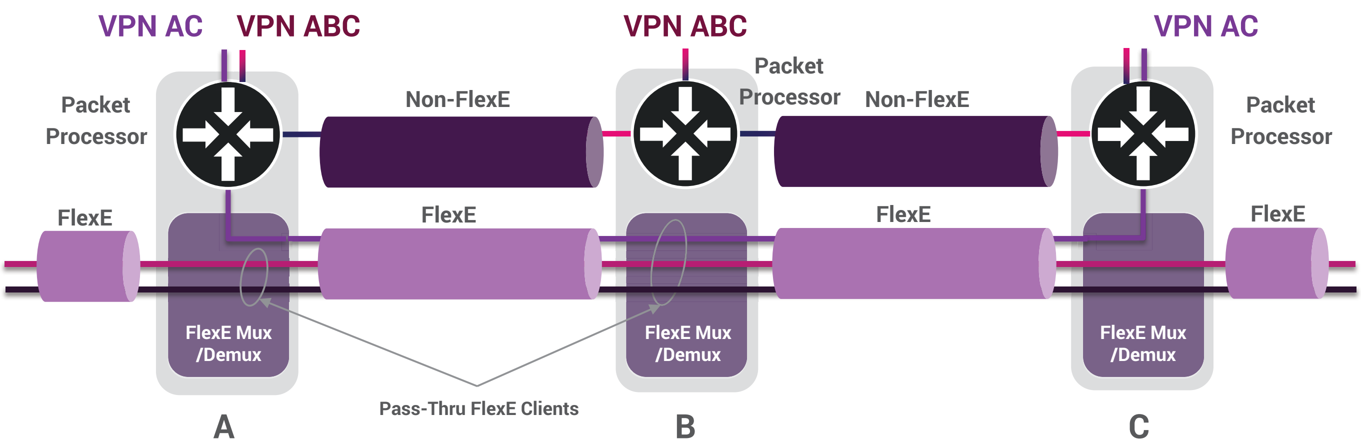

TDM frame structure

FlexE uses TDM principles to multiplex, transmit, and switch multiple Ethernet streams (FlexE clients), instead of using the flow control, address lookups, and packet buffers of traditional Ethernet transport. In the TDM approach, FlexE clients are allocated fixed timeslots in the transmission path, which mimics SDH behavior. Hence, delay from FlexE processing is very minimal compared to packet processing, and keeps the packet delay variation deterministically low. This is optimal in networks for the transport of interactive real-time applications/services. Such applications include real-time control systems, industrial machine-to-machine applications, and glitch-free real-time interactive multimedia.

FlexE Delivering Deterministic Services

There has been an industry push to develop solutions that support deterministic services. Much of this trend is due to the rise of real-time applications in automation and control systems, critical infrastructures, multimedia streams, and 5G networks. Purpose-built constant bitrate networks were never cost-effective enough. On the other hand, Ethernet and packet-based networks are not deterministic enough. A balanced mix of both would be optimal when deterministic services have to co-exist with best-effort services in the same network.

FlexE's deterministic properties makes it possible for an Ethernet-based network to deliver deterministic services, which were traditionally delivered via SDH, OTN, and WDM. FlexE was conceptualized with Ethernet in mind from the very start, so it can easily be integrated in existing Ethernet/packet networks. Network operators also have the flexibility of deploying FlexE, only on parts of the network, where deterministic properties are required. Therefore, much of the network can remain purely packet-based, where statistical multiplexing makes more sense, making the overall solution cost-optimized. The VPN ABC and VPN AC example in the previous example is a typical scenario showing the existence of a non-FlexE connection together with a FlexE connection on the same network element.

FlexE in Building 5G Transport

Since the initial FlexE standard in 2016, FlexE has already been deployed in networks delivering flexible Ethernet rates. It is continually being tested by both telecom equipment vendors and telecom operators for new applications, especially in 5G. Industry acceptance of FlexE is also progressing with some key events in 2018 and 2019 as the industry prepares for the upcoming 3GPP Release 16, due in 2020. In late 2018, OIF announced the launch of the FlexE 2.1 project, which enhances the FlexE 2.0 capabilities to include 5G use cases, adapting 50GBASE-R PHY in FlexE Groups.

At about the same time in 2018, ITU-T initialized the study work on G.mtn (ITU-T's recommendation for metro transport networks), which covers considerations in D-RAN and C-RAN transport. This G.mtn study is considering the slicing packet network (SPN) technology proposal, which aims to provide low latency, high bandwidth, flexible management, and control using FlexE in network slicing. G.mtn also enhances FlexE by adding protection and OAM features on top of OIF's FlexE 2.0, making it fit for services requiring high availability. Ribbon adapts FlexE, according to the proposed G.mtn recommendation and incorporates the solution as one of the options in Ribbon's portfolio for building 5G Transport.

Transport Network Slicing with FlexE

Network slicing is the partitioning of network resources into logical networks, called network slices. These are assigned to specific services or customers, each having different network requirements and topologies. While this network principle is driven mainly by 3GPP for 5G mobile networks, it is also becoming a practical approach for transport networks that carry diverse types of services for different purposes. Examples include fixed and mobile converged transport, wholesale transport networks, Utelcos, and private networks that require high differentiation of services.

The gain in popularity of FlexE has led some service providers to consider the use of FlexE's ability to separate traffic at the physical layer, while coupling different Ethernet rates to the optical layer. It is now possible for an Ethernet-based transport to separate different types of services within the same transmission path or port.

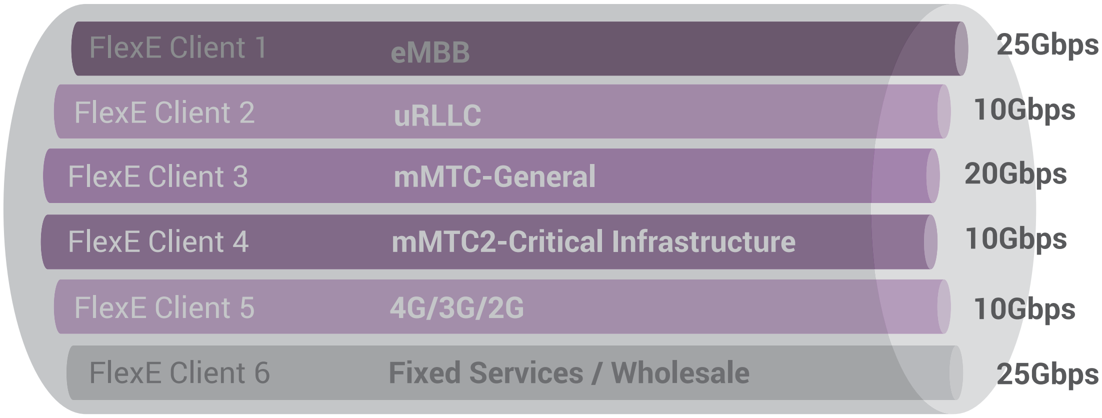

100G FlexE

For example, we can take one 100G FlexE instance and divide the capacity into 6 slices corresponding to 6 FlexE clients. Each slice carries one type of service. Each of these service slices can have different network quality requirements. FlexE can ensure that traffic from one FlexE client will not influence the traffic of the other FlexE clients. This is called hard isolation of network resources.

FlexE is not intended to be deployed on its own. It is part of an ecosystem of layers and technologies, which can be assembled to comply with network requirements while optimizing costs. FlexE functions as an underlay to provide a very high degree of isolation of network resources.

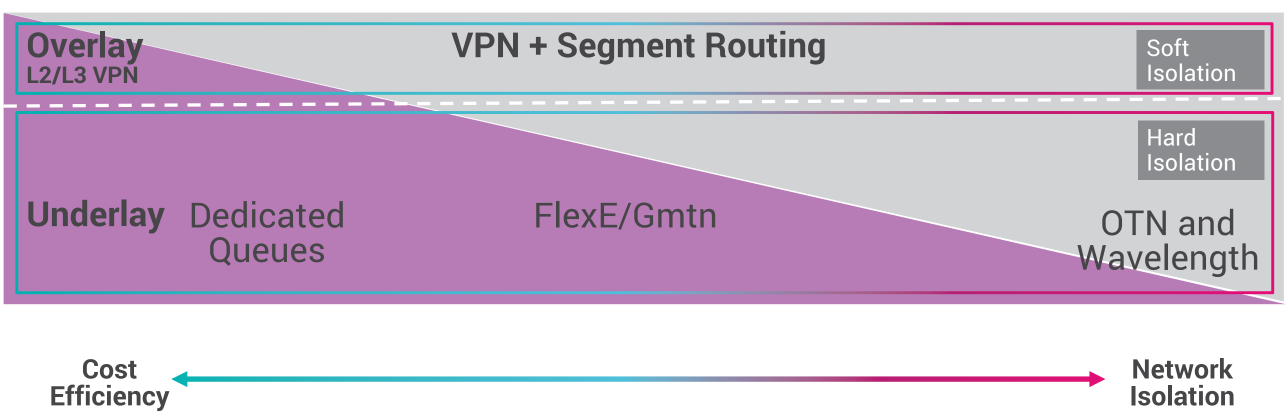

It also provides higher granularity compared to wavelength, which is at the extreme end of the network isolation scale. By packing in more slices on FlexE, the capacity of a wavelength can be optimized. Further optimization can then be achieved when coupling FlexE with advanced Layer 2 and Layer 3 technologies on the overlay, leveraging traffic engineering and QoS techniques. The example below illustrates how a 100G FlexE instance on a wavelength can carry 4 network slices, while maintaining a pragmatic balance between hard isolation and statistical multiplexing.

	Slice A	Slice B	Slice C	Slice D
Shared Resources Statistical Multiplexing	L2/L3VPNs (*Enhanced VPN)	L2/L3VPNs (*Enhanced VPN)	L2/L3VPNs (*Enhanced VPN)	L2/L3VPNs (*Enhanced VPN)
	*SR Tunnels	*SR Tunnels	*SR Tunnels	*SR Tunnels
Dedicated Resources Hard Isolation	FlexE Client: 25G	FlexE Client: 25G	FlexE Client: 25G	FlexE Client: 25G
	FlexE Instance: 100G			
	Wavelength			

*SR - Segment Routing
*Enhanced VPN based on IETF's draft-ietf-teas-enhanced-vpn

Ribbon's Approach to Transport Networks

As services become more and more diverse, transport networks have to cope with very different requirements, sometimes on the same path. To solve these challenges, Ribbon devised a versatile transport network solution to serve all types of services by incorporating technology building blocks, which can achieve best-in-class transport. The example above demonstrates how different types of resources, operating in different layers, can achieve optimal delivery of services.

On top of Ribbon's FlexE layer, Ribbon's transport network solution uses topology-aware segment routing tunnels and resource-aware enhanced VPN, based on IETF's draft-ietf-teas-enhanced-vpn to transport the services. In establishing end-to-end connectivity, Ribbon's network intelligence sits on a software-defined networking platform, hosting centralized network intelligence with an advanced traffic engineering mechanism. Merged with FlexE, it is now possible to have a cut-through view of the transport network, from Layer 3 down to Layer 1. This achieves the desired multilayer approach for transport networks.

Conclusion

The telecom industry is once again experiencing a new wave of technology changes triggered by the development of new service types. This technology trend is being experienced, not only among service providers, but also in the vertical industries. These brand new services require new ways to build networks, since existing networks, at some point, will no longer be able to support the new requirements.

Efforts from different sectors in the telecom industry are being made to create new sets of recommendations and standards, so that network operators will have options and toolkits to help build the networks of the future. FlexE is one these toolkits. Starting from OIF, FlexE made its way into ITU-T and is slowly making its way into IETF, ensuring continuous development of the technology.

FlexE's deployment has initially appeared in the more aggressive and innovative sectors in the market and industry, where the transport networks are built in advance, anticipating new services. However, as these brand new services become more of a norm, we will see FlexE making its way to the rest of the world, thanks to its versatility and flexibility.

Contact Ribbon today for more information about the FlexE solution at rbbn.com

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.