NETWORK SLICING

5G will enable carriers to expand their service offering to include a complete range of telecoms services to individuals, businesses, enterprises and critical industries. However, each of these sectors has a different range of set of services each with their set of policies and parameters.

Mission critical services, for example, may need permanently "always-up", deterministic, low latency, low bandwidth connectivity. Whereas broadcasting footage of an event, to people at the event, so that they can look at replays and close-ups requires ultra-high bandwidth connectivity, at the arena, during the event itself. Video streaming 'on the go', on the other hand, requires the connectivity and capacity to follow users.

5G envisions a universal services platform capable of supporting these diverse services with a single, unified network.

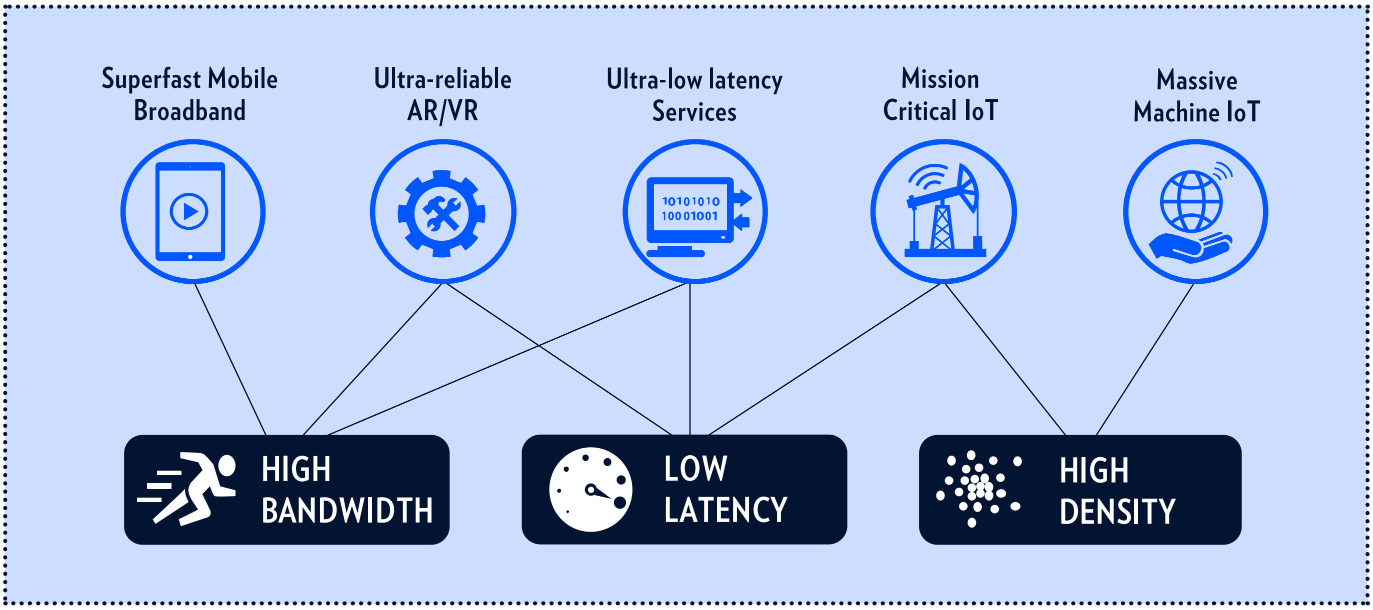

Network slicing is an approach proposed to allow a single network to support services with completely different operational parameters and policies.

In effect the network is viewed an asset pool of physical network resources (PNFs), virtual network resources (VNFs), connectivity, bandwidth and compute. A network 'slice' combines these assets to form a virtual network.

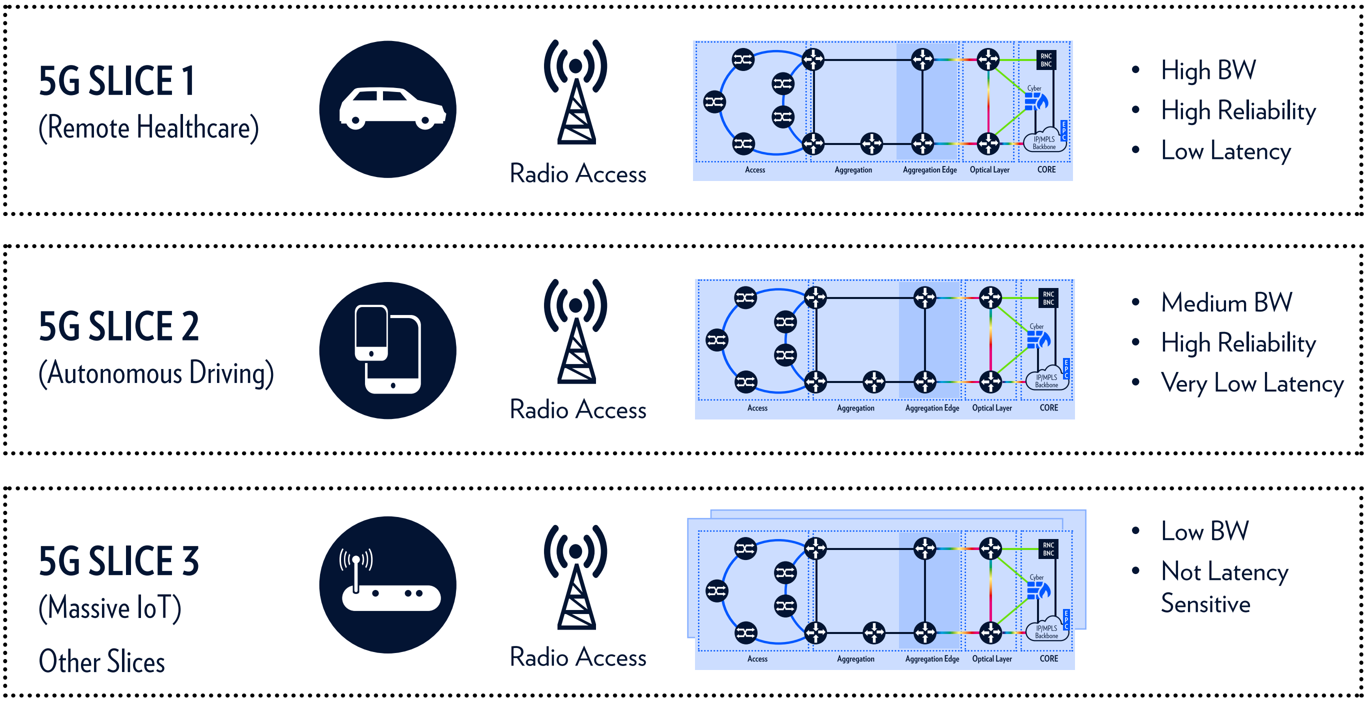

Different network slices will have different operational parameters and hence a different combination of assets and attributes. The slices may share network assets or they may have assets specifically allocated to them, depending on the service policies.

Service orchestration provides end to end orchestration of how the services cross the domains that make up the 5G network. The domain controller provides the linkage between the service requests coming from the orchestration layer and the configuration required in the transport layer. The domain controller is intent driven and provides autonomous configuration of the transport assets, as it receives service requests.

Multiple options exist for providing network slicing and, depending on the network dynamics, these can be used in isolation or together.

Hard Slicing: Elastic OTN and Wavelength Slicing

OTN slicing and wavelength slicing is used to create network slices directly across the optical transmission network. Router bypass can be applied for services travelling directly from the edge to the core, providing the lowest possible latency and freeing up router ports.

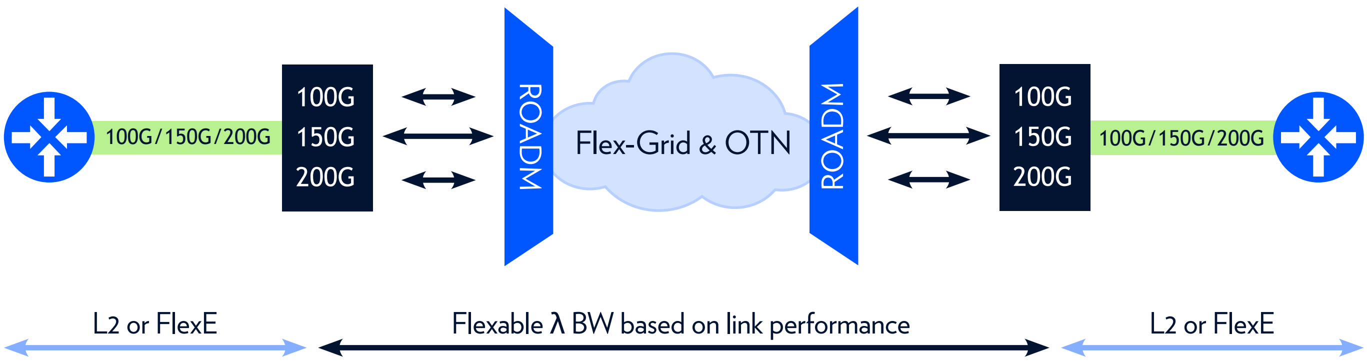

Service Provisioning (IP Slicing): Segment Routing (TE) and Enhanced VPN

Segment routing uses a centralized control function (PCE) to calculate the route across the network. The control function has full knowledge of the transport resources. So the path calculated can be as defined as rigidly as necessary to meet service policies and parameters, hence providing the determinism required for many of the new mobile services.

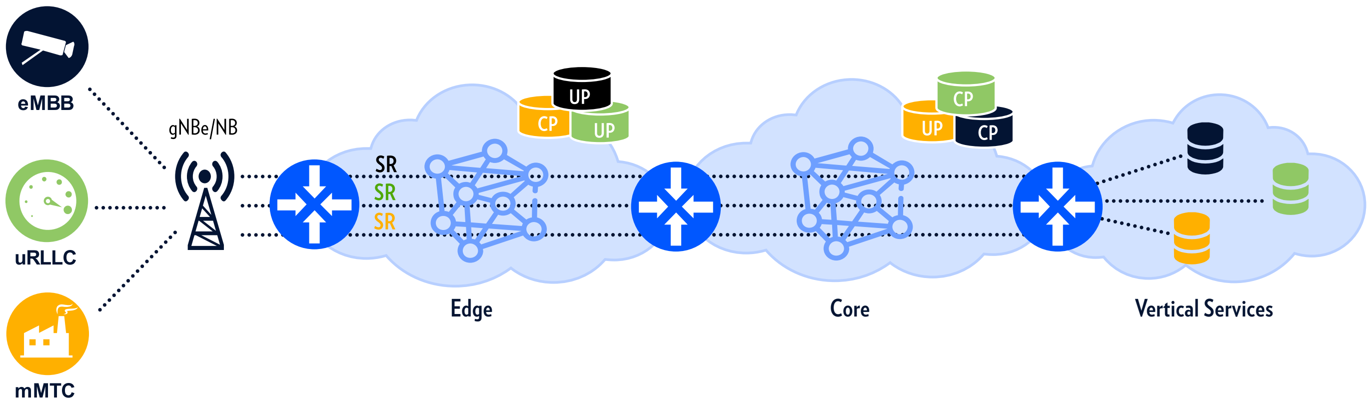

Elastic Slicing: Flexible Ethernet

Flexible Ethernet (FlexE) provides deterministic transport by using TDM principles to remove the need for flow control when mapping packet services onto optical transmission paths.

The FlexE pipes are used to create fixed bandwidth connectivity to link together specified network assets for each required network slice. It essentially acts like a clutch mechanism between the IP slices and the OTN slices.

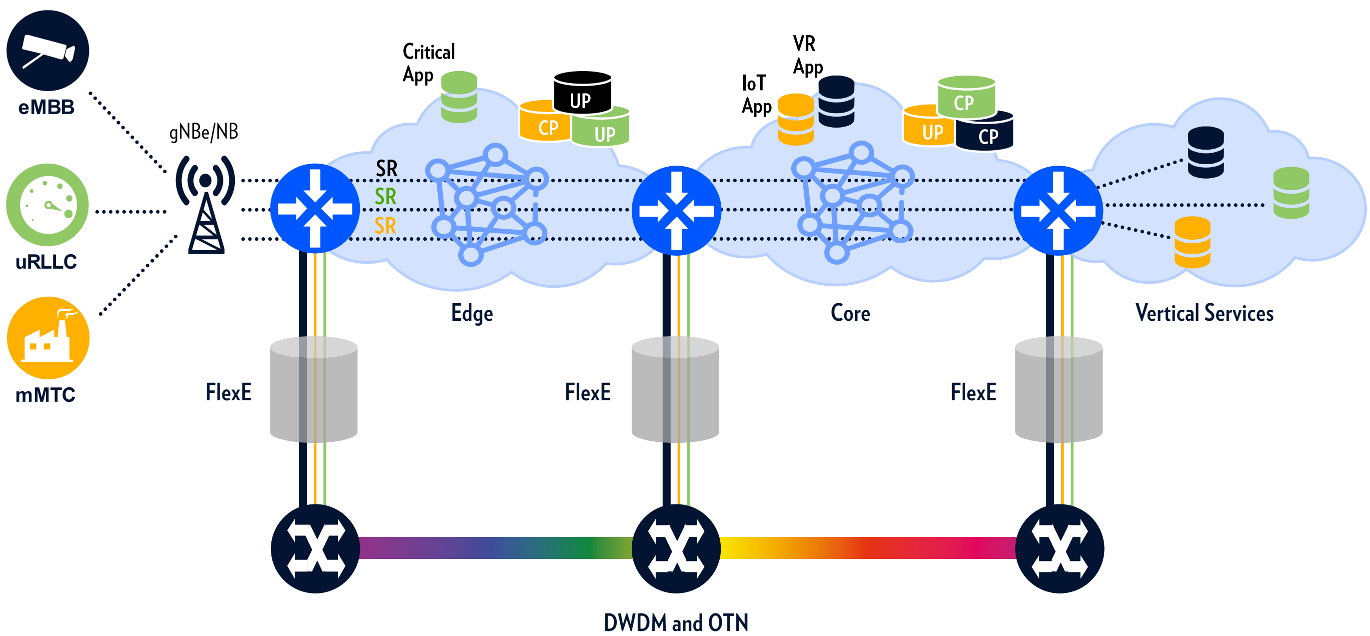

BRINGING IT ALL TOGETHER

The mobile network operator will initially configure the network to meet the service types in their portfolio and then as services are activated these will be activated on the relevant slice, or a new slice will be created if the existing slice does not meet the service parameters and policies.

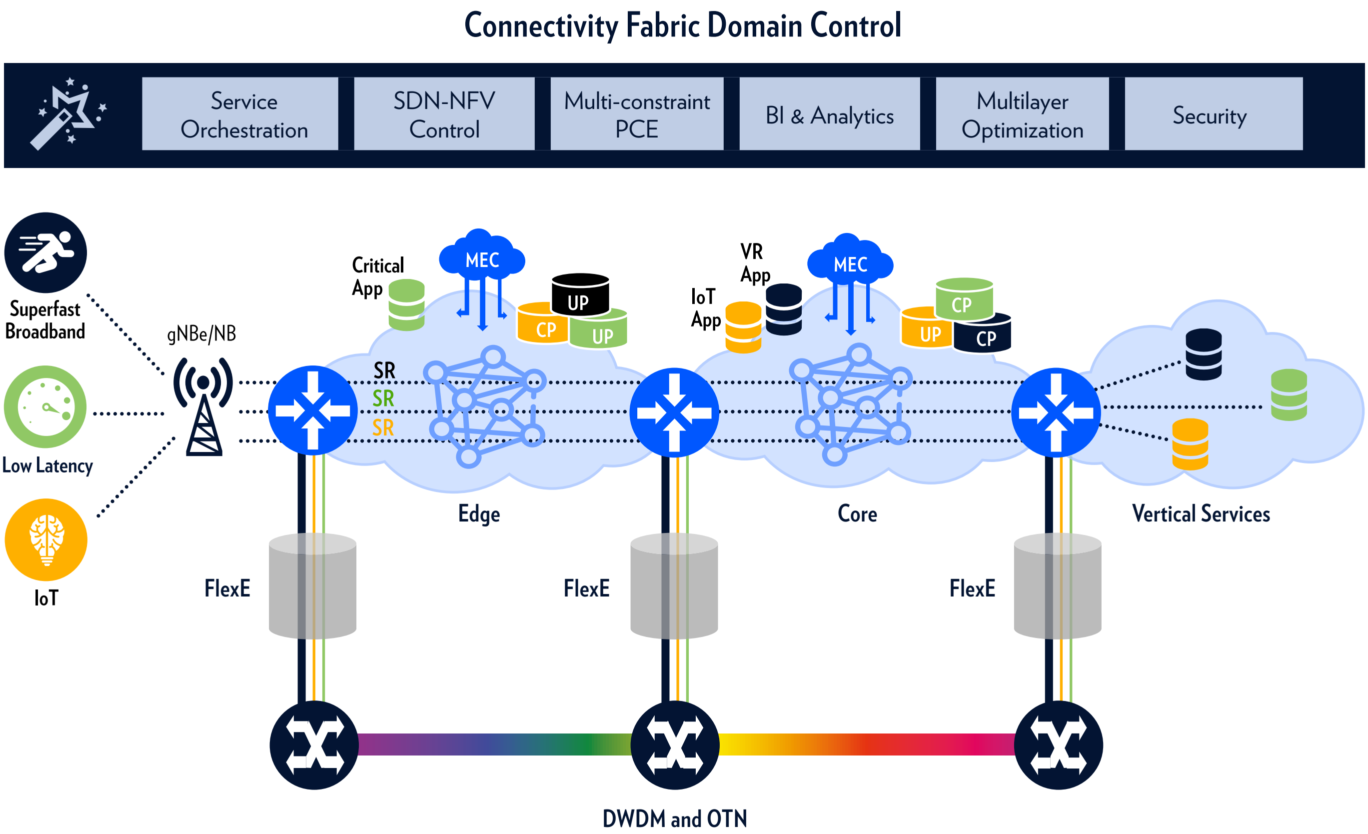

So for the sake of simplicity let's assume that the MNO initially offers service types based around the service types identified below:

- **Superfast Mobile Broadband**
- **Basic Voice Services**
- **Video Call**
- **Ultra-reliable services**
- **Ultra-low latency services**
- **Mission Critical IoT**
- **Massive IoT Services**

A slice is created for each of the service types using the assets available from the common network infrastructure. The resources and assets assigned to each slice are reserved for the use by services that run over that slice and cannot be used by other slices. Each slice is isolated from each other slice such that the performance and loading of one slice will not impact the performance and loading of the other slice. In essence a virtual private network is created for each service type, much like building a separate network for each service type, but assets may be shared across multiple slices as long as they can be isolated. So for example the compute asset offered by the network maybe used by multiple slices, with one slice running two virtual machines (VMs) and another running three VMs on this asset. Or one slice may require 100G capacity through a node and another may require 200G capacity, as long as the capacity is reserved then this is acceptable.

So for example let's look at the massive IoT services slice. This could be a very simple slice providing connectivity for water meters and electricity meters:

So as new meters are added to the network, the "meter service" is allocated to the correct slice, which is ensured by the domain controller. It will also ensure the service is routed correctly across this slice and that the user plane and control plane data go to the correct UPF, AMF and SMF functions.

In planning the connectivity fabric for this slice, assets and resources will be allocated to the slice, based on the anticipated service profile (growth, geographical distribution, etc.). The domain controller will monitor resource utilization and identify when new resources are required.

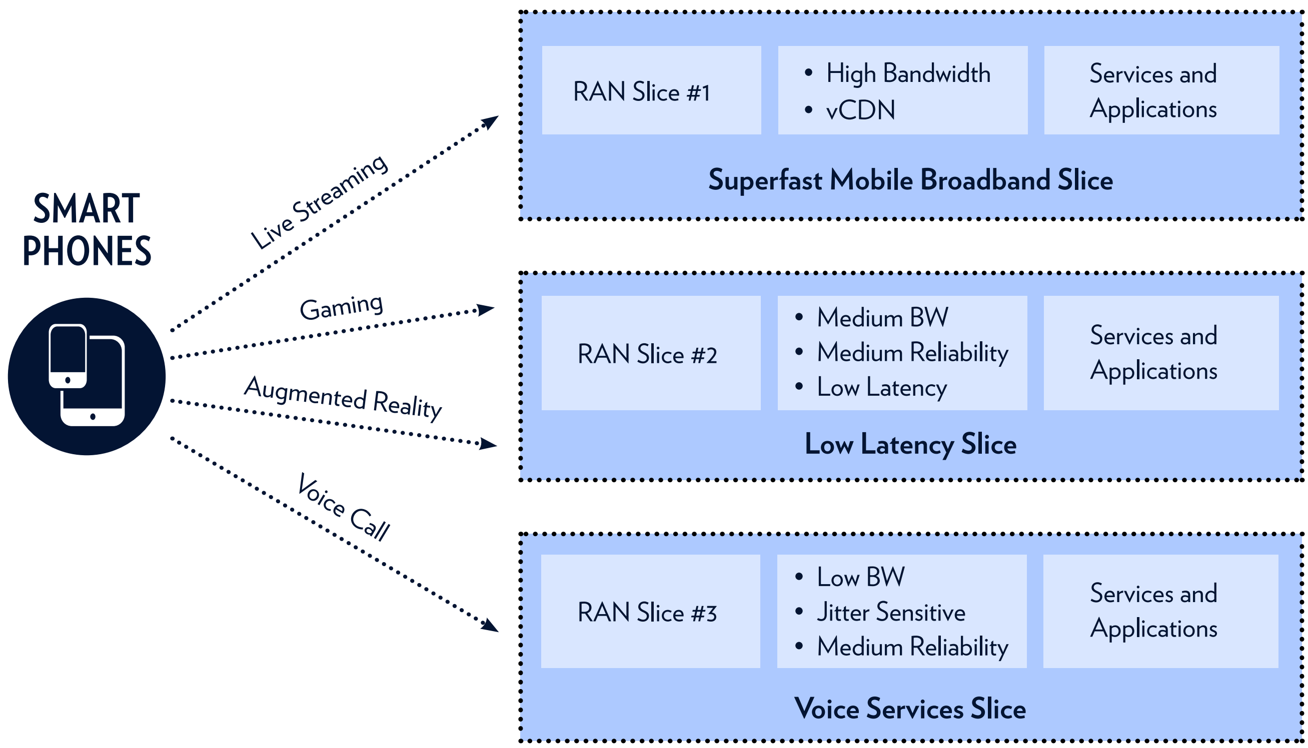

The "meter service" is an extremely simple, but real, service example, but it helps explain the principals involved. So let us now examine a more complex example, such as a smart phone.

The smart phone will connect to multiple network slices. The services the smart phone is running will determine which slices it is connected to, however let us assume that it is connected to a "superfast mobile broadband slice" for streaming video, a "low latency slice" for gaming and a "voice services slice" for voice calls

CONCLUSION

The success of 5G depends on the ability to support a diverse set of services with varied requirements for latency, throughput and availability. The 5G network needs to be highly adaptable to support new services and business models that will emerge with 5G. Today's "best effort, one-size-fits-all" approach to mobile networks is just not viable.

Network slicing enables operators to create and support multiple virtual networks on a common network infrastructure. This approach allows operators to deploy only the specific functions needed to support specific use cases. In addition, for the first time, network slicing gives operators the ability to provide the service differentiation which is not available in traditional mobile networks.

Network slicing uses the agility provided by SDN and NFV, the determinism provided by Segment Routing and FlexE and optimized transport (both packet and optical) to support a full range of services for the full range of customers.

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