Managing the Future of Cellular: What 5G Means for the Radio Access Network (RAN)

White Paper

Radio Access Networks (RANs), which link end-user equipment and the core network, are an essential part of new 5G network architectures. Having a high-level understanding of what 5G entails can help make it easier to manage expectations when developing a 5G RAN. This paper explains the current status of 5G standards and rollout, summarizes the three categories of use cases that 5G RANs will need to support, and highlights the considerations for 5G RAN design. We also explain how the Arm Neoverse supports the optimization of 5G RAN architectures.

The Factors That Influence 5G RANs

Ask a dozen people what 5G is, and you're likely to get a dozen very different answers. That's because 5G is such a multi-faceted subject, intended to address a complex technology that needs to evolve in many different areas.

Any way you look at it, 5G is a massive undertaking. Designed to accommodate every type of electronic device that will eventually access the network – today, tomorrow, and years from now – 5G addresses everything from tiny sensors and smartphones to industrial robots and autonomous vehicles. With 5G, tens of thousands of devices can access individual cells, and edge devices will perform increasingly complex processing tasks that involve large amounts of data. The network will need to be resilient enough to handle all this, yet flexible and responsive to manage traffic while maintaining the extremely low latencies required for real-time operations.

The Radio Access Network (RAN), which manages traffic and transmits data between end-user equipment and the core network, is an essential part of this new way of operating. To place the RAN into a 5G context, and anticipate what next-generation RANs will be asked to do, it can help to step back and take a high-level look at 5G – the standards involved, the use cases addressed, and the network architectures required.

"Introducing 5G functionality in stages provides a stable platform for development while ongoing refinements are made."

That's what we aim to do in this paper. We summarize where we are in the phased rollout of 5G standards, define the use cases that 5G standards are designed to enable, summarize the network architecture that will support those use cases, and describe the role of RANs in that architecture. We also introduce Arm's approach to 5G, which includes technologies that enable the RANs to perform in this new 5G environment.

5G Has Arrived. And It Will Keep Arriving For Years To Come.

The 3GPP (3rd Generation Partnership Project), the global industry organization developing the technical specifications for 5G, is itself a complex entity. Made up of representatives from regional standards bodies, as well as people from roughly 400 companies in 39 countries, the 3GPP identifies present and future use cases for cellular, defines the technical features needed to deploy those use cases, and manages the details that ensure interoperability on a global scale.

The 3GPP model uses planned "Releases" to introduce functionality in stages. This lets network providers and device manufacturers move ahead with development, based on a stable platform, while some of the details are still being worked out.

2019 was a big year for 5G, since it saw the main drop of Release 15, the first Release to directly address 5G operation. But Release 15 is just the beginning:

3GPP Timeline

2019 = Release 15
2020 = Release 16
2021 = Release 17

- Release 15, which defines much of the basic operation of 5G, was completed in June 2019. It defines the 5G New Radio, or 5G NR, along with the 5G core network, or 5G system, and several important use cases, including the Internet of Things (IoT).
- Release 16, which is the second phase of 5G definitions, is expected to be completed by the end of 2020. Along with further definitions of 5G NR and core network operation, Release 16 will define services for latency-sensitive devices in applications like factory automation, autonomous driving, and remote surgery (more on that in the next section).
- Release 17, which will include various incremental enhancements and support features, is still under discussion, but is scheduled to be completed by the end of 2021. As of this writing, Release 17 topics under consideration include nearly 30 possible refinements to the standalone 5G architecture and more than 15 RAN features.

The 3GPP in Numbers

- Representation: North America (22%), European Union (36%), Asia (42%)
- 9 regional standards bodies
- ~400 companies from 39 countries
- 50,000 delegate days per year
- 40,000 documents per year
- 1,200 specifications per Release
- New Releases every 18 to 24 months

What's Included in Release 15 and Release 16

Release 15

- **NR**
- **The 5G System – Phase 1**
- **Massive MTC and Internet of Things (IoT)**
- **Vehicle-to-Everything Communications (V2x) Phase 2**
- **Mission Critical (MC) interworking with legacy systems**
- **WLAN and unlicensed spectrum use**
- **Slicing – logical end-2-end networks**
- **API Exposure – 3rd party access to 5G services**
- **Service Based Architecture (SBA)**
- **Further LTE improvements**
- **Mobile Communication System for Railways (FRMCS)**

https://www.3gpp.org/release-15

Release 16

- **The 5G System – Phase 2**
- **V2x Phase 3:** Platooning, extended sensors, automated driving, remote driving
- **Industrial IoT**
- **Ultra-Reliable and Low Latency Communication (URLLC) enhancements**
- **NR-based access to unlicensed spectrum**
- **5G Efficiency:** Interference Mitigation, SON, eMIMO, Location and positioning, Power Consumption, eDual Connectivity, Device capabilities exchange, Mobility enhancements
- **Enhancements for Common API Framework for 3GPP Northbound APIs (eCAPIF)**
- **FRMCS Phase 2**

https://www.3gpp.org/release-16

What Use Cases Will 5G RANs Need To Support?

As one of their first steps in defining 5G, the International Telecommunication Union - Telecommunications (ITU-T) standards identified the ways that cellular is and will be used by consumers, enterprises, and industry. 3GPP then set about implementing the standards to support this. As part of what the 3GPP called their "SMARTER" project (Study on New Services and Markets Technology Enablers), they identified high-level use cases for cellular, now and in future, and identified the features and functionality 5G would need to deliver to enable those use cases.

In addition to a category called fixed broadband, they came up with three categories of mobile use cases: Massive Machine Type Communications (mMTC), Enhanced Mobile Broadband (eMBB), and Ultra-Reliable, Low-Latency Communications (URLLC).

Now, admittedly, these category names aren't particularly catchy and don't exactly roll off the tongue, but they've become industry-standard terms so it's best to keep using them. Here's a closer look:

- **Massive Machine Type Communications (mMTC)**

  This category covers machine-to-machine interactions happening on a very large scale and includes battery-operated devices operating in the IoT. By and large, these devices need relatively low latency and highly reliable connectivity while they operate on efficient power capacities. The challenge is to provide scalability and consistent connectivity for the billions of IoT devices that communicate relatively infrequently and in short bursts. Wide coverage and deep indoor penetration are important, as are low cost and energy efficiency.

+ **Enhanced Mobile Broadband (eMBB)**
  If mMTC mostly addresses how machines use cellular, eMBB mostly addresses how people use cellular. It includes activities like 8K video streaming, immersive Augmented Reality/ Virtual Reality (AR/VR), connected transportation infotainment, and enterprises supporting mobile broadband connectivity. The key requirements in this category are ultra-high spectral efficiency, extreme data transmission rates, and ultra-low interrupt time. All these things are addressed by the 5G NR defined in Release 15. As infrastructure support for the 5G NR begins to expand, these use cases will become more widely available. This category can be considered a mix of evolutionary and revolutionary, since laptops using cellular for connectivity isn't exactly new, while immersive AR/VR and other data-intensive applications haven't really been possible with previous generations of cellular.

+ **Ultra-Reliable, Low-Latency Communications (URLLC)**
  This is, for many, the truly revolutionary aspect of 5G, since it delivers a level of performance not yet seen in real-world applications. It's essentially the IoT on steroids. It covers things like intelligent transportation, including vehicles that can navigate complex road situations and avoid collisions by cooperating with each other, and use cases associated with the fourth industrial revolution, including time-critical factory automation. It also includes remote healthcare, which includes devices that measure vital signs and either automatically or semi-automatically respond as needed, as well as remote treatment, including surgeries performed in ambulances, during disaster situations, or in remote areas while using real-time guidance from an offsite doctor. In all these situations, the connection needs to be exceptionally stable and needs to operate with end-to-end latency rates on the order of a millisecond or less. The features needed to support URLLC are still being defined and will be included in Releases 16 and 17. In other words, URLLC is the future of 5G, even if that future is only a few years away.

The various features added to each Release are intended to address different aspects of these three categories. The specific use cases that are already in use today or close to arriving are addressed in earlier Releases, while use cases that are farther in the future are addressed in later Releases. It's all part of the ongoing evolution of 5G.

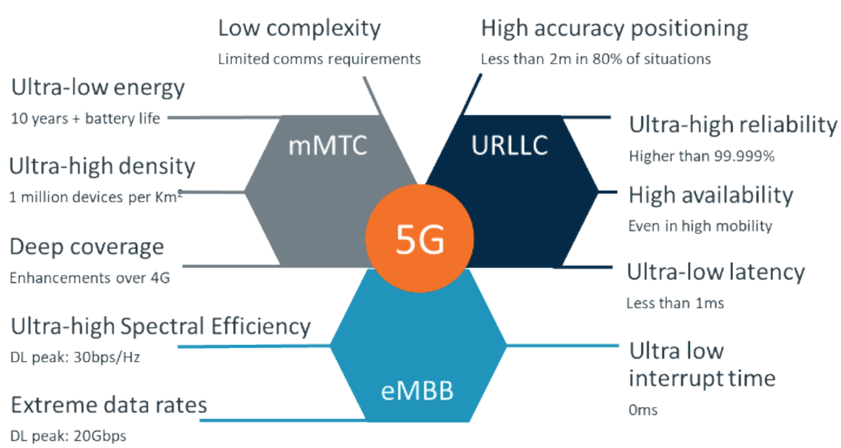

NGMN Association

3GPP Focus Areas

- Release 15 = mMTC, eMBB
- Release 16 = URLLC with enhancements for mMTC, eMBB
- Release 17 = Ongoing enhancements for mMTC, eMBB, and URLLC

Sample 3GPP Use Cases

mMTC	eMBB	URLLC
Lighting & road sign control Smart waste management Asset tracking Structure and environmental monitoring	Immersive AR/VR Advanced gaming 8k video streaming Enterprise broadband connectivity Connected transportation infotainment	Industrial automation Intelligent transportation Remote healthcare

"5G is both evolution and revolution. Release 15 brings upgrades to IoT, enterprise, and consumer applications, while Releases 16 and 17 will enable entirely new applications."

How Will 5G RANs Support 5G Use Cases?

RANs will continue to be an essential part of the 5G network, but they will take on new functionality to support the broader 5G cellular architecture, which covers everything from the client device to the cloud data center. Multiple pieces of equipment will be required to support this connectivity.

Today's IoT is a good example of how the network architecture needs to evolve to support 5G use cases. With current setups, data from a connected device is typically transferred, using backhaul equipment, to a hyperscaler data center, for application processing, and then transferred back to the device. This proves to be inefficient and costly to the IoT provider while delivering sub-par performance to the IoT device user.

As we approach a trillion connected devices, the steady increase in data traveling through the network forces a change. If, for example, today's cellular network carries a total of 150 Exabytes per month of data. Assuming, in the future, the network was to evolve to support a billion connected image sensors over 5G, this would, on its own, increase the bandwidth required to service these sensors and the resultant image data to 400 Exabytes. Simply adding capacity to the backhaul channel won't keep pace with this need and would be prohibitively expensive for network operators to keep adding additional capacity to their backhaul network.

This is why there will be a move to push computing power towards the edge of the network. Having processing capability closer to the client device connected through 5G means there's less need to transfer data to backend data centers. Intelligence to make decisions on the images captured will be provisioned closer to the edge of the radio network.

5G RANs will enable these newer use cases at the edge, by supporting the necessary features, including higher performance levels whilst meeting strict power requirements. At the same time, as URLLC capabilities come online, with their strict requirements for exceptionally fast response times, 5G RANs will filter workloads and react to edge-based applications – with latency rates around 1 or 2 milliseconds.

Round-Trip Latency Requirements for Emerging Applications

- Robotics/Medical — 2.2 ms
- V2X Augmented Apps — 10 ms
- Smart City, IoT Control — 15 to 20 ms

The transition to 5G requires existing building blocks to perform new functions. As shown in the diagram, the RAN needs to support legacy devices running 4G as well as new devices that use the 5G NR radio, and the network-access mechanisms must support local decision-making.

Edge processing is used to filter and react, using content caches in smaller, more localized data centers referred to as edge clouds. The edge cloud acts as an intermediary, forwarding data as needed to the core network and hyperscaler clouds so it can be analyzed and stored.

The core network, managed by the MNO, provides supervision of edge-based network components, including the RAN.

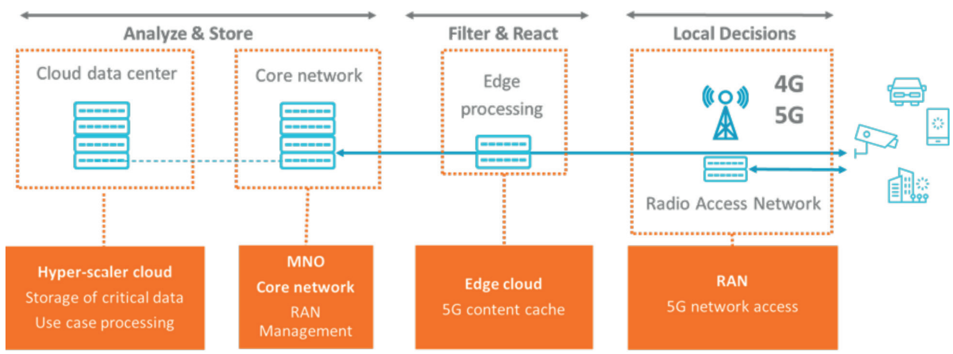

Building Blocks for 5G Network Architectures

Network Numbers

Global Mobile Traffic	2018 = 28 Exabytes/Month	2025 = 130 Exabytes/Month
Unique Mobile Broadband Subscribers	2019 = 8 Billion	2025 = 9 Billion
Mobile Data Traffic	2017 = 46% Video	2025: 74% video
IoT Connectivity	2019 = 7.5 Billion	2025 = 25 Billion

Sources: Ericsson Mobility Report, ReThink, Cisco

"With 5G, RANs need to take on new tasks, and they need to perform those tasks with remarkable efficiency."

Considerations influencing the evolution of 5G

Because 5G is relatively new and covers such an extensive list of features and operating requirements, it can be hard to know what to expect from the development effort for those features. At Arm, our work developing and deploying 5G technology has uncovered five issues that are likely to need attention as hardware and software teams begin defining the next generation of 5G RANs.

The 5G Learning Curve

- New algorithms
- New architectures
- New spectrum
- New suppliers
- New standards

1. **Data Processing** – To achieve higher spectrum efficiency and meet end-to-end latency requirements, a 5G RAN will need to perform more complex algorithms for data processing. When considering the requirements for these algorithms, it's important to find the right balance between hardware and software tasks, so the system meets its goals for performance, power, and cost.
2. **Deployment Scenarios** – The specific use cases to be supported by a given RAN have a strong influence on the system as a whole. Since each use case (mMTC, eMBB, URLLC) has its own unique characteristics, one size probably won't fit all. Deciding how to divide network functionality across different pieces of equipment, to support a given set of use cases, can impact the RAN design.
3. **Radio and Spectrum** – 5G uses more of the spectrum, with devices operating in low bands (below 1 GHz), mid bands (between 1 GHz and 2.6 GHz or between 3.5 GHz and 8 GHz), and high bands (between 24 GHz and 40 GHz). Each band has its own set of requirements for edge performance, capacity, speed, and latency. As new spectrum assets become available, these various requirements will need to be addressed by the RAN system.
4. **Supply Chain and Ecosystem** – 5G is disrupting the supply chain in several ways. There are initiatives that aim to reduce vendor dependence and there is growing availability of proprietary and open software platforms. The level of infrastructure support also varies from region to region. OEMs may need to reassess and revise their ecosystem partnerships.
5. **Emerging Standards** – In addition to the 3GPP, there are independent industry organizations, such as the Telecom Infra Project (TIP) and the Open RAN Alliance (O-RAN), that are working on aspects of 5G operation and deployment. There is a growing convergence towards O-RAN alliance as the key industry body that drives interface specifications.

"Working with an experienced partner makes it easier to navigate the newness of 5G."

Arm Neoverse: The Foundation for 5G Cloud-To-Edge Infrastructure

Building on our industry-leading success in markets like mobile and smart IoT, Arm has developed Neoverse, a comprehensive portfolio for the 5G infrastructure that provides end-to-end coverage from the cloud to the edge. As 5G evolves and expands, our infrastructure business is developing, and in many instances inventing, the technologies needed for a 5G world that includes a trillion edge devices.

A Modern Architecture

The Arm Neoverse considers the current and long-term trends influencing 5G development and addresses the need for a strong 5G ecosystem:

- **High Performance, Secure IP, and Architectures** – We deliver the highest levels of performance for each power/area budget and support competitive, single-thread performance with exceptional power efficiency. Security is a crucial aspect of 5G operation, which is why Arm Neoverse platforms are built around a very specific security model supported by our own secure platforms and compliance programs.

- **Diverse Solutions and Ecosystem** – We give designers the freedom to integrate custom accelerators and to supplement operation with open-source software capabilities that are pre-integrated into the Arm environment. We enable different approaches, so Proof Of Concept (POC) designs can accurately verify and validate capabilities. We also work with a broad range of silicon and software vendors to push innovation, with features like multi-chip module (MCM) capability, FPGA integration, and chip-to-chip (C2C) interconnects.
- **Scalable from Hyperscale to the Edge** – We offer the performance per socket needed to deploy servers on standardized, Arm-based blades, and the performance per thread for throughput processing in the RAN. What's more, all this IP is scalable across all the points in between.

"Arm Neoverse solutions are uniquely designed for unprecedented levels of performance, security, and scalability."

An Expanding Ecosystem

Beyond the tools, software, and solutions that we supply to support Neoverse IP, we're working with a strong ecosystem that enables every kind of design for 5G deployment. That means everything from modeling and emulation to closer integration with leading software vendors.

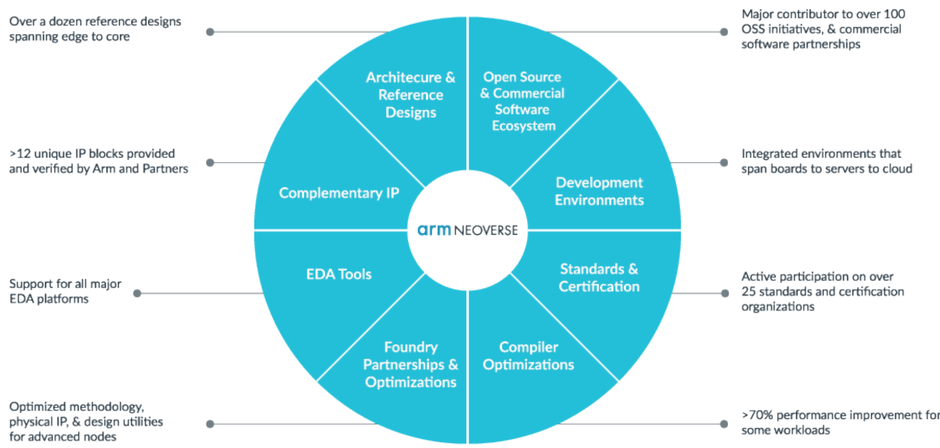

Realizing the full potential of 5G requires new thinking and new system architectures. This is especially true for 5G RANs, because they will need to manage increasingly complex processing tasks that involve an unprecedented number of cellular-connected devices. The Arm Neoverse is specifically designed to help developers meet this challenge. To learn more, visit www.arm.com.