

DPI AND 5G:

NETWORK VISIBILITY AND REAL-TIME APPLICATION AWARENESS

ROHDE & SCHWARZ



TABLE OF CONTENT

1.	Introduction	3
2.	DPI-empowered application awareness in 5G	4
2.1	5G service classes – networks within the network	6
2.2	5G network slicing	9
2.3	5G V-RAN and localized traffic management.....	10
2.4	Programmability and virtualization	11
2.5	5G network edge.....	12
3.	DPI as a service	13
4.	DPI use cases	14
5.	Conclusion	15

1. INTRODUCTION

5G presents a huge leap in mobile broadband capacity, speed and latency, and brings the promise of innovative and exciting new use cases. Part of the excitement comes from the convergence of new technologies such as big data, machine learning and artificial intelligence. These new technologies, combined with the gigabit speeds, enhanced capacity and ultra-low latency promised by 5G, are set to massively redefine cities, industries and our day-to-day living.

Ericsson Mobility Report June 2019* predicts the number of 5G mobile broadband subscriptions to grow from 10 million subscriptions in 2019 to a massive 1.9 billion by 2024, accounting for 20 % of all mobile broadband subscriptions. The report also expects 4.1 billion cellular IoT connections within that timeline. A total of 45 percent of these connections will use 5G and other technologies which use the 5G spectrum, such as NB-IoT and Cat-M.

This whitepaper looks at the network- and service-level enhancements brought about by 5G and determines the need for application awareness at various points throughout the network. Moreover, it assesses how application awareness is delivered by deep packet inspection (DPI),

an advanced technology for real-time classification of IP traffic. With application awareness, networks can distinguish different types of traffic they carry. It is a capability that lends intelligence to today's networks, allowing effective application monitoring and dynamic traffic management for enhanced network performance. In a 5G world, application awareness will be even more important. This paper examines in detail DPI's granular classification capability and how this helps network operators to manage an array of new applications and services introduced by 5G. These include much-touted use cases such as evolved mobile broadband (eMBB), ultra-reliable low latency communications (URLLC) and massive machine type communications (mMTC), which will see applications like remote surgeries, autonomous driving and industrial automation being delivered on mobile networks.

More importantly, the paper uncovers the need for application awareness and detailed traffic analytics across various network services to support 5G's new features, such as network slicing and edge computing. It looks at how DPI-empowered real-time traffic classification information feeds into the network to support the implementation of these advanced features. This contains the rise of programmable and virtualized networks in 5G. Enhancing both capabilities, application awareness allows the creation of agile and responsive networks, delivering enhanced performance, security and quality of experience for end users.

*Source: <https://www.ericsson.com/en/mobility-report/reports/june-2019>

2. DPI-EMPOWERED APPLICATION AWARENESS IN 5G

About 5G

5G boasts new technologies such as MIMO and beam-forming, which open up new frequency bands in the sub-6 Gigahertz (GHz) and the mmWave range. Current standards classify 5G networks and devices into 2 categories: Standalone and non-standalone 5G. While standalone (SA) 5G implements 5G end-to-end, non-standalone 5G uses the existing LTE infrastructure for the control plane.

At the forefront of 5G is the combination of an advanced architecture and service classification that corresponds to a high-performing network operating in an optimized, virtualized and programmable environment. These characteristics, however, need to be supported by the breadth and depth of traffic analytics that enable traffic management decisions to be implemented effectively across a wide range of applications.

This inherently gives rise to the need for an application-aware network. This capability is critical in 5G, where applications vary tremendously in their needs for speed, latency, security and scalability. While data from a UHD video application requires more bandwidth, traffic from a remote surgery application requires ultra-low latencies and data from a moving train requires highest mobility. Accordingly, overall network visibility and real-time application awareness is key in 5G for real-time traffic management decisions and flexible allocation of network resources, corresponding to application-specific key performance indicators (KPIs).

What is deep packet inspection?

DPI is a traffic recognition method that classifies IP traffic in real time. DPI identifies protocols, applications and application attributes (for example VoIP in a chat app). Commercial DPI engines leverage continuously updated traffic libraries established through analysis of traffic patterns and application behavior, and current trends.

In addition to IP traffic classification, some DPI engines extract protocol- and application-based metadata, providing insight into user behavior and application usage. The metadata extraction functionality enables

detailed understanding of network transactions and behavior. These insights can be used for a wide array of applications, such as customer experience management, network planning, policy management, network security and many others. Examples of metadata that can be extracted from IP traffic include the following:

Metadata category	Example metadata
Traffic volume	Per user, per protocol, per application, per flow, per direction
Service detection	Differentiation between for example Skype audio and video calls
Quality of service	Jitter, throughput, latency, roundtrip time, ramp-up time, packet loss, retransmissions
Security and data leakage	File up- and downloads, entropy-based DNS tunneling detection
Client information	HTTP/QUIC user agents, operating system

On the security front, DPI-powered tools or systems check anomalies and suspicious traffic patterns against updated libraries of latest cyber threats, such as DDoS attacks, malware and ransomware. With advanced techniques such as flow tracking, byte-pattern matching and behavioral and statistical analysis, DPI can extend application and protocol classification even to encrypted traffic.

DPI is deployed either as a network service or a network capability (deployed within a network service such as a firewall) to facilitate end-to-end traffic management and also to read and analyse traffic patterns. Providing network visibility and real-time application awareness, DPI facilitates analyzing and managing IP traffic and securing IP networks in real time.

Network analytics for network providers

Communication service providers (CSP) identify eroding profitability as a top challenge in building the next-generation Internet. More personalized, intelligent and intuitive service offerings will be key to add value to basic connectivity. To generate such offerings, CSPs depend on fine-grained insights into user behavior and demands.

Apart from application-specific information, DPI provides insights on users, including their subscription plans, device types and location, as well as information on overall network performance and network security. Besides influencing bandwidth and traffic management decisions, these insights enable decisions relating to the following:

- ▶ Policy and charging control – for example, providing guaranteed QoS for specific social-media applications
- ▶ Network optimization – for example, further improving handover mechanisms based on VoLTE/Vo5G accessibility or VoLTE/Vo5G drop call ratio data
- ▶ Monetization – for example, delivering real-time offers to customers who are running out of data and are forced to terminate their session; and cross-selling video plans to customers based on their video content consumption patterns
- ▶ Service assurance and service monitoring – assessing current QoS levels against service level agreements (SLAs) on a service, application and user basis

These insights, especially when gathered over a longer term, provide the analytics required for strategic decision-making. This is how CSPs gain the business intelligence to tackle the 5G challenge and sustain overall profitability. This includes but is not restricted to the following:

- ▶ Enhancements in customer experience management – for example, caching and compressing video content on the most consumed video applications to improve customer experience on these applications
- ▶ Improvements to 5G mobile plans and services – for example, introducing app-based plans that correspond to the most consumed mobile applications
- ▶ Predictive maintenance – for example, using line quality parameters (e.g. transmission control protocol KPIs such as timing, entropy or throughput) to identify potential bottlenecks in the network
- ▶ Security and threat management – for example, the implementation of NGFWs with an intrusion prevention system to arrest attacks by IoT botnets known to launch DDoS attacks and spread ransomware

Traffic management decisions enabled by application classification

Traffic management decisions are based on a set of rules built into the network to manage the delivery of IP traffic. These decisions are then implemented by routing traffic through the right network tools and security appliances.

Real-time application awareness coupled with network analytics enable these decisions to be implemented at application level, where each application is sent to a different set of tools and appliances depending on the classification provided by DPI. Network resources such as bandwidth, computing power and storage are then automatically allocated (optimized) to support these decisions.

Real-time application awareness enables traffic management decisions to be implemented at application level, where each application is sent to a different set of tools and appliances.

While DPI has been providing IP application awareness in the core of 4G networks and their predecessors for some time, under 5G DPI's capability extends from the core to across the entire network, supporting 5G's new and enhanced features, namely:

- ▶ 5G service classes defined by key use cases
- ▶ 5G network slicing
- ▶ 5G vRan
- ▶ 5G network virtualization and programmability
- ▶ 5G network edge

The following sub-chapters discuss these features and how DPI enables network visibility and real-time application awareness in each of them.

2.1 5G SERVICE CLASSES – NETWORKS WITHIN THE NETWORK

One of the key attributes of the new 5G architecture is the delivery of distinct classes of service categorized by network performance attributes such as capacity (coverage and speed), latency, mobility and scalability. The current classification by 3GPP outlines the three major classes as follows.

Enhanced mobile broadband (eMBB)

eMBB is characterized by higher throughput, higher coverage and mobility. eMBB is often touted as the successor of LTE, providing enhancements on all fronts to the current mobile broadband experience. eMBB boasts gigabit speeds of 1Gbps and mobility of 500km/hour. Within 5G, eMBB will be providing connections to smart 5G devices and delivering rich content such as ultra-high definition (UHD) videos and immersive content such as AR/VR. All early 5G deployments are essentially deployments of eMBB.

Additionally, 5G fixed wireless access (FWA) extends from the 4G FWA to provide broadband connectivity to households and business premises in underserved, often sparsely populated areas. Deployed via an external receiver, 5G FWA promises speeds of 1 Gbps to enable similar use cases as eMBB.

Ultra-reliable low latency communication (URLLC)

URLLC is a new class of communication defined by 5G's ultra-reliable low-latency feature. URLLC delivers a latency of 5 ms to 1 ms, enabling the creation of 'Tactile Internet' with edge computing and haptics delivering mission-critical applications such as remote surgeries, real-time interactive games (for example cloud gaming) and V2V communications for autonomous driving. Early deployments of URLLC are expected to begin in 2021.

Massive machine type communication (mMTC)

mMTC is the 5G variant of LTE-m and NB-IoT, boasting the ability to simultaneously connect to massive numbers of connected things, often requiring consistent connection but much less bandwidth. Examples include smart meters, signage, CCTV cameras and IoT sensors. mMTC is hence characterized by scalability, and powers the IoT and M2M communications on 5G networks, enabling up to 1 million devices to be connected simultaneously in a space of one square kilometer. Just like URLLC, the deployment of mMTC is expected to begin in 2021.

Within this multi-service architecture, network visibility and real-time application awareness provided by DPI plays the following roles:

1. Classification of traffic by service classes

Identification and classification of IP traffic in real-time using advanced filtering and classification technology such as DPI, enables traffic routing to be executed seamlessly through the different service classes. This is important in the 5G architecture which manages traffic from millions of smart devices, IoT nodes and customer premises equipment (CPE).

2. Optimization of each service class

With DPI, optimization of each service class is enabled by continuous reallocation of resources between them and implementation of traffic management policies that correspond to the applications delivered within each class.

This is due to the fact that while broad traffic management policies can be assigned to eMBB, URLLC and mMTC use cases, application-based policies are much more optimized and effective for the delivery of application specific KPIs.

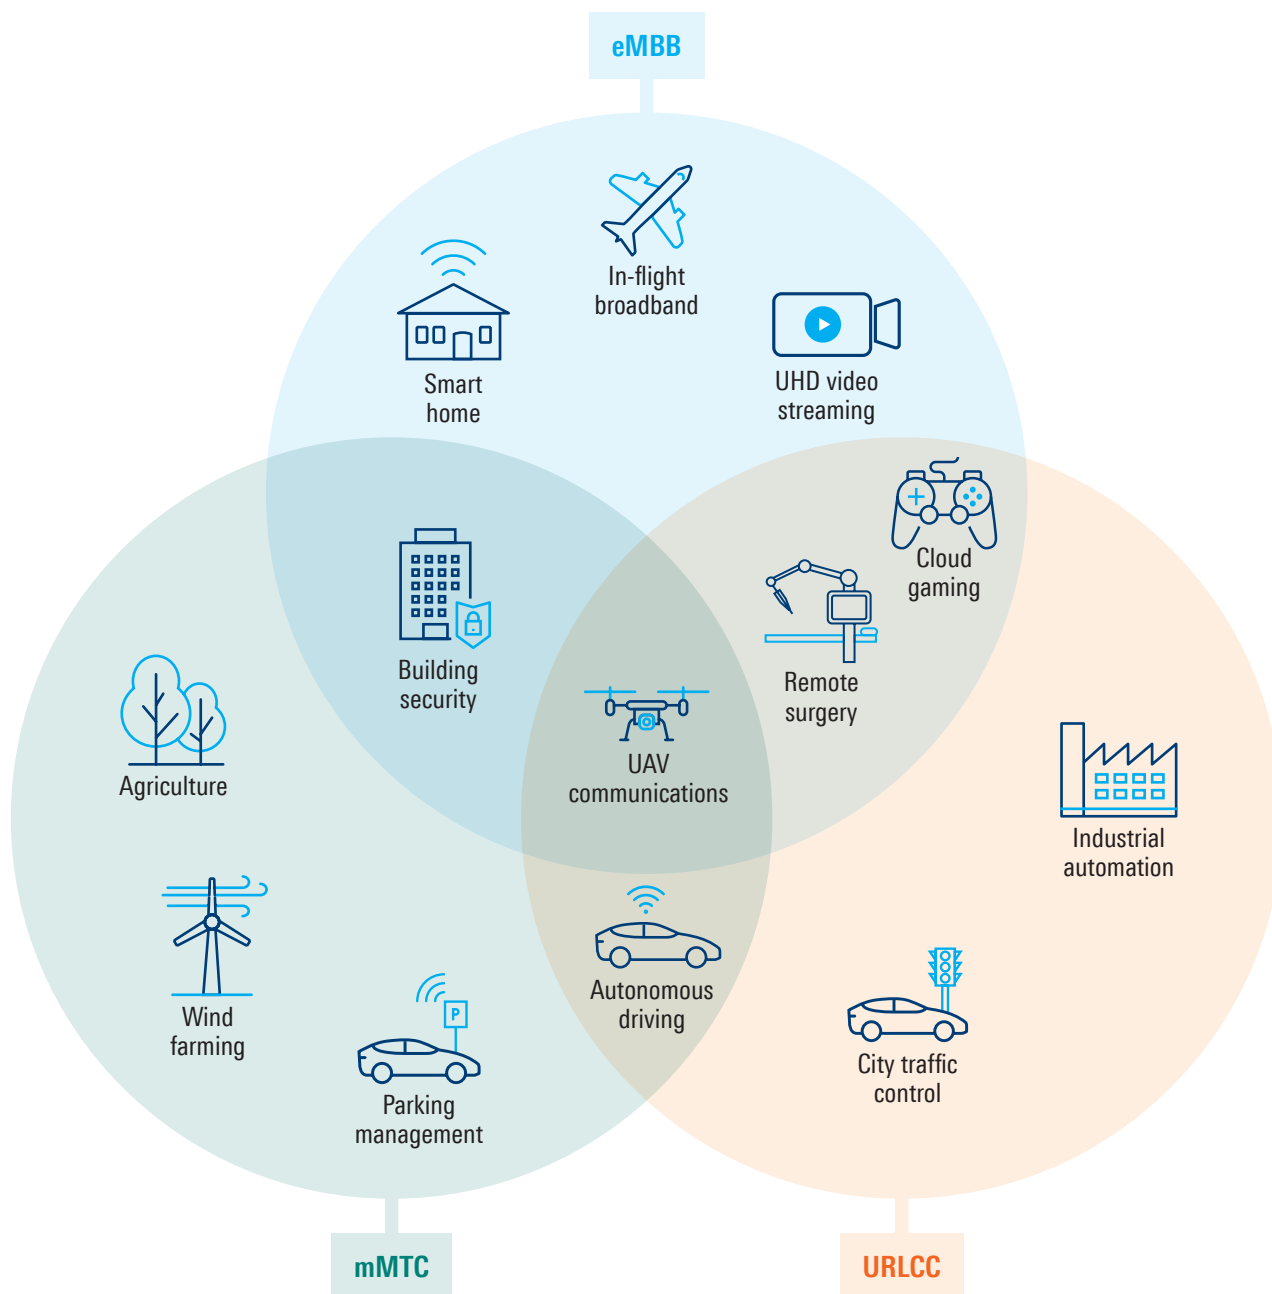
In general, eMBB aims for high throughput, speeds and bandwidth. URLLC aims for low latency, low enough to power Tactile Internet applications. mMTC aims at connecting densely deployed IoT nodes, often transmitting low amounts of data consistently over very long periods of time. However, an application may require performance KPIs that involve more than a service class. For example, a large number of connected security cameras in the city center may be relaying high-definition images in real time during a high-profile event. This requires both the bandwidth to cater for transmission of video content, and the 'always on' connection that connects to hundreds of cameras in a single locality, combining the attributes of eMBB and mMTC.

KPIs may also vary within a service class. For example in eMBB, some applications such as corporate emails, online banking sites and mobile wallets require filtering through additional network firewalls, while applications such as YouTube can be delivered directly, bypassing these ele-

ments. With application awareness, security KPIs can be adjusted individually on an application level.

The graph below presents some 5G applications with regard to required service classes.

COMPLEX APPLICATIONS EMPOWERED BY 5G SERVICE CLASSES



Network visibility and real-time application awareness hence enable each service class (eMBB, URLLC and mMTC) to be optimized according to overall network conditions as well as the applications and the application attributes (including security), ensuring highest performance and QoS on each service class and each application.

The following are the common traffic management decisions that correspond to specific applications and their attributes with reference to 5G use cases.

Application/traffic	Classification	Traffic management policies
Remote surgery	Ultra-low latency	Guaranteed bandwidth, highest QoS class, low latency with low jitter and low packet loss
UHD video streaming	High bandwidth, low latency	Burstable bandwidth, medium QoS class, traffic load balancing, video optimization, video caching
Cloud gaming	Ultra-low latency	Guaranteed bandwidth, highest QoS class, edge traffic steering
Heavy file transfer	High bandwidth	Burstable bandwidth, best effort QoS class, data offloading, content compression, content caching
Augmented reality	High bandwidth, ultra-low latency	Guaranteed bandwidth, highest QoS class, edge traffic steering
Smart metering	High scalability, low throughput	Medium QoS class, no bandwidth guarantee
Industrial automation	Low latency	Low latency path, no bandwidth guarantee, medium QoS class
Video call on a high speed train	High mobility	Burstable bandwidth, video optimization
Mobile banking	High security	Secure path, real-time threat management, SSL inspection
Corporate email	High security	Anti-phishing, anti-spam and malware detection

2.2 5G NETWORK SLICING

One of the key features in 5G is network slicing, available only on standalone 5G networks. Network slicing technology enables operators to virtually ‘slice’ the network (RAN, transport and core) by optimizing resources and network topology to create logical networks or partitions that correspond to different market segments, radio access, plans and customer types (for example, enterprises and third-party service providers). Network slicing leverages technologies such as network functions virtualization (NFV) and software defined networking (SDN) to create layers of different virtual networks on a physical infrastructure. Network slicing enables each ‘slice’ to be provisioned dynamically in terms of resources and network services, scaling a slice up or down according to traffic needs. This partitioning renders the 5G network a highly optimized one, where resources (such as computing, storage and bandwidth) and network services (such as firewalls, session border controllers and intrusion protection) are invoked only when necessary. Again, network visibility and real-time application awareness plays a vital role.

Slice development

For example, there may be a gradual growth in traffic from IoT end nodes coming from smart city applications. That would signal the need for a standalone IoT slice that is scalable and can handle more connections and data.

Resource reallocation between slices

DPI-powered classification of applications, in combination with wider network analytics and traffic management policies, can establish the real-time load on a slice. This information enables resources to be moved from one slice to another, especially during peak traffic hours.


Slicing within a slice

DPI can further optimize traffic by classifying the traffic within a slice into sub-service groups defined by a hierarchy of speeds, latency and network service requirements. Combined with SDN-powered dynamic service chaining, this will allow allocating additional network resources and services to a portion of that slice’s traffic without affecting the rest.

Multi-slice, single end-service

With network slicing, some enterprises will require more than a single slice to serve their data needs. For example, a city parking management company may require its massive IoT terminals to gather small bits of data on parking space availability. Simultaneously, it would require sufficient bandwidth in the ‘mobile broadband’ slice for its security cameras to relay parking bay images to app users. DPI enables the operator to use the most appropriate slice for each traffic type, offering their customers optimized 5G connection at a lower cost.

5G NETWORK SLICING



2.3 5G V-RAN AND LOCALIZED TRAFFIC MANAGEMENT

The quest to further enhance and optimize the RAN saw operators moving to the centralized RAN architecture, and later to the virtualized RAN (V-RAN) architecture. In 5G, the V-RAN is expected to move a significant portion of traffic management decisions from the core to the RAN.

In a V-RAN architecture, pooled baseband units (BBUs) at the edge of the network are deployed as virtual functions via virtual machines (VMs) on standard commercial servers in a cloud data center. The virtualized BBUs are front-hauled to the remote radio heads (RRUs) and back-hauled to the 5G core.

The pooling of BBUs in a data center allows flexible allocation of computing and storage resources based on the traffic handled by each base station to run RAN control functions and service delivery optimization. This in turn calls for traffic management decisions, which require real-time network insights at each end node.

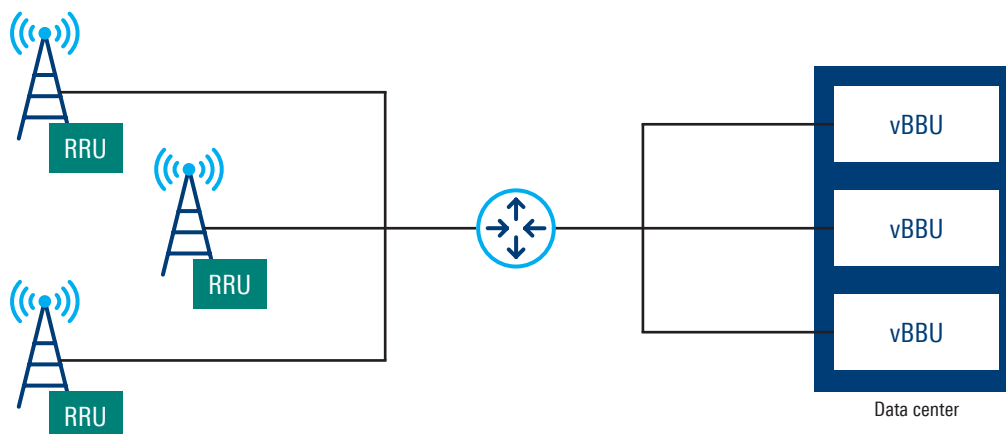
Supporting these needs within the V-RAN is the DPI functionality, which inspects, detects and classifies traffic at each BBU, allowing better management of network capacity and better end user experience in a given locality. This is particularly important during peak traffic or congestion times, where prioritization by application (for example, mission critical applications) is necessary, but required only in the congestion area (hotspots).

The 5G V-RAN will also require similar intelligence to manage network densification in 5G, resultant of deploying a higher number of macro sites, small cells and in-building wireless as well as the use of MIMO and sector splitting technologies.

Network densification in 5G will require more handover decisions as users move from cell to cell. For example, a user watching a YouTube video on a smartphone at the bus stop who needs to be moved to a closer small cell from a larger macro cell, or else be connected to the operator's WiFi hotspots serving the town's center. Matching protocol and application metadata with information on network capacity and performance allows smarter hand-over decisions in real time.

Generally, convergence between multiple access networks (WiFi, fixed networks, 3G, 4G and 5G) becomes more imminent in 5G, driven partly by the densification via small cells, and the centralization of base stations enabled by the centralized RAN and V-RAN architectures. This will see major improvements to current convergence technologies, namely LAA (license assisted access) and LWA (LTE-WLAN aggregation) especially for deployments in high-density areas and in-building deployments. The result is here, too, a need for more enhanced real-time, granular data on local traffic flowing through multiple networks for seamless movement of traffic from one access node to another.

5G V-RAN



2.4 PROGRAMMABILITY AND VIRTUALIZATION

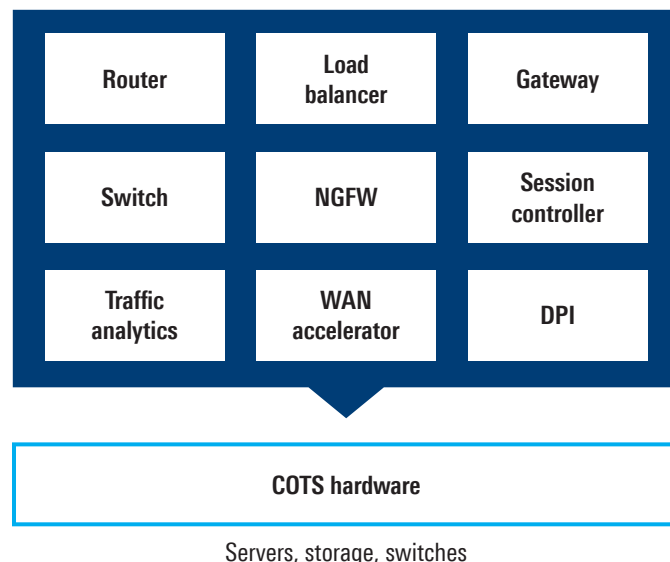
SDN in 5G aims to create programmable networks and network hierarchies that correspond to different end services and applications by separating the forwarding (data) plane from the control plane, with an SDN controller managing and controlling the network. Programmability is crucial for 5G, primarily due to the traffic load and the diverse range of use cases it will handle, which requires intelligent networking for a more efficient use of network resources. At the same time, network hierarchies will allow content to be cached and stored locally, with the SDN controller managing the orchestration between different network devices. Both processes lead to highly optimized networking.

The widespread adoption of NFV will make 5G networks both highly programmable and highly virtualized. In an SDN-NFV environment, the SDN controller implements service chaining where virtual network functions (VNF), such as encryption software, firewalls, NAT, routers and load balancers are linked together to create a service chain optimized to each application, service and customer.

Within virtualized networks, DPI will be playing a vital role. Network visibility and real-time application awareness provide the intelligence required by the SDN controller and the VNFs (for example, routers, switches and gateways) for automated, programmable decision making. This translates to an application-aware network that intelligently

NETWORK FUNCTIONS VIRTUALIZATION (NFV)

NFV involves decoupling network functions from their proprietary hardware, resulting in deploying them as virtual machines on generic high volume switches, storage and servers known as commercial off-the-shelf hardware (COTs). This improves network agility, scalability and customizability while significantly reducing costs, as deploying new network services merely involves installing or de-installing and upgrading software, instead of deploying completely new hardware. Altogether, NFV allows operators to continuously optimize their networks.



classifies and manages IP traffic end-to-end. Whether DPI is collocated with the SDN controller for centralized control of traffic management (with intelligence being relayed to network services via APIs) or collocated alongside network services (closer to the user/source), its role will be crucial in feeding the virtualized network with real-time information for traffic management decisions.

The use of DPI in network virtualization extends further to support the shift from automated to cognitive network management. The shift is enabled by machine learning that leverages historical data from the network, again a function enabled by the advanced analytics provided by DPI. Machine learning enables the network to dynamically self-optimize, self-heal, self-configure, self-protect and self-relocate, matching current network resources with shifting variables related to traffic, users, content, location and performance SLAs based on historical correlations between them.

In addition to these layers of intelligence, DPI is expected to play another important role in virtualized 5G networks, namely to support the deployment of microservices. Microservices enable networks to become more efficient, agile and responsive. The sharp increase in microservices deployment in 5G is expected to promote the use of containerization via tools such as Docker. However, containers are highly susceptible to cyberattacks which is why DPI plays a key role here to detect threats, validate application access, identify sensitive data and minimize the vulnerable area.

2.5 5G NETWORK EDGE

The development of multi-access edge computing (MEC) is expected to create a powerful new network edge in 5G.

The MEC will be powering 5G's service-based architecture (SBA), where user planes are deployed at the edge, retaining only the control planes in the 5G core. Complete control and user plane separation (CUPS) will see the user plane function (UPF) undertake functions that involve packet routing and forwarding, including packet inspection and QoS handling with the MEC steering traffic to edge applications. This architecture will enable edge applications such as cloud gaming, V2V communications for autonomous driving, face/vehicle recognition, remote surgeries and AR/VR applications to be delivered from the edge itself without navigating the core, enabling ultra-low latencies.

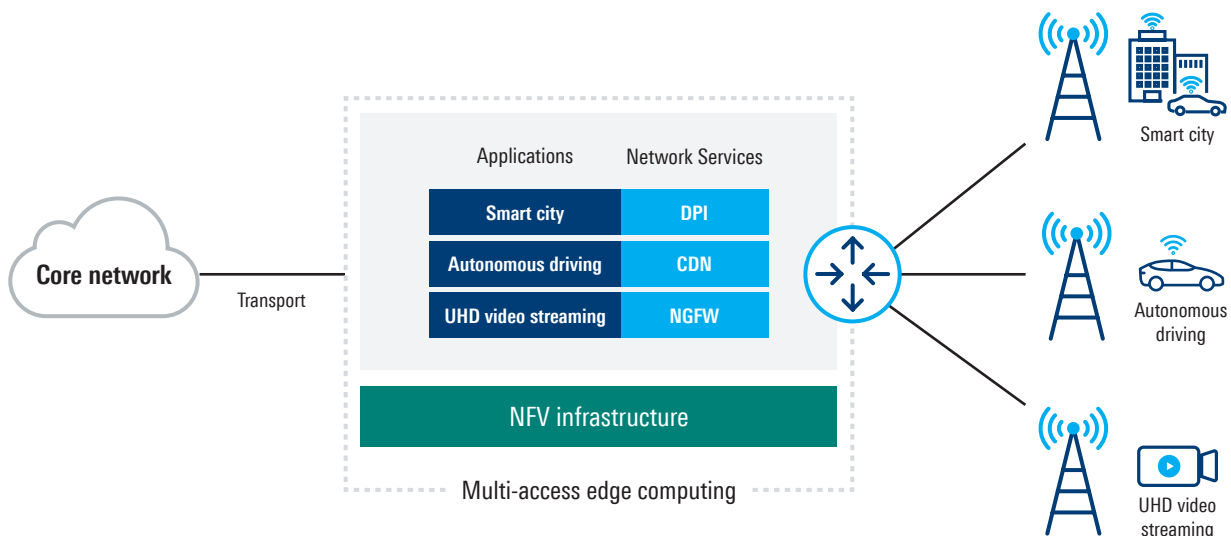
The 5G network edge, via MEC, will also optimize the delivery of rich content on 5G networks. MEC performs caching, compression and video optimization at the edge, reducing the resources consumed in the core network. UPF and MEC will essentially require real-time analytics on the network (RAN, transport and core), as well as user analytics (plan type, device type) and application-specific information (for example, latency and type of content delivered) on traffic passing through the edge. Edge analytics

EDGE COMPUTING

Edge computing moves computing processes from a centralized cloud to a distributed cloud with computing now being done as close as possible to the user. The edge computing node can reside in a regional data center, a smaller data center close to the user, an end node connecting to small cells in a stadium, the CPE or even in the user's end device itself. Edge computing is touted as a major evolution in 5G networks, due to the rise in applications that require processing and delivery from the edge, especially those which call for ultra-low latencies.

provided by DPI therefore become key in traffic routing and traffic optimization decisions. Likewise, they contribute to enforcing of security policies at the edge to ensure QoS on operator services and network optimization. They also help to sound out new monetization opportunities.

5G NETWORK EDGE



3. DPI AS A SERVICE

Network visibility and real-time application awareness thus become major tenets in the 5G network, creating the need for a DPI-as-a-service solution that is capable to perform inspection and classification of IP traffic in a highly virtualized and programmable environment, and which can be deployed at any point within the network.

OEM IP network analytics solutions by Rohde & Schwarz

Rohde & Schwarz is a global leader in the world of IP network analytics software. With ipoque, a Rohde & Schwarz company, the technology group leverages deep domain expertise to create customized software solutions that empower the communication industry to transform network data into intelligence. Customers include network equipment and software vendors, as well as service providers that enhance their network solutions with the market-leading DPI software R&S®PACE 2 or the technically advanced IP probe R&S®Net Sensor OEM.

R&S®PACE 2

Continuous evolution of web protocols and applications requires software libraries that are frequently updated and register the latest changes in IP traffic trends. Companies opting to license DPI software from a DPI specialist such as Rohde & Schwarz benefit not only from the expertise, but also from weekly updates of application signatures included in the libraries. Combined with continuous performance and reliability testing, this significantly increases traffic detection rates and traffic classification accuracy.

R&S®PACE 2 is a ready-to-use DPI software library that enables network software and equipment vendors to reduce costs and risks associated with developing and maintaining this highly complex technology. The DPI engine classifies thousands of applications and protocols regardless of whether they use advanced obfuscation, port hopping techniques or encryption. Additionally, it provides content and metadata extraction. R&S®PACE 2 boasts the most efficient memory and CPU utilization in the industry, featuring the smallest processing footprint. The software requires only around 400 bytes per flow, while using very little processing power and no memory allocation during runtime.

R&S®Net Sensor OEM

R&S®Net Sensor OEM is the ideal choice for companies looking for a ready-to-use DPI network probe. The OEM-ready solution comes as a passive probing software based on the market-leading DPI engine R&S®PACE 2. It provides fast packet processing to reveal a clear picture of the entire network and its subscribers. The IP probe classifies both plain and encrypted IP traffic to offer detailed visibility. Additionally, it can correlate control and user plane to keep track of specific subscriber sessions and experiences in a network. R&S®Net Sensor OEM also provides aggregated information of the collected data. This includes information on applications and protocols by user, time, duration, frequency and usage.

Customizable to individual demands, additional modules can further enhance the solution: aggregation and correlation functions, a database and graphical user interface or integration of third-party products and solutions offer additional network insight and flexibility. APIs enable customers to extend R&S®Net Sensor OEM with their own modules.

OEM customer benefits:

- ▶ Customizable to individual requirements
- ▶ Focus on core competencies to become more efficient and profitable
- ▶ Speed up time-to-market by optimizing the development schedule
- ▶ Reduce and optimize development costs by outsourcing DPI and IP probing
- ▶ Maximize return on investment (ROI)

4. DPI USE CASES

One of the key highlights of R&S®PACE 2 is the flexibility it provides in terms of how and where it is deployed and how classification results are presented to support different use cases. As a software library, R&S®PACE 2 can be deployed as a network capability or as on-demand resource via APIs and other interfaces.

This enables R&S®PACE 2 to cater to specific real-time network visibility and application awareness needs from vari-

ous layers, parts and functionalities within the network, in line with the 5G network features and 5G services discussed in sections 2.1 to 2.5 of this whitepaper.

In its software form, R&S®PACE 2 is a perfect fit for SDN and NFV. In 5G, it provides the scalability and flexibility required to cater for a highly agile network environment, providing DPI functionalities across various use cases, the most popular of which are summarized below.

Application	Use case details
Network and traffic management (QoS/QoE)	Provides traffic management vendors with real-time insights on network performance by user, application and services, along with trend analytics. Provides protocol, application, application attributes and application metadata such as delay, packet latency, jitter and call completion on applications such as carrier VoIP. By integrating R&S®PACE 2, vendors can keep up with the dynamic changes in protocols and applications, ensuring a high rate of detection for traffic management.
Policy and charging control	Provides policy control and charging software vendors with the capability to define bandwidth guarantees, priorities and limits, offer fine-grained QoS for an additional fee and deliver real-time charging and billing support. R&S®PACE 2 software offers a high detection rate (> 95 %) and accurate application identification for policy and billing purposes.
Network security (firewalls, IPS/IDS, SIEM, UTM)	Provides next generation firewalls (NGFW) vendors with application-aware capabilities to accurately distinguish applications (for instance, Hulu vs. Salesforce) and apply policies based on business rules. By integrating R&S®PACE 2, security vendors quickly gain accurate application visibility and control, becoming better able to manage security threats and prevent network attacks.
Network and subscriber analytics	Provides analytics vendors with accurate identification of application usage by user groups, usage and behavior patterns, end devices, geography and other rich data, including heavy users and popular applications.
Mobile data offload	Provides operators with the capability to distinguish traffic by application and device, enabling operators to offload data intelligently and thus maintain high QoE on services such as VoIP and video streaming. R&S®PACE 2 software enables vendors to reliably detect applications and implement smarter offload strategies based on total flow visibility. Intel has integrated R&S®PACE 2 into their "Smart Pipe" server, used in small cells as a mobile Internet access gateway.
WAN optimization	Provides WAN vendors with real-time protocol and application visibility that enables them to enhance the performance of their WANs. R&S®PACE 2 can identify thousands of applications for every IP flow in real time and can handle complex applications that aggregate several multimedia content types (video, images, VoIP, etc.) from different servers.
SDN/NFV environments	R&S®PACE 2 can migrate from being embedded in many network appliances to becoming a shared function hosted on standard servers. R&S®PACE 2 has no external dependency, works on standard servers and OS and can be used in all environments, regardless of whether it is a physical environment, a virtualized one or even an SDN architecture.

5. CONCLUSION

Over months to come, we will start seeing major milestones in the deployment of 5G services. eMBB rollouts will intensify, and commercial deployments of URLLC and mMTC will begin. This will bring about a surge in the different types of applications and services delivered on mobile networks, which in turn will demand dynamic traffic management policies that are highly responsive to each traffic type. 5G's ability to deliver the speeds and latencies it promises hinges on network visibility and real-time application awareness. This intelligence enables:

- ▶ Different service classes defined by eMBB, URLLC and mMTC
- ▶ Implementing network-level enhancements such as:
 - ▶ Network slicing
 - ▶ Edge computing
 - ▶ Virtualization of RAN
 - ▶ Network programmability and virtualization

Without network visibility and real-time application awareness, 5G networks will not be able to implement differentiated policies resulting in all applications being routed, prioritized, filtered and delivered in the same way. This would lead to overconsumption of network resources, poor performance of business-critical and latency-sensitive applications and loss of monetization opportunities.

DPI thus becomes key in 5G – as it provides real-time network visibility and real-time application awareness from the core to the edge and across all layers and nodes within the network. DPI integrates seamlessly across the network, feeding crucial insights that enable intelligent traffic management. This helps network operators to maintain high levels of network performance, security and customer experience across all 5G applications and services. Likewise, it ensures high network efficiency and, in the long run, a lower cost of ownership for the network.

ipoque

ipoque, a Rohde&Schwarz company, is a global leader of network analytics software for the communications industry. We leverage our deep domain expertise to create software solutions that empower customers to transform data into intelligence. As a subsidiary of Rohde&Schwarz, we take advantage of potential synergies, yet act independently.

Rohde & Schwarz

The Rohde&Schwarz technology group develops, produces and markets innovative information and communications technology products for professional users. Rohde&Schwarz focuses on test and measurement, broadcast and media, cybersecurity, secure communications and monitoring and network testing, areas that address many different industry and government-sector market segments. Founded more than 80 years ago, the independent company has an extensive sales and service network in more than 70 countries.

Rohde & Schwarz GmbH & Co. KG

www.rohde-schwarz.com

ipoque GmbH

Augustusplatz 9 | 04109 Leipzig

Info: + 49 (0)341 59403 0

E-Mail: info.ipoque@rohde-schwarz.com

www.ipoque.com

R&S® is a registered trademark of Rohde&Schwarz GmbH&Co. KG

Trade names are trademarks of the owners

PD 3608.1717.52 | Version 01.00 | October 2019

White paper | DPI and 5G

Data without tolerance limits is not binding | Subject to change

© 2019 Rohde&Schwarz GmbH&Co. KG | 81671 Munich, Germany

© 2019 ipoque GmbH | 04109 Leipzig, Germany



3608171752