Independent market research and competitive analysis of next-generation business and technology solutions for service providers and vendors

HEAVY READING
WHITE PAPER

Converged Packet Transport for 5G

A Heavy Reading white paper produced for Cisco Systems

AUTHOR: STERLING PERRIN, PRINCIPAL ANALYST, HEAVY READING

INTRODUCTION

Network operators and governments around the world are betting heavily on the 5G future. While industry conversations revolve around new 5G services, the network – including the wireline transport network – plays a fundamental role in delivering them.

As 5G New Radio (NR) standardization progresses, operators are increasingly understanding that transport networks must be upgraded and re-architected to support 5G applications at scale. This white paper addresses 5G evolution from a transport network perspective, emphasizing the role of a converged packet network capable of addressing all 5G use cases, as well as diverse mobile and fixed services. Such a converged packet network is critical for making the 5G business case work.

STANDARDIZATION STATUS

While trials and isolated deployments can take place pre-standard, full standardization is required for operators to roll out 5G services on a massive scale. The 5G NR draft specifications for the first phase were approved in December 2017 (3GPP Release 15), and work on the second phase (3GPP Release 16) is scheduled to be finalized in December 2019. Release 15 introduced non-standalone 5G NR, in which 5G radio shares the existing Long-Term Evolution (LTE) radio and core network. A Release 15 standalone implementation specification was completed in July 2018.

Release 16 focuses on the full spectrum of 5G use cases – beyond enhanced mobile broadband (eMBB) – to include ultra-reliable low latency communication (URLLC), unlicensed spectrum (5G NR-U), new spectrum sharing (5G NR-SS), vehicle communications for autonomous driving (5G NR C-V2X), and more. New interfaces between the radio access network (RAN) and 5G core are being developed to support the expanded use cases.

Among the most ambitious countries in the race to 5G are the U.S., South Korea, Japan, and China, with the U.S. leading the early rollouts. In October 2018, Verizon launched residential fixed 5G services in four cities using 28 GHz millimeter wave (MMW) spectrum, with a fifth market (Panama City, Florida) planned for 2019. In addition, AT&T has announced 17 U.S. cities for initial mobile 5G NR launch using MMW-based "puck" hotspots to reach consumers. AT&T service turnup began in late 2018 and will continue throughout 2019. T-Mobile has announced that it will have hundreds of U.S. cities ready for its low-band 600 MHz 5G services by the first half of 2019, including six of the 10 biggest U.S. markets.

The race to early 5G deployments is encouraging, but the bulk of commercial activity will wait until after full standardization with Release 16. If the industry stays on its current course, 2020 is looking to be a key year for big launches – in the U.S. and abroad.

5G IMPLEMENTATION OPTIONS

The diversity of 5G's use cases and deployment options sets it apart from previous mobile technology generations, but it also adds challenges and complexities that did not exist before. Many radio spectrum options are available under the 5G umbrella. Although not officially defined in the industry, Heavy Reading broadly categorizes these options as low-band (sub-2 GHz), mid-band (2 GHz-6 GHz), and high-band (anything above 6 GHz) spectrum.

Capacity and coverage ranges vary widely according to spectrum band used. For example, low-band spectrum provides the greatest geographic coverage but also delivers the lowest data capacity. In the U.S., T-Mobile has announced plans for nationwide coverage using its existing low-band 600 MHz spectrum, but data rates will be limited.

Mid-spectrum ranges are suitable for metro coverage areas, and channel sizes in the 100 MHz range allow operators to increase data rates beyond 4G. Operators can also use carrier aggregation for even higher data rates, though single carrier is the ideal. Many countries have assigned spectrum in the 3.3 GHz to 3.8 GHz range. There is also a growing trend toward using 3.8 GHz to 4.2 GHz spectrum. Several countries, including China and Japan, plan to use spectrum in the 4.5 GHz to 5 GHz range. In the U.S., less mid-range spectrum is available for use, but unlicensed Citizens Broadband Radio Service (CBRS) is one example.

The highest data rates will be delivered in high-band spectrum, with particular interest globally in spectrum above 24 GHz, also known as the MMW bands. For bands in the MMW range, channel sizes range from 50 MHz up to 400 MHz, thus providing maximum 5G data rates without the added complexities of carrier aggregation. The tradeoff is that the coverage range is limited and achieving wider coverage requires additional cell sites (densification). High-band spectrum is planned for dense urban areas and is most closely associated with small cells for densification. Verizon launched initial 5G fixed wireless access services in the fall of 2018 using 28 GHz spectrum and AT&T has plans for MMW-based mobile 5G hotspots throughout 2019. Promised download speeds are as high as 1 Gbit/s.

Many 5G network architecture decisions – including transport network decisions – must follow from the 5G spectrum used. In the low bands, capacity is scarce and maximizing spectral efficiency will be paramount (i.e., operators need to squeeze as much data as possible from each megahertz). The spectral efficiency requirement is driving some operators toward centralized RAN architectures that provide tight coordination between macro and street cells. Centralized RAN leads to a set of decisions for handling new fronthaul and, to a lesser extent, midhaul segments.

High-band spectrum options will have limited coverage, thus driving cell site densification to meet urban coverage requirements. Densification will come primarily in the form of small cells deployed on light poles, on buildings, within buildings, and on new fixtures – all places that do not have existing telecom infrastructure (including fiber connectivity). Delivering coverage and performance as economically as possible will be key to success for operators.

Finally, different architectures will not be deployed exclusively or in isolation. While operators will start 5G from their points of strength (based on spectrum that is most readily available), mobile operators will ultimately need a range of bands to address all the 5G use cases that will emerge – enhanced MBB, URLLC, massive machine-type communications (mMTC), mobile access, and fixed wireless access.

BUILDING THE CONVERGED TRANSPORT INFRASTRUCTURE

5G services will be built around three major use cases, each with a distinct set of requirements for capacity, latency, reliability, and other factors. For example, high capacity is a hallmark of the eMBB use case, which promises downlinks rates up to 1 Gbit/s. mMTC describes Internet of Things (IoT) applications in which data rates to individual sensors can be very low (measured in kilobits per second), but numbers of connected devices scale into the billions. URLLC describes mission-critical and extreme precision applications in which end-to-end latency may be 1 ms, jitter less than 1 μs, and reliability measured to six nines.

Operators are tasked not just to meet the performance requirements of each diverse use case, but also to meet those requirements as economically as possible. Although building three separate transport networks – each tuned to different use case requirements – is technically suitable, doing so is not economical. Operators must share network infrastructure wherever possible, including sharing the transport network. Thus, it is extremely likely that at a single 5G cell site, the device will need to support fronthaul, midhaul, and backhaul connectivity. Meeting diverse use case requirements over a shared infrastructure is also the primary driver for network slicing (discussed in the "Key Packet Technologies for 5G" section).

Significantly, converged packet transport extends beyond the 5G use cases themselves. High-band spectrum requires cell site densification, and it is likely that new remote radio heads will be deployed very close to urban enterprise customers. Operators have an opportunity to serve enterprise services over the same access network. Doing so spreads buildout costs and helps justify deployment for the site. In the early days, much of 5G will be speculative, so it may be difficult to justify builds based on 5G consumer uptake alone.

As a final point on convergence, there is a role for the transport network in connecting the compute and storage clouds that will be deployed at edge locations, such as central offices or aggregation points. Many operators are also thinking strategically about developing a universal aggregation network that can serve many access types and customer segments. (The term "multi-access edge compute" [MEC] reflects the understanding that edge compute will serve applications beyond mobility.) In this sense, the RAN transport decision is strategically important beyond 5G itself. The transport network should provide meshed connectivity to allow the operator the flexibility to deploy multiple models and, critically, to adopt new deployment models over time.

Figure 1 illustrates a converged packet transport network supporting 5G, fixed access, and MEC.

Figure 1: Converged Transport Network Diagram

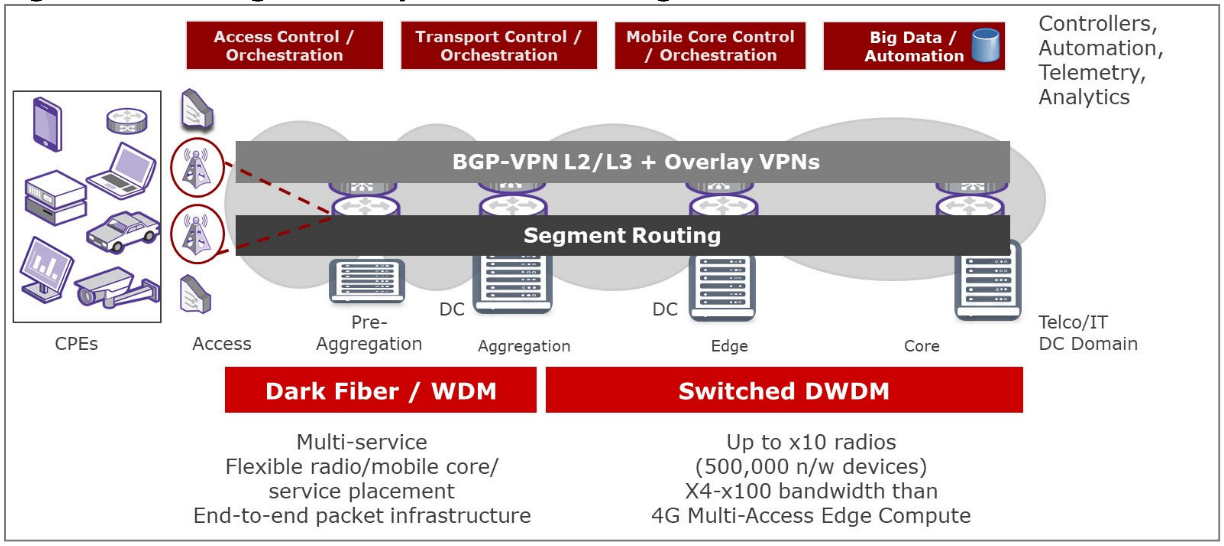

Source: Cisco and Heavy Reading, 2018

KEY PACKET TECHNOLOGIES FOR 5G

Industry debate continues over the best options for 5G at the physical layer – whether transmission should be over dark fibers, active wavelength division multiplexing (WDM), next-generation passive optical network (PON), or even wireless. There is strong consensus, however, on the use of packet technologies at the protocol layers, particularly in backhaul and midhaul segments. There are also potential benefits of using packet switching in fronthaul. This section addresses packet technology innovations that will be key to success in 5G transport.

Segment Routing

Segment routing is a variation of source routing. This is a routing technique in which the sending router specifies the route that the packet will take through the network, rather than the path being chosen based on the packet's destination only. In segment routing, a node steers a packet through an ordered list of instructions called "segments." A segment can represent any instruction, whether based on topology or service. As with other source routing techniques, the full instructions for the path through the network are embedded in the packet header, and this is applied at the source node. In segment routing, these are Multiprotocol Label Switching (MPLS) headers on Internet Protocol Version 4 (IPv4) packets today (and directly on IPv6 packets in the future).

Segment routing was officially introduced in the Internet Engineering Task Force (IETF) in 2013. Under the standard name Source Packet Routing in Networking, or SPRING, there are roughly 50 IETF segment routing drafts as of now.

Figure 2 provides a simplified representation of segment routing through a MPLS network. In the diagram, segment information is embedded in the segment router header at the ingress router (R1), defining an explicit path to the destination router (R6) via intermediate routers R2 and R3.

Figure 2: Segment Routing through an MPLS Network

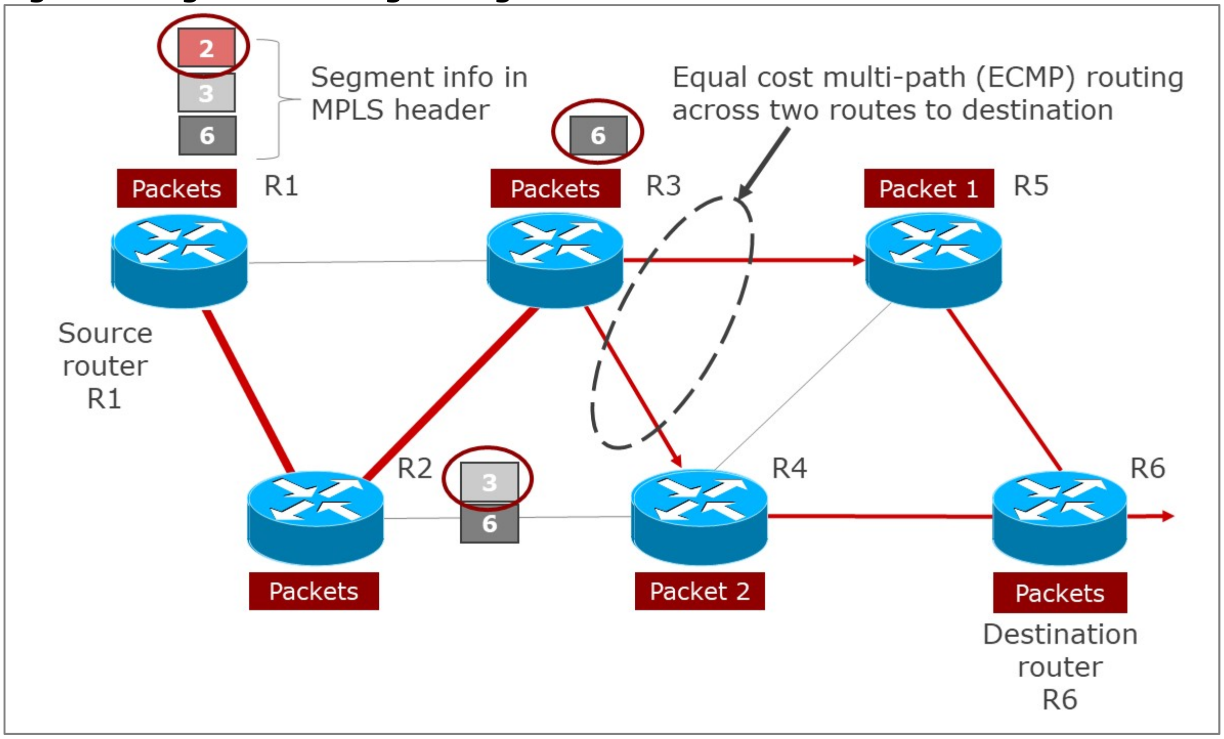

Source: Heavy Reading, 2018

Benefits for 5G

Segment routing was not developed specifically for 5G. Rather, it was developed to scale next-generation packet networks more efficiently and simply than legacy MPLS techniques and to provide a path to centralized, software-defined networking (SDN) control. As the biggest packet networks in the world (by node counts), mobile networks are well-suited to the benefits of segment routing. With massive investments beginning now, 5G is the ideal mobile generation for introducing segment routing to metro networks, including backhaul.

Benefits specific to 5G include the following:

- **Network slicing:** Network slicing is needed in 5G to segment the physical network into multiple virtual networks that can support different service types and different radio access technologies. Segment routing can be used to create different paths and priorities for different types of traffic based on the priorities assigned to those traffic flows (e.g., SDN slices). Through a combination of technologies, segment routing can build end-to-end slices in a simpler, more scalable and controllable way compared to MPLS Traffic Engineering (MPLS TE). As a final point, there are two types of network slicing as defined by the IETF. With hard slicing, resources including routers, control planes, and links are physically partitioned. With soft slicing, although slices are partitioned and cannot interfere with one another, resources are shared. Segment routing can be used to implement either hard slicing or soft slicing, as needed.

- **Network scale:** Segment routing helps scale networks for 5G as capacities increase compared to 4G and as node counts increase (due to small cells and densification).
- **Simplification:** Segment routing simplifies packet networks and network engineering by removing protocols from the network (e.g., Label Distribution Protocol [LDP] and, with IPv6, MPLS itself), removing state from the network, and adding SDN for centralized inter-domain control and global network views. Simplification is important for design and operations. It contributes to scale as well since networks that are easier to design and run also scale quicker and more efficiently.
- **Automation:** Segment routing simplifies automation in the network. By reducing the number of touch points of devices down to only the edge, segment routing reduces the complications of automation of services management.

Network-Based Timing and Synchronization

As operators plan for 5G, network timing and synchronization has become a hot topic. 5G creates a new set of demands and challenges that must be addressed.

Some countries (most notably North America) rely on global positioning systems (GPS) as their timing source for existing mobile networks, requiring an antenna at each cell site for communication with satellites. As operators deploy small cells in dense urban environments, keeping a reliable signal becomes a problem due to obstructions (such as buildings) as well as a wide range of interferences, including commodity GPS jamming devices that can take down cell sites. While North American operators will likely continue to use GPS as their primary reference source, interest is increasing in adding a network-based reference source as a backup in case of failures. In fact, U.S. operators drove the development of the ITU Telecommunication Standardization Sector (ITU-T) G.8275.2 Telecom Profile specifically for backing up GPS synchronization with Precision Time Protocol (PTP).

The other major challenge for GPS as a timing source in dense small cell environments is cost, including capital expenditures and operational expenses. Antennas and associated cabling and amplifiers are costly and require specialist skills to deploy. Access to roof space to deploy antennas may carry additional costs. As a result, operators may prefer network-based timing as their sole reference source in some cases – even if GPS is provided to macro sites.

Outside North America, most operators already use network-based timing and synchronization. Nonetheless, they too face synchronization challenges in moving from 4G to 5G. The primary reason is that most of the network-based synchronization deployed to date is frequency synchronization, using either the ITU-T SyncE standard or the ITU-T PTP Profile for Frequency, named G.8265.1.

Frequency synchronization was suitable for 3G networks as well as early 4G networks, but 5G brings new requirements for phase synchronization that are not covered by frequency-only implementations. These requirements are driven by new 5G services, technologies, and architectures that will require higher accuracy time synchronization requirements. Examples include tight radio coordination requirements due to carrier aggregation and coordinated multipoint and the use of spectrum techniques such as time division duplex (TDD), which needs phase synchronization to operate.

The discussion above relates to backhaul and midhaul networks. Packet-based fronthaul brings in a new set of timing and synchronization requirements due to additional restrictions imposed by separating the remote unit (RU) from the distribution unit (DU) processing. Packet-based fronthaul is of interest for some operators, though others will prefer to keep fronthaul networks circuit-based for relative simplicity. Regardless, more work needs to be done in packet fronthaul. The IEEE Time Sensitive Networking (TSN) for Fronthaul standards work seeks to address timing in packet-based fronthaul networks.

CONCLUSION

5G promises a communications revolution for consumers and businesses, with countless new application possibilities enabled by massively scalable networks, unprecedented flexibility, ultra-low latency, and ultra-high reliability. But leading-edge operators understand that the wireline transport network plays a fundamental role in delivering future 5G services, and Heavy Reading believes that an end-to-end converged packet transport network will be critical to doing so economically. End-to-end converged packet transport can provide the following:

- A unified network where all use cases can be addressed from a single cell site and a single device.
- The ability to share infrastructure costs across consumer and business and mobile and fixed applications.
- A unified meshed connectivity for compute and storage clouds deployed at edge locations (i.e., MEC).

New innovations are required to make packet technologies suitable for 5G midhaul and backhaul. Most notably, IETF-standardized segment routing delivers scale, simplicity, and SDN-based automation. And network-based timing and synchronization will be needed for economics and reliability, either alone or in addition to GPS-based timing. Finally, beyond midhaul/backhaul, interest is growing in packet-switched fronthaul, though additional standards work is required to meet 5G performance demands.