www.pipelinepub.com
Volume 17, Issue 8

Smoothing the 5G Transition with Traffic Engineering

By: Irit Touitou

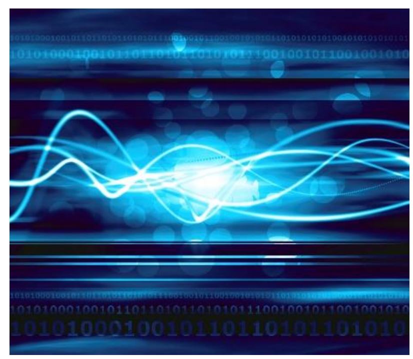

Communications service providers around the world have experienced an unprecedented explosion in demand for bandwidth. This is the result of several driving forces, including more video streaming, the growing popularity in areas like cloud gaming, and the increasing move to the cloud, to name a few. The increased demand for bandwidth has only accelerated with the pandemic, which triggered a dramatic increase in the number of people working from home, a spike in remote learning, a surge in video meetings, and other online activities.

In addition, a vast number of current communications networks have physical limitations and just cannot keep pace with the many other bandwidth-hungry demands of today's consumers and businesses. These communications networks are also often hampered by high maintenance costs and lack the ability to quickly roll out new services to customers.

These challenges—and many others that service providers face—are only magnified with the rollout of 5G networks and the expectation that they will deliver much higher capacity per end-user and lower latency. 5G will enable service providers to deliver highly diversified services, which will in turn introduce diverse and dynamic network utilization patterns and higher timing demands to ensure limited latency. This means these networks will need to be optimized to adjust to increased network utilization patterns and ensure a better end-user experience.

Currently however, many service providers are relying heavily on their 4G networks, with 5G being a part of their future but not necessarily their immediate plans. So how can service providers leverage their existing infrastructure to overcome bandwidth challenges and at the same time ensure a smooth, gradual transition to 5G?

The good news is that service providers can meet these challenges with their existing network infrastructure. They can capitalize on higher capacity and lower latency, making their transition to 5G more gradual and evolutionary. Moreover, they can do this in a manner that allows them to achieve faster return on investments and potentially accelerate their 5G rollouts. By optimizing current network utilization through a process called traffic engineering, and via improved network automation, service providers can extract more resources from their existing network. In addition to providing them with a more efficient network, traffic engineering and enhanced network automation also allow them to create the foundation for a 5G-ready network that will enable a smooth transition to a full rollout.

A primer on traffic engineering

As background, let's look at why traffic engineering is essential in network performance, optimizing changing network traffic patterns, and supporting 5G rollouts. The concept of traffic engineering is actually not new, but it was less scalable and less effective in previous iterations (pre-IPv6 and pre-segment routing-traffic engineering). As service provider networks have become more complex, traffic engineering has become simpler to implement and is a much more effective tool. Conversely, network resources have become more limited as bandwidth demand has spiked. Dynamic traffic with diverse services only adds to the network complexity.

Service providers run the risk of overwhelming their networks with these spikes in demand, while at the same time underutilizing network resources in other areas. These challenges only intensify with 5G and its special requirements for reserved bandwidth, deterministic latency, enhanced isolation, and ultra-reliable low latency.

For example, for massive IoT use cases like managing irrigation systems or smart cities, bandwidth may be most important. For gaming, augmented reality (AR) and virtual reality (VR), deterministic latency is important. For defense systems, isolation is expected for maximum security. And for remote surgery, ultra-reliable low latency is a must. With traffic engineering and network automation, we have an opportunity to look at different types of network behavior and requirements (like load hours, mobility, and latency) and adjust the network to traffic patterns based on these different scenarios. Traffic engineering using segment routing (SR) is designed to provide service providers with the ability to dynamically differentiate and steer traffic and optimize the use of network resources. The source routing principle used in segment routing traffic engineering (SR-TE) enables steering of the packet by simply changing the instructions on the packet header at the head end network device; there is no change needed at the transit nodes. This improves both the service and network performance.

Consider the following example as it relates to mobile communications services. A group of mobile consumers commuting from the city center to the suburbs on a bus are streaming a live football game during rush hour. They are paying a premium for this streaming capability and expect an uninterrupted user experience, even during heavy user periods.

The ability to steer network traffic dynamically enables improved mobility and justifies the premium service charge for the consumers. In network terms, traffic has to be steered to adapt to the shifts of the load in the network. The best effort traffic will have to be shifted to take longer routes in the presence of livestreaming traffic, decongesting the premium route, so latency figures can be maintained.

In general, dynamic adjustment to changes in traffic patterns is required for an efficient and reliable network. If such change in the network requires manual configuration, there is a high probability that the network adaptation will be too late, causing intermittent service degradation and negatively impacting the user experience. If the network adaptation is too early, network resources could be unnecessarily wasted. This illustrates the importance of automation of traffic steering by dynamic and closed-loop systems coupled with a programmable network infrastructure.

Traffic engineering benefits

Traffic engineering provides increased bandwidth and network reliability assurance by computing an end-to-end path for routing packets. The route is pre-planned by offline optimization algorithm or by real-time computation as opposed to locally finding the best next-hop to the destination. Leveraging the pre-planned routes allows service providers to preserve valuable network resources such as bandwidth on each port or link and ensures protection against node and link failures by pre-computing back-up paths.

Both service and network performance are dramatically improved by having guaranteed bandwidth, regardless of if there are failures, heavy traffic loads, or congested nodes or links. Traffic engineering allows service providers to seamlessly manage and adjust to changing traffic patterns. In addition, the closed loop automation of the network allows service providers to redirect traffic to new paths to react to network utilization or prepare for future demands. These traffic flows can have different paths and be assigned different levels of protection based on the scenario.

Traffic engineering is not only applicable to 4G environments but is also instrumental in improving 5G network operations. In a 5G environment, service providers will be introducing more services, managing higher bandwidth demand, and experiencing accelerated rates of change in traffic patterns—meaning the need for traffic engineering exponentially intensifies.

The next-level benefits of 5G network slicing

Network slicing is one of the many highly anticipated benefits of 5G. With network slicing, service providers can assign each service the exact network it needs. Combining 5G network slicing with segment routing enables a higher level of traffic engineering and can offer better network and service performance. Traffic engineering enables service providers to assign the appropriate network resources to each traffic pattern in each slice.

At a high level, 5G network slicing takes traffic engineering to a new level. The end-to-end concept enables multiple tenants over single physical infrastructure, so service providers can have independent operation and slice resources management per tenant. The slicing capabilities enable strict resource isolation, multi-tenancy, independent management, and orchestration per each slice along with enhanced security. Network slicing allows service providers to manage service level specification (SLS) and quality of service (QoS) per slice, whether it is via hard slice (such as IP/MPLS over FlexE) or soft slice. And, network slicing enables an increased level of sophistication by delivering customized traffic engineering policies for each 5G service that can run independently in each slice, assigning the appropriate network resources to each traffic pattern. Combining slicing with traffic engineering creates opportunities for new revenue streams for service providers, as the ability to assign slices per flow on demand opens new domains. The slice itself becomes a resource that service providers can sell or lease. This enables new applications, offering slices separately, as a bundle, or even grouped in a hierarchical manner.

Examples of this include large enterprise units (increased bandwidth) or military bodies (isolation) seeking independent networks within the organization. It also introduces infrastructure-as-a-service where slices can be assigned to potentially different customers on demand and ad-hoc. So, new business model innovation and use cases across all verticals are enabled by the ability to deliver services faster with high security, isolation, and agreed-upon service level agreements (SLAs).

Traffic engineering offers service providers a number of benefits that can all lead to a premium user experience and a more reliable, efficient network. It also grants service providers the ability to leverage existing 4G infrastructure and gradually expand with 5G-ready capabilities that can easily be upgraded to 5G, as well as invest in network automation that enables faster deployment of services and dramatically reduces the possibility of human error. Once service providers transition to 5G, the combination of slicing and traffic engineering opens the door to new revenue streams and business models—essentially enabling them to future-proof their networks.