Tackling the network challenges for XR

Understanding, developing and delivering an E2E XR connectivity solution

White paper

This paper provides an overview of today's network capabilities when serving different types of extended reality (XR) traffic. It outlines the network deployment strategies, the technologies, and the key enablers for providing even a greater XR experience and higher network capacity in the future.

It further describes how Nokia enables XR end-to-end (E2E) through collaboration with key global players and plays a driving role in key standardization bodies affecting XR. It outlines Nokia's unique offering, providing products, services and technologies to enable E2E XR solutions.

Contents

Introduction

Extended reality (XR) is an umbrella term for virtual reality (VR) and augmented reality (AR). The global XR market is estimated to reach US$357 billion by 2030 [1]. The rise in demand for immersive entertainment, industrial training, and remote assistance by utilizing augmented, virtual, and mixed reality is gaining traction, paving the way for immense business opportunities in the metaverse.

While the XR market is projected to grow significantly, the global metaverse market is expected to reach approximately $426.9 billion by 2027 with a CAGR of 47.2% [2]. XR will be a crucial enabler for many use cases in the industrial, enterprise, and consumer metaverses [3]. Its devices will be essential for connecting the virtual and physical worlds.

Figure 1. Metaverse categories

Industrial

- Large range of use cases of full physical/ digital fusion from any industrial sector (manufacturing, supply chain and logistics, power and utilities etc)
- Virtual facility optimization
- Maintenance, robotics etc

Enterprise

- Digital and collaborative 3D design
- Augmented worker
- Immersive team collaboration
- Training and simulations

Consumer

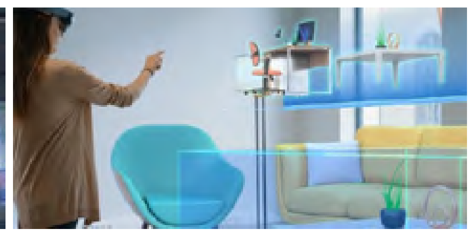

- Gaming & entertainment
- Social interactions
- Shopping
- Immersive Social Virtual Reality
- Smart City augmentation

Through XR, every aspect of our lives will change, bringing many benefits to the consumers and industries that adopt it [4]. The opportunities for integrating digital information into our physical world are endless, ranging from using artificial intelligence (AI) and machine learning (ML) to leveraging the internet of things (IoT) to enable physical-digital fusion and human augmentation [5].

Like in the early days of smartphones, VR headsets and AR glasses will become faster, lighter, and more affordable, hence, more consumer friendly. This may have heavy implications for the network; in the past, even the sudden popularity of faster smartphone models has led to unprepared network crashes [6].

Implementing consumer or business XR applications at scale will demand consistently high network performance. It will also require high computation power at or near the network edge while supporting various application types from multiple application providers. If communication service providers (CSPs) or enterprises want to provide a good user experience for XR applications and services, it's essential to enhance the network accordingly.

To reach economies of scale, standards will be critical for interoperability, thus supporting XR device and application proliferation, and ensuring user safety, security and privacy.

XR use cases and traffic model

From recent study results we have seen that there is a strong momentum for the adoption of industrial and enterprise XR that will drive commercial adoption [1]. Initial commercial deployments will allow users to try out the technology, experiment with different use cases, and innovate new solutions to be ready for the larger consumer market. As XR evolves, it will make a profound impact on all socio-economic sectors.

AR and VR are clearly different in the nature of the experience and are often deployed in different ways. VR is mainly consumed indoors while AR lends itself to both indoor and outdoor experiences. Nonetheless, they are based on similar technological components. We will therefore use the umbrella term XR in this paper and call out differences between AR and VR where needed.

All XR applications exhibit a combination of typical traffic patterns:

- Sustained high-throughput, low latency video streams of cloud rendered content to the user device
- Momentaneous high throughput downloads of 3D assets for local rendering when launching an application or moving to a new scene
- Low latency upstream activity and pose information to a cloud rendering process
- Low latency upstream and downstream activity and pose information streams in multi-user XR experiences
- Upstream and downstream voice communication in multi-user experiences.

AR applications, uniquely, can offload heavy world understanding tasks to the cloud; the associated upstream traffic can range from the upload of a few pictures to a continuous video stream.

Due to the different nature of VR and AR applications there are, however, differences in network usage. VR applications are mainly consumed indoors; thus they predominantly rely on fixed access in combination with Wi-Fi. AR applications can be consumed both indoors and outdoors, which adds cellular access to the usage patterns.

When downloading 3D assets and cloud rendered video streams, the expected downstream traffic for VR is substantially higher than for AR, because VR completely replaces the users' physical world environment with virtual background and foreground objects covering their complete field of view. In contrast, AR only positions a few virtual objects in the real-world environment, which in some use cases are rather lightweight (e.g., text and arrows for AR navigation).

The 2023 IEEE paper Extended Reality (XR) over 5G and 5G-Advanced New Radio: Standardization, Applications, and Trends [8] introduces the traffic model adopted in 3GPP as the basis for 3GPP Rel.18 capacity utilization and user equipment (UE) power saving enhancements. Although this traffic model can be considered generic and accurate enough for analysis and network performance evaluation, the real XR traffic requirements strongly depend on the specific use case as they have different characteristics and quality of service (QoS) requirements.

Figure 2 shows different latency and throughput requirements of some representative XR use cases and cloud gaming. Though cloud gaming is not necessarily an XR use case, it shares some of the VR and AR traffic patterns and, with its increasing popularity, can be considered as a precursor to the XR networking requirements.

Figure 2. Applications and network requirements

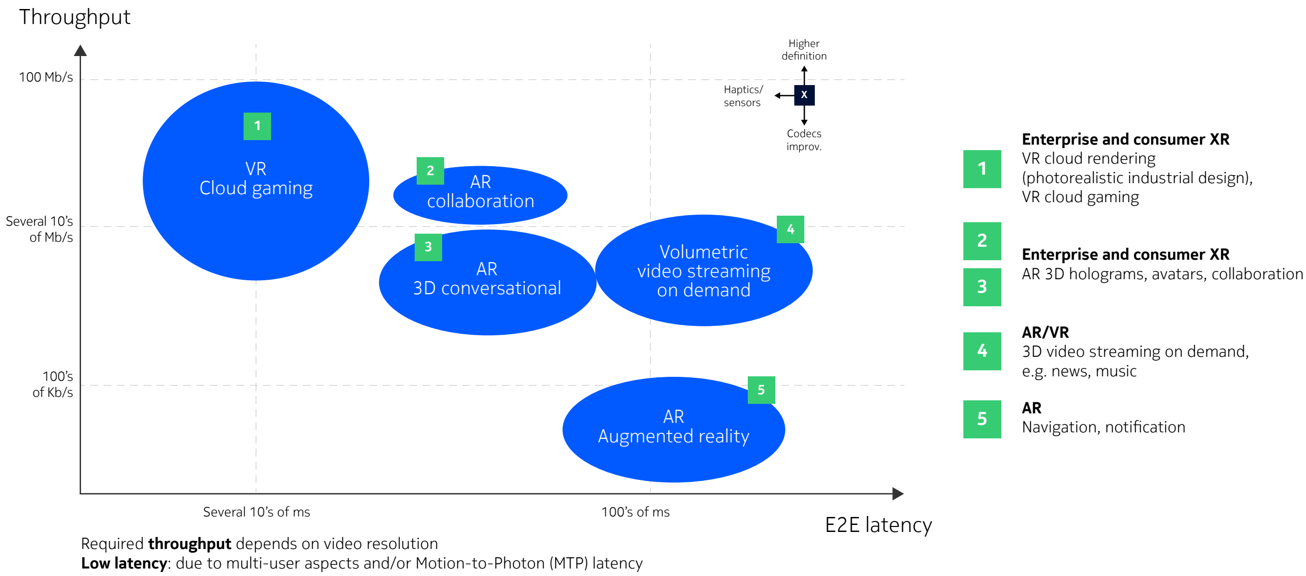

It should be noted that in addition to the XR traffic, the network will also have to deliver all other types of traffic such as voice calls and IP/Internet applications. This means XR traffic can't be strictly prioritized over other traffic but must be served in concert with them.

Meeting the highly demanding requirements of ubiquitous XR experiences will nonetheless remain a challenge in future networks (depending on the use cases). Examples include 30-45 Mbit/s downlink 4K video, tens of Mbit/s for uplink video capture streaming, and a delay budget of some tens of milliseconds for round trip time (RTT) for split rendering with pose correction.

XR E2E connectivity

Technologies to support XR use cases involve multiple domains and system components spanning applications, devices and network technologies along with optimization technologies for improving the end user experience. Although the key technologies for XR and Metaverse include content and tools, platforms and enablers, the focus of the paper is on infrastructure and devices and deals with the key technologies to provide XR end-to-end connectivity and achieve the best user experience. The logical E2E architecture is shown in Figure 3.

Figure 3. Network logical architecture

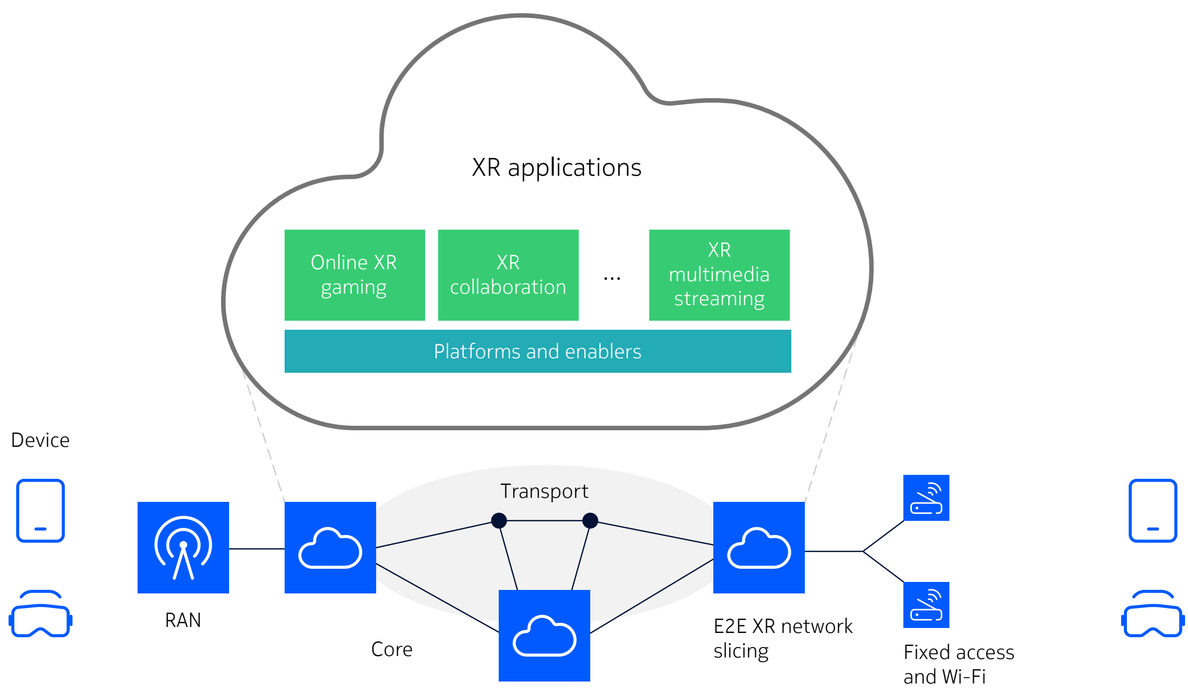

In this chapter, we describe the capabilities of current networks when carrying XR traffic. To evaluate current and future network requirements, it is important to understand XR devices and their future evolution, as well as network deployment scenarios.

Processing modes for XR devices

XR applications require heavy processing loads for the rendering of virtual 3D assets and for the geometric and semantic understanding of the user's environment. Split rendering over cellular networks is described in 3GPP TR 26.998, TR 26.928 and TS 26.565 and depicted in Figure 4.

Figure 4. 5G Split-rendering architecture

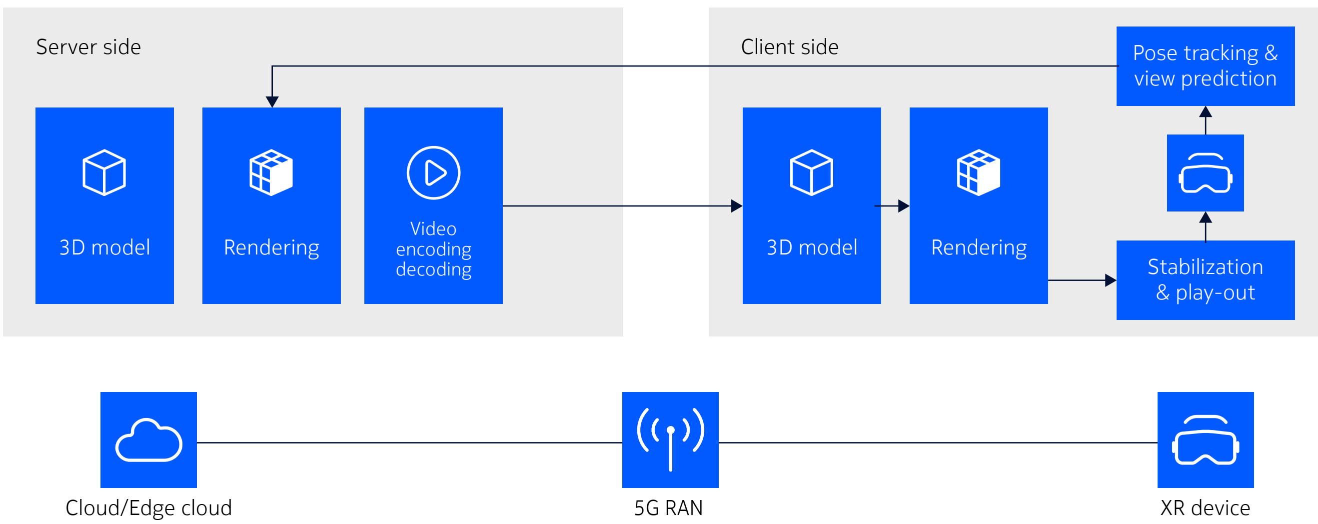

XR devices can operate in different modes, depending on where the processing takes place: local processing, cloud processing and split processing. These various operational modes for XR processing generate fundamentally different networking requirements.

Local processing is when all XR-related processing is executed on the head-worn XR device itself. Examples include standalone VR head-mounted devices (HMDs), which are currently the most popular consumer option, and standalone AR headsets, mainly used in enterprise settings. The obvious advantage is the simplicity of a single all-in-one device, but the solution suffers from several drawbacks: the inherent processing limitations of mobile chipsets, the limited battery life, heating issues and a bulky uncomfortable form factor.

A more complex implementation of local processing is a partial offload of XR processing from the head-worn XR device to another device in the immediate vicinity of the user:

- VR headsets are tethered (wired or wirelessly over Wi-Fi) to powerful PCs that provide high-quality rendering capabilities, however at a substantially higher total solution cost
- AR glasses are tethered (wired or wirelessly over Wi-Fi) to smartphones or dedicated portable compute packs, which makes for more lightweight and comfortable AR glasses, but is merely moving the mobile graphics processing unit (GPU), rendering and battery consumption challenges to a secondary portable device. Future AR glasses with higher resolution and wider field of view (FoV) will generate even higher processing loads.

Cloud processing moves the heaviest XR processing loads from the local processing device(s) to the cloud. The main use case is cloud rendering of high-quality XR experiences for standalone VR headsets and standalone/smartphone-tethered AR glasses. Cloud rendering uses GPUs in the cloud to render 3D assets in high quality and deliver the information as a video stream to the XR device. Current cloud gaming services use similar mechanisms to stream games to TV sets, laptops, smartphones and tablets.

Additional complex processing for semantic understanding of the user's environment (e.g., object recognition) that are beyond local processing and memory storage capabilities can also be moved to cloud resources.

Cloud rendering offers a set of clear advantages:

- Possibility to visualize photorealistic renderings of 3D assets and avatars
- Much lower processing needs on standalone headsets and support for local offload devices (eventually eradicating the need for additional local compute devices) and the associated battery life increases
- Device-agnostic application deployment
- Additional security because confidential and copyrighted 3D models are not delivered to unauthorized user devices.

Cloud rendering demands challenging networking support: 30 to 70 Mbps downstream video streams and latency below 30ms. Optimizations like viewport-based streaming and foveated streaming (using eye tracking in XR devices) can bring down throughput requirements. These will be needed to support future XR devices with higher resolutions of up to 8K per eye.

Split processing is a future, more advanced hybrid processing mode that splits XR processing loads between the XR device, additional local processing devices, and the cloud depending on device capabilities, device status (e.g., battery state and number of concurrently running applications), and the network status (e.g., latency from user to cloud processing, network load, radio conditions and cloud load status). The rendering and AI-based world understanding processes are split between the device and the cloud thus providing a tradeoff between network throughput, latency requirements and device power consumption. Additionally, in split processing, the network status and device capabilities can determine the best processing split, which may dynamically change with network and device conditions such as poor radio signal or low device battery.

The evolution from standalone to split processing is shown in Figure 5.

Figure 5. XR processing from device to cloud

Summarizing, XR can be delivered across different architectures, standalone, cloud or hybrid, depending on device and network capabilities:

- Stand-alone HMD to the 5G system (in the future)
- Tethered HMD via smartphone or USB stick to the 5G system (mobile scenario)
- Wi-Fi Tethered HMD to the 5G system via Fixed Wireless Access (FWA)
- HMD connected via Wi-Fi and/or local area network (LAN) to the fixed access network (fixed scenario)

Deployment scenarios

While Figure 3 represents the logical architecture, Figure 6 shows some industrial, enterprise and consumer XR network deployment examples. Depending on device type, device operating mode and network capability, the application stack and related processing can be distributed in different ways.

Figure 6. Deployment examples of industrial, enterprise and consumer use cases

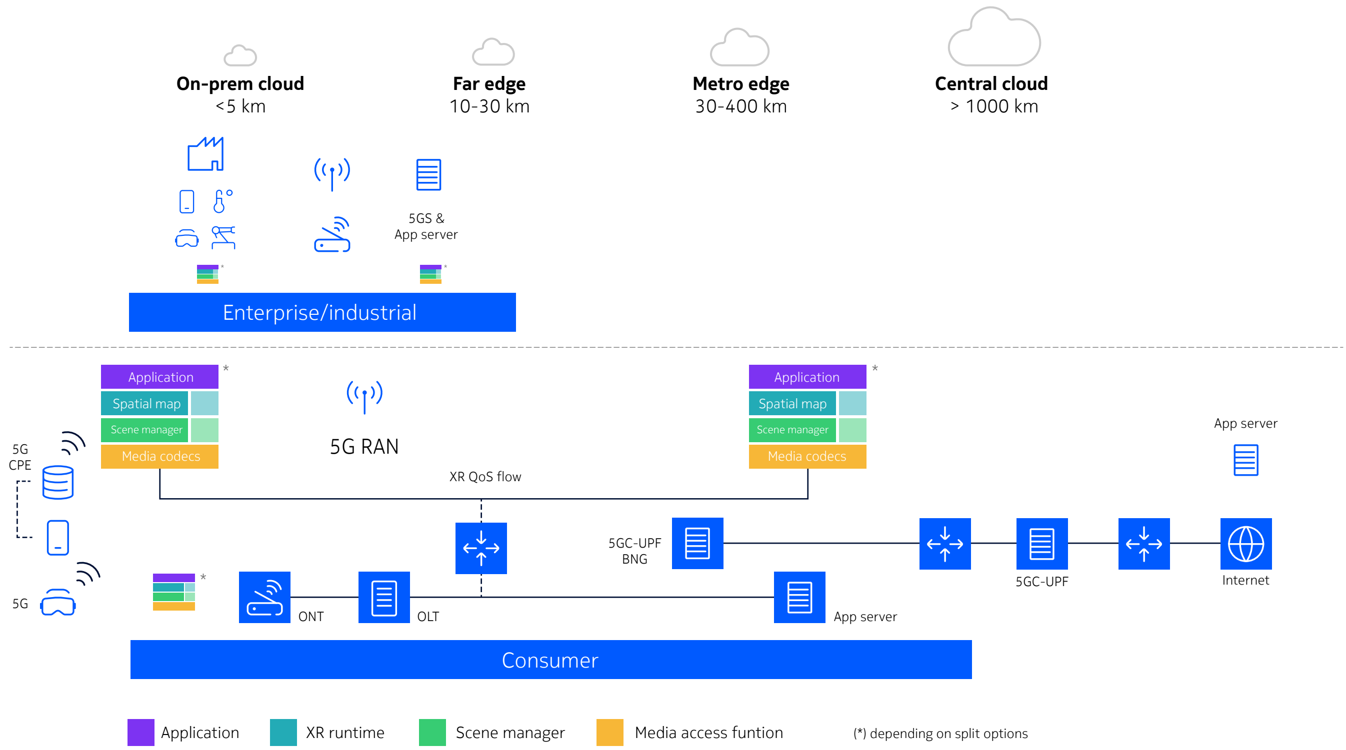

Industries and Enterprises: Digital transformation requires connectivity for a variety of devices such as sensors, robots, cameras, tablets and head-mounted displays. Industrial workers need to operate in real time over machines with HMDs and tablets. Security and latency requirements dictate that both the entire connectivity infrastructure and data processing be placed very close to the devices, either on-premises or at the far-edge. Furthermore, while today most of the processing takes place on the device, in the future lighter glasses with higher video definition will offload even more of the processing, such as cloud rendering and split processing.

Consumer: Smart glasses are the device of choice for social media and entertainment applications. In contrast to Industrial XR, where the server is deployed either on-premises or at the edge due to requirements like security and video analytics processing, the edge server distance in consumer deployments is mainly dictated by the need of use cases for cloud and/or split processing requirements to drive the HMD. In addition, the compute resources needed to do the rendering should be dynamically and intelligently selected depending on the application requirements and environment resources available.

Capabilities of current networks

5G

Simulations by 3GPP [7] have proven that 5G Rel.15 can support 4-10 simultaneous XR users with high-quality real-time video in each cell of the network deployment. The number of users depend on the Packet Delay Budget (PDB) and the requested throughput. Figure 7 shows the DL and UL capacity in frequency range one (FR1) in outdoor dense urban (DU) and frequency range two (FR2) in both Indoor hotspot (InH) and DU scenarios. While the available bandwidth is 100 MHz in both cases, InH deployment shows higher capacity than DU scenarios because of better propagation conditions: shorter distance to the base station (BS), use of three antenna panels mounted on the UE, and higher beamforming gain. The DL is, in many use cases, the limiting factor for network capacity because of transmission of high-definition video, e.g., 45 Mbps, while head and hand pose estimation requires only a few hundred Kbps in the UL. The UL can instead be the limiting factor when applications transmit UL video, e.g., as shown in the figure, the 10 Mbps UL video in DU-FR1 limits the number of users to five. Finally, simulations show better performances for cloud gaming (CG) compared to AR/VR because of less stringent delay requirements, i.e., 15ms against 10ms AR/VR PDB.

In addition to network capacity, the device power usage in a 5G Rel.15 network is not optimized for XR traffic patterns.

Figure 7 - XR capacity evaluation results in FR1 and FR2

DL capacity

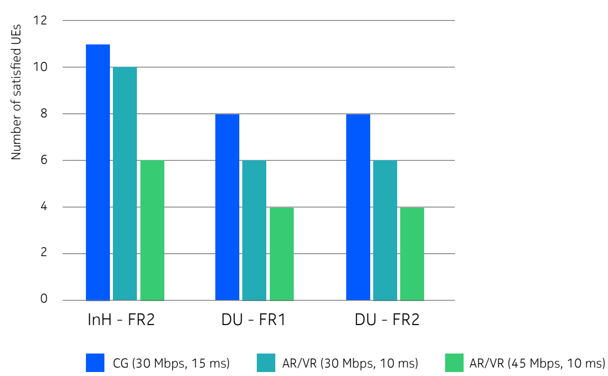

UL capacity

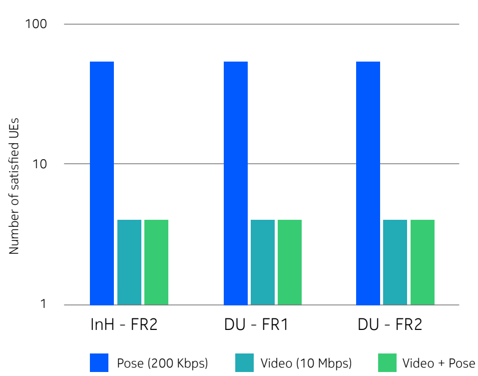

Wi-Fi

XR requires a certain PDB to be fulfilled. This is difficult to obtain with pre-Wi-Fi 6E generations in scenarios with multiple users and/or applications, because XR requires prioritization and admission control to avoid congestion. Wi-Fi 6 does add target wake-time (TWT) capabilities for treating multiple real-time traffic streams in parallel for collision avoidance, and it also uses basic service set (BSS) coloring for throughput increase. However, it still operates, as with previous generations, on congested 2.4 and 5GHz spectrum. In contrast, Wi-Fi 6E, as well as operating on 2.4 and 5GHz, additionally operates on the 6GHz frequency band, where there are currently no interference problems. As Wi-Fi is designed for short-range, low-mobility use, outdoor use-case scenarios will need to be served with a cellular technology.

DSL

DSL technologies use twisted pair networks for broadband access. VDSL2 17a can reach up to 100 Mbit/s down, 30 Mbit/s up. This looks sufficient if you consider a single XR application. As a peer-to-peer (P2P) technology, there is also no competition for this bandwidth, which keeps latency lower than 10 ms, even with retransmissions. The broadband consumption profile over a fixed network, however, is an order of magnitude higher than a mobile network due to the amount of sustained traffic. A fixed network is typically designed with the ability to handle heavy users with multiple parallel high-demanding streams (P99 model). Therefore, operators and governments are pushing for gigabit fixed networks, which VDSL2 can no longer provide.

Summarizing, today's fixed and wireless networks can, depending on the use case, support certain types of XR traffic. VR is mainly used indoors and often needs a fiber connection. AR can offer a more varied set of application experiences, both in usage and location. Many AR applications are already available on smartphones/tablets, such as Snapchat Lenses and shopping applications that add a digital overlay to the view in a smartphone camera in real time, allowing the user to change how they appear or to visualize digital versions of goods before buying them online. These experiences occur both indoors and outdoors, but their effect on the network is negligible in terms of both throughput and latency needs.

XR at scale

3GPP and other standards development organizations (SDOs) foresee a set of network enhancements (as shown in Figure 8) to support XR at scale. These enhancements will increase the network capacity and minimize the user equipment (UE) battery consumption when consuming XR applications.

While providing XR, the operator's network needs to ensure that other types of traffic are served according to their needs. One effective mechanism is network slicing.

Figure 8. Network infrastructure enhancements

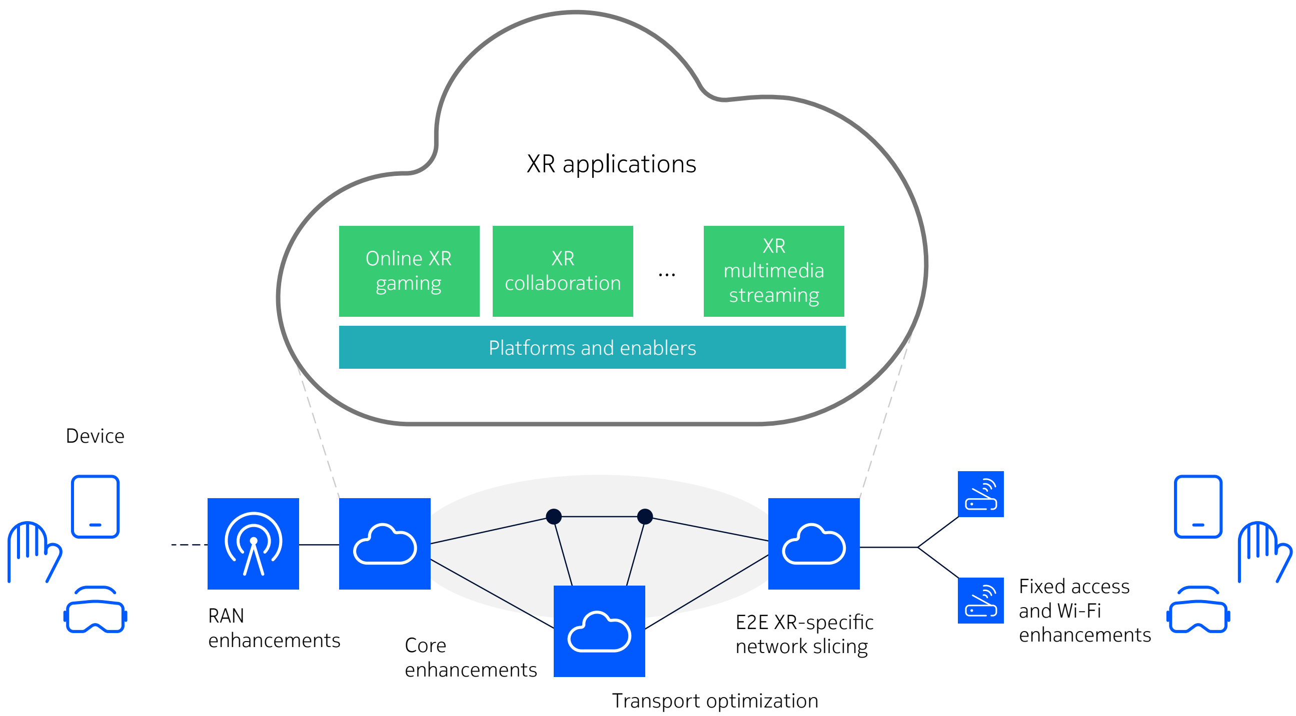

Access and transport technologies

Low and consistent E2E latency is fundamental to providing a good quality of experience (QoE) for all interactive services, whether it is in mission-critical services (e.g., remote control of straddle carriers at ports), business services (e.g., virtual meetings), or leisure (e.g., online gaming). To support XR at scale, the fixed network — the wireless access and the backhaul infrastructure — need to evolve to offer more throughput and deliver suitable QoS end-to-end.

The evolution of fixed access (Figure 9) shows a continued push to increase throughput higher, while also taking care that the UL/DL ratio increases. VDSL2's throughput is capable of hundreds of Mbps with the latest 35b profile. The gigabit-capable G.fast adopts time-division duplexing (TDD) that allows flexible, or even dynamic, allocation of the DL/UL ratio. Cable networks were initially designed for pure downlink transmission, but they have recently developed the fourth generation of the Data Over Cable Service Interface Specification (DOCSIS 4.0 or D4.0) allowing both upstream and downstream to transmit over the full spectrum. Passive optical network (PON) may be the ultimate future-proof fixed access solution. Its virtually unlimited capacity allows operators to move towards full symmetry with 10-Gigabit symmetrical (XGS-PON), 25-Gigabit symmetrical (25GS-PON) and beyond.

Figure 9. Fixed access evolution for DSL, cable and PON

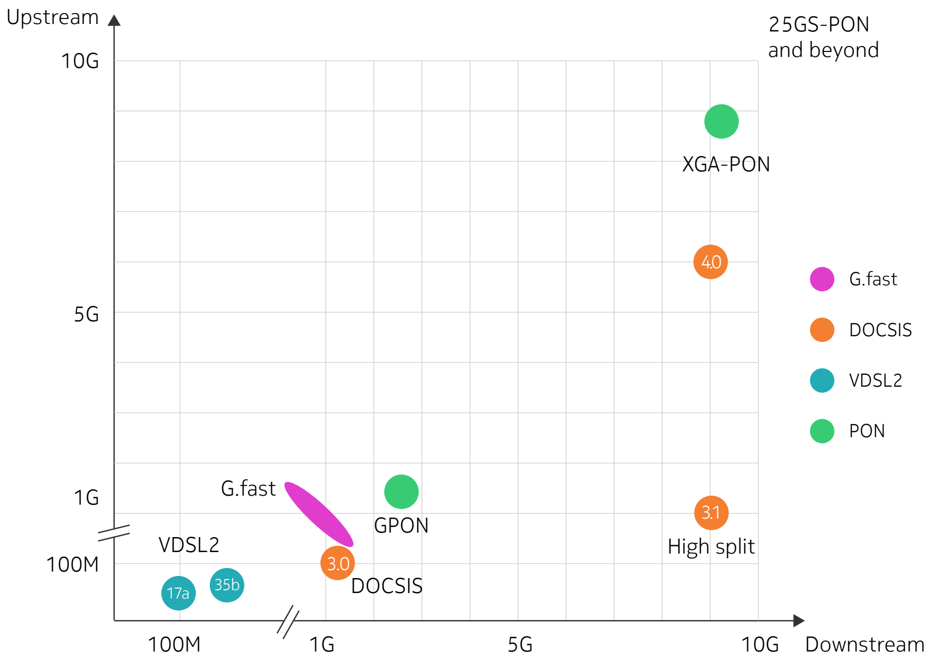

During the past three years, the adoption of fiber networks has increased exponentially for both existing telcos and new, mainly fiber operator entrants. Cable operators have been a bit more reluctant to make this transition. They may be forced to, however, because D4.0 allows no evolution path to higher/symmetrical bit rates at a reasonable cost. In many cases, a D4.0 transformation will cost more than a complete fiber-to-the-home (FTTH) overbuild that brings additional benefits like increased energy efficiency, or convergence benefits that allow operators to run several services over a single network.

The fiber network also serves as backhaul for small cells and/or Wi-Fi access points connecting XR devices. We expect that all networks will migrate towards a single, converged energy and cost-effective fiber network [9] supporting all services and future applications, including XR .

The wireless connection will either be cellular, Wi-Fi, or both. While the debate between Wi-Fi and cellular is often posed as a competition, we believe both to be essential and complimentary solutions.

Wi-Fi is predominantly used in enterprise and residential use cases for indoor (short-range), non-mission-critical nomadic applications. Today, Wi-Fi offload enables increased QoS for higher-margin and mobile use cases [10]. This is only possible, however, when the underlying cellular and Wi-Fi technologies can support the same set of use cases. Wi-Fi 7 addresses loaded environments with more transmission capacity and multi-link operation with flexible traffic allocation, allowing an order of magnitude improvement in throughput, latency, and jitter. Wi-Fi 8 will include further enhancements to realize ultra-high reliability.

5G and 5G-Advanced are needed for outdoor and FWA to 5G and mission-critical scenarios, where many users compete for the network resources and the network needs to ensure QoE for XR and other services such as voice, enhanced broadband IP applications, and industrial IoT. In industrial and enterprise scenarios, 5G and 5G-Advanced will assure QoS for users consuming XR and other services as well as serving many other device types like IoT. The 5G system (5GS) includes both radio access and core network, and natively handles traffic with different characteristics and QoS requirements. Its radio bearers enable different schemes for scheduling radio resources in DL and UL.

While most applications require high throughput in DL, only some generate high data rate in UL, e.g., for video capture streaming in concerts, sport events, or holographic calling. Although there are different throughput characteristics, the UL and not the DL is emphasized in Figure 10. The two main reasons to focus on the UL are because there is less UL bandwidth available compared to DL and the limited device UL RF power. Enhancements are needed to improve the available UL throughput and coverage for a given guaranteed latency. Figure 2 shows the network requirements and their dependency on the XR use case.

For the expected XR DL throughput on the contrary, it is possible to leverage the optimum 5G Network modulation type and frequency bands. For example, for sub-6GHz (FR1), TDD systems have frequency bands offering up to 100MHz bandwidth and multiple carrier components. Alternately, for even higher XR capacity and other demands, the high throughput and bandwidth capabilities of mmWave (FR2) would be a good option.

Figure 10. Network throughput requires a 4x boost in UL to guarantee 20 Mbps

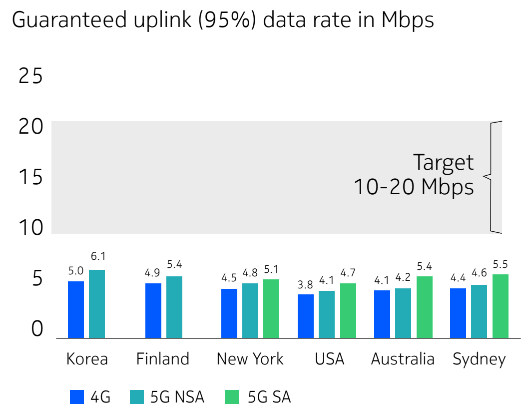

Source: Nokia, crowd-sourced data collected from Tutela in December 2022.
Latency measured as roundtrip time to Amazon, Google or Akamai CDN test server

The core limitation that restricts network capacity is the combination of low latency requirements — e.g., head, hand, or even body pose information or control actions — with high bandwidth requirements. Leveraging a local UPF, introduced with 5G, enables the use of an edge cloud. Compared to a central cloud, the latency from an XR device to an edge cloud is much lower. This allows lightweight XR devices as the processing can be offloaded to the edge cloud.

Figure 11 shows the latency of current networks against the required values for mid- to long-term use cases as well as network deployment enhancements.

Figure 11. Network latency requirements

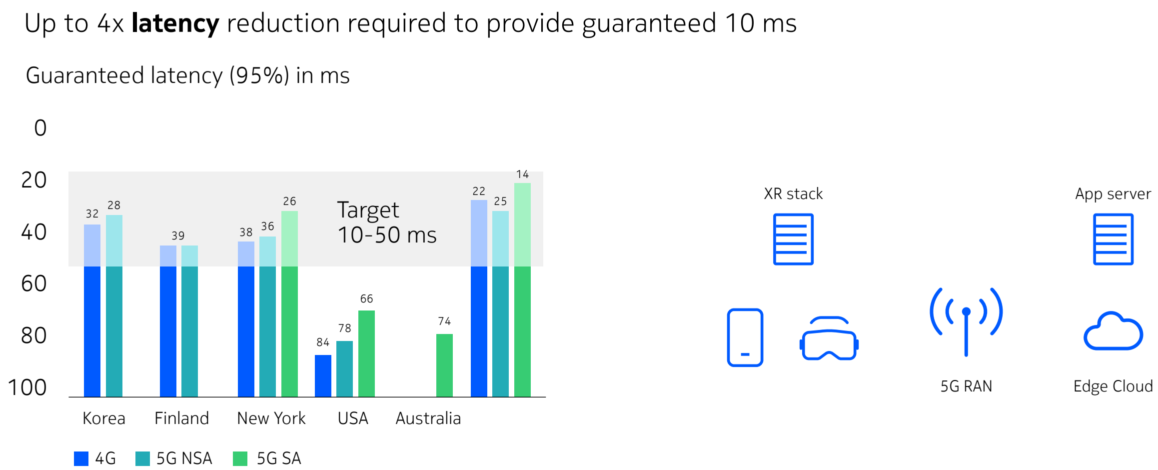

Source: Nokia, crowd-sourced data collected from Tutela in December 2022.
Latency measured as roundtrip time to Amazon, Google or Akamai CDN test server

3GPP has proved that 5G Rel.15 and 16 can support XR services for a limited number of users. Rel.17 introduces dedicated guaranteed big rate (GBR) 5G QoS identifiers (5Qis) for XR traffic and application interaction with the network to request specific QoS (delay requirements) to the 5GS.

In addition to previous releases, 3GPP Rel.18 (5G-Advanced) is defining a framework and feature-toolset, represented in Figure 12, for bridging the gap between the network and consistent XR experience while providing XR at scale and enabling device's power saving.

Figure 12. Rel.18 toolset for XR at scale

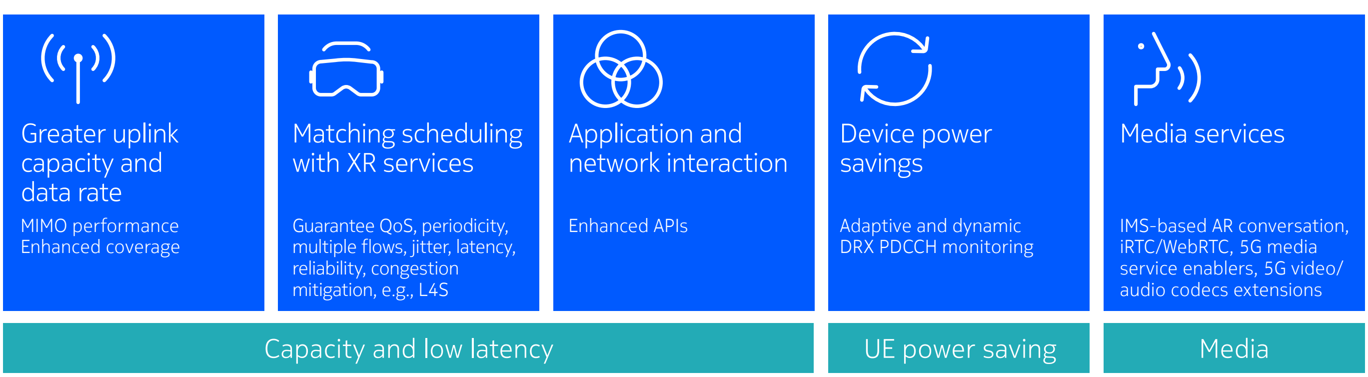

In general, transport networks, e.g., aggregation, metro, and core, are not a discard point since they are dimensioned to support the radio access network (RAN) spectrum (i.e., the most expensive resource). There are, in any case, options to mark the packets at the optical line terminal (OLT) and then map them to MPLS or other labels for a preferential service-level agreement (SLA) treatment along the path, for example for providing transport slices.

Network slicing is an important tool to guarantee the XR QoS while assuring other non-XR services. The QoS needs for XR applications are unique as they typically demand both high bandwidth and low latency in the same slice. A network slice that could guarantee the desired QoS for a given XR application can be used in a dynamic manner. The UE route selection policy (URSP)-based mechanisms of the core network can push the dynamic policies from the policy control function (PCF) to the UE, enabling it to choose the right QoS user plane for a given XR application. In a similar but UE-agnostic approach, a user plane function (UPF) enabled with a UL classifier (ULCL) could be used to select a user plane on a slice with closer and hence reduced device-to-application latency for user plane offloading for the specific XR application.

E2E enablement

Any type of XR application running partially or completely in the cloud poses strict requirements to the network. The network must:

- Maintain the device-to-server (D2S) or device-to-device (D2D) latency and jitter within the required limits of a given XR application
- Ensure a consistent throughput for the D2S or D2D in both UL and DL direction as needed by a given XR application
- Ensure that non-XR users are not starved for network resources as a result of prioritizing and maintaining the required QoS levels for XR users
- Support split-processing architectures to enable the XR UE to run cordless for multiple hours before recharging thus reducing battery size and keeping it light.

These requirements entail an end-to-end XR session being able to reach and maintain the desired QoS goals. Being the transport network, one of the main ways to reduce D2S2D latency, is to use an edge cloud to lower the distance between the access and application server. Figure 13 shows the breakdown of the E2E delay.

Figure 13. E2E latency

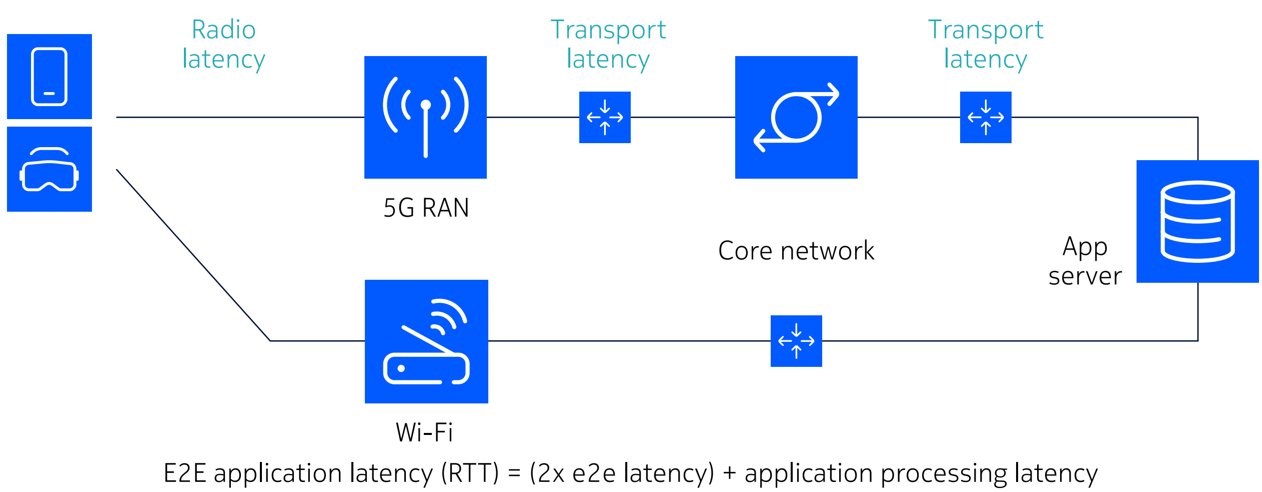

How Nokia enables XR E2E

Nokia develops technologies, services and offers a full portfolio of products for end-to-end connectivity, including radio, Wi-Fi, fixed access, core network, IP and optical transport, as well as products for orchestration, assurance, and network slice controls for radio, transport, core network and the end-to-end network. These products can be used to build local or wide area networks in single or multivendor environment.

Figure 14. Nokia portfolio

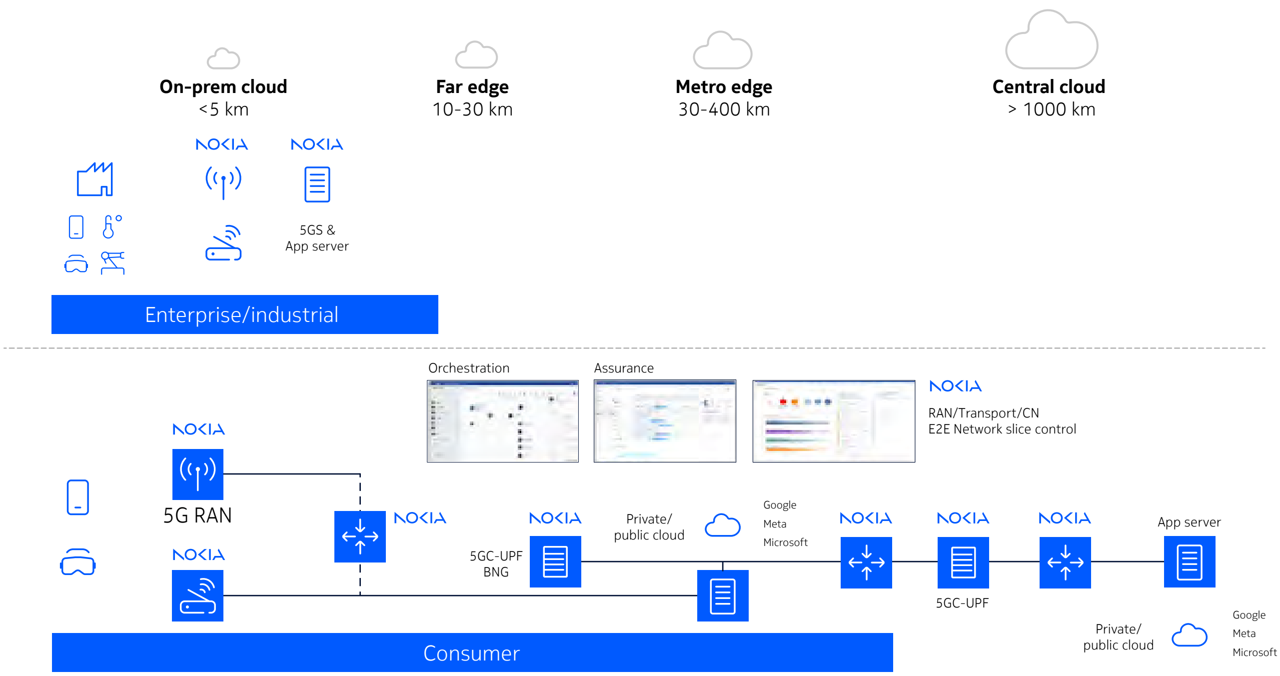

Network programmability is enabled by network capabilities exposure by means of network-as-code APIs, enabling developers to create network aware applications. Technologies to support XR use cases will involve multiple domains and system components spanning devices, radio core and transport network technologies along with optimization technologies that improve end user experiences. The Nokia portfolio is shown in Figure 14.

Ecosystem

Although connectivity and devices are key elements, the XR technology stack also needs platforms, content and tools such as codecs/media, digital twins and 3D engines. While developing technologies, Nokia is working with key global players to experiment, develop and integrate networks, devices and applications to ensure the best user QoE for XR and other services. Further standardization will also be needed to achieve interoperability among key XR components to enable XR at scale. Nokia is driving significant work items in key standardization bodies.

Standardization and industry Initiatives

Nokia has driven many fundamental technologies used in fixed and mobile devices and networks and holds leadership position in ITU-T, BBF, Wi-Fi Alliance, 3GPP 5G & 5G-Advanced standardization. Being rapporteur of XR key specifications and among the top three contributors in 3GPP, Nokia is shaping XR to drive network evolution of both operators as and vertical industry sectors. Nokia is driving Immersive Video-based visual volumetric coding as well as Immersive Video and Audio standards in both MPEG and 3GPP and holds key positions in other relevant groups. Nokia is among the initial board of the Metaverse Standards Forum and drives the "Network Requirements and Capabilities to support Metaverse Applications" Working Group.

Successful evolution of Extended Reality and Metaverse requires a combination of technology stack, integration, and interoperability amongst the various technology components and companies. Nokia believes that collaboration across companies and development following the principle of openness are keys to making XR successful. Standards will play an essential role to ensure interoperability including running applications across different devices, infrastructure, and platforms.

As device modes move from standalone mode to split processing, the end-to-end processing pipeline from cloud to device becomes critical for technology evolution, integration, and interoperability. The Metaverse Standards Forum (MSF) has been recently incorporated with 2400 member companies from Standards Developing Organizations and advocacy organizations to accelerate the mission of XR and Metaverse throughout the technology stack. MSF will coordinate interoperability prototypes, POCs, and testbeds with industry partners, and coordinate requirements & insights with existing SDOs. Figure 15 describes the interaction between MSF and the different industry partners and SDOs.

Figure 15. The XR ecosystem

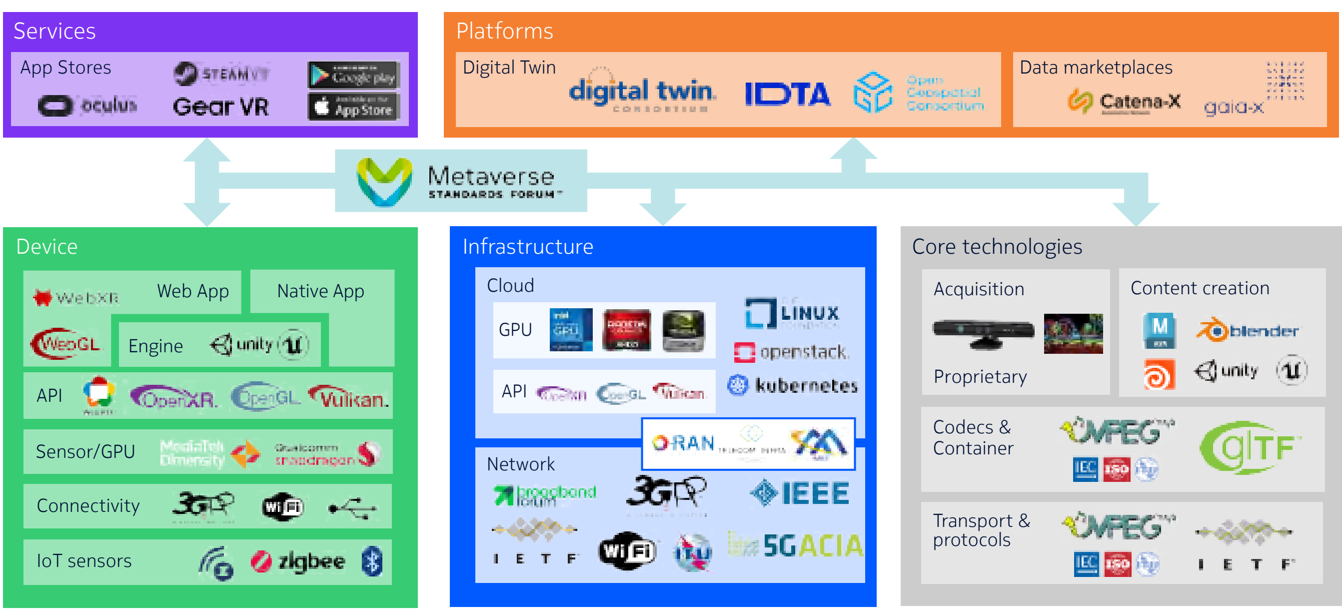

XR use cases piloted at Nokia

Nokia Real-time eXtended Reality Multimedia (RXRM)

Nokia RXRM [11] is a software-based XR media processing and streaming solution for industrial use, especially for situational awareness and efficiency. Use cases include remote supervision, teleoperation, security and surveillance, quality control, pre-emptive maintenance, remote inspection, and expert aid, as well as remote training.

RXRM leverages ground-breaking 360° video streaming technology and spatial 3D audio that is perfected for real-time streaming over low-latency networks. It provides a scalable real-time XR media streaming solution, especially when deployed on top of private wireless networks.

RXRM provides throughput-efficient “viewport-dependent” 360° video streaming technology that is guided by a low delay control data flow (e.g., head pose information). This low-latency interaction between the viewer and media server enables up to a 90% reduction of required streaming throughput for real-time 360° videos. This reduction helps to relax the network requirements for XR and significantly improves scalability for massive numbers of XR devices. This elevated level of interactivity can be achieved with a high throughput and ultra-low latency network, as well as edge processing.

Processing of the RXRM video streams at the edge enables scalability, so that multiple viewers (be it a human or an AI task) can benefit from the desired visual information with a best level of latency and smallest load on the underlying network resources. This aspect is especially important for on-premises industrial deployments where network resources are shared between different real-time information sources and vital for business continuation and productivity. Figure 16 shows a simplified XR media streaming flow of RXRM.

Figure 16. Simplified XR media streaming for RXRM

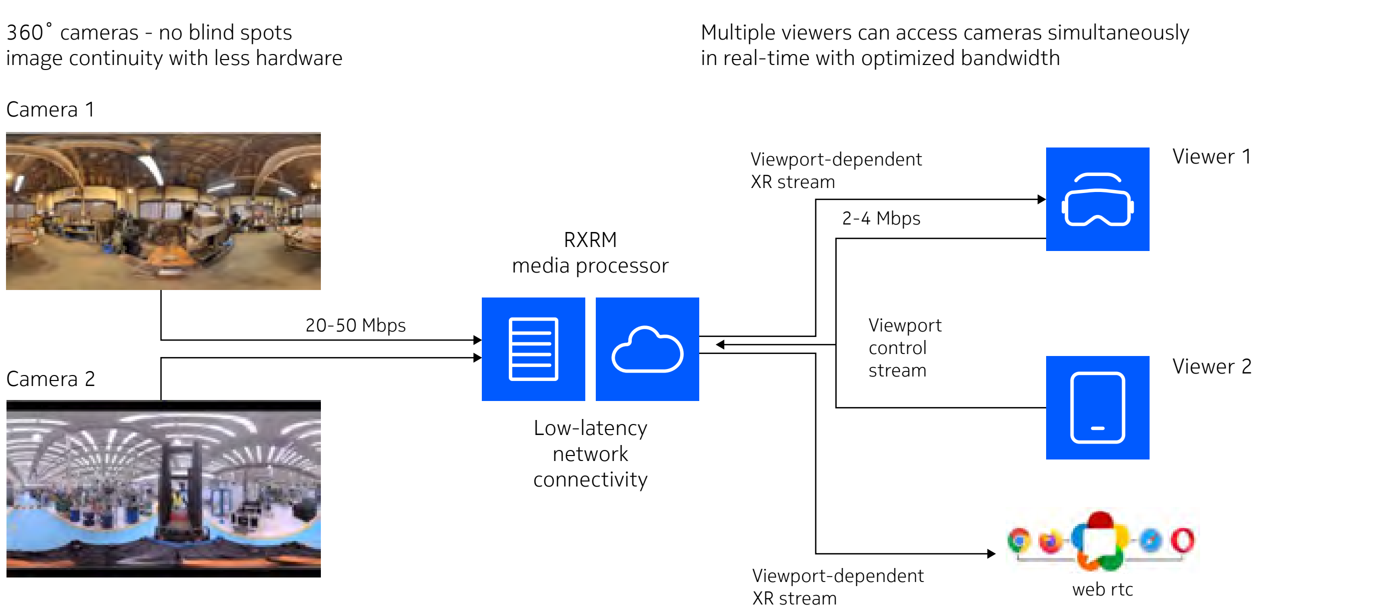

RXRM can use multiple legacy and 360° cameras in real-time. The RXRM media processor ingests the source videos and then provides throughput-efficient XR media streams to RXRM viewers. Viewers only receive the parts of the real-time media that are of interest (i.e., viewport), which drastically reduces the throughput and provides scalability to multiple users simultaneously. RXRM viewers can use XR and AR headsets, as well as mobile phones, tablets and PCs.

RXRM can be deployed as a stand-alone service, on-premises or on cloud, even on embedded devices that are on vehicles. Its flexible architecture is designed to support many use cases from teleoperations to visual monitoring.

The Nokia MX Industrial Edge (MXIE)

Industry 4.0 requires widespread digitalization and connectivity of equipment, machines and other industrial environments' assets (Figure 17). Many companies need a private network and, since they run in remote locations or temporary sites, e.g., mines, power plants, offshore oil platforms, factories, warehouses, or ports, the connectivity is challenging and use cases require typically very low latency favoring edge processing, sometimes even of core functions. 75% of all enterprise generated data will be processed at the edge. The edge cloud is not only placing computing resources closer to the source of data, but it also enables real-time collaboration between cyber-physical systems. And it drives actionable intelligence for industrial automation, situational awareness and worker safety.

Figure 17. Industry 4.0 in production

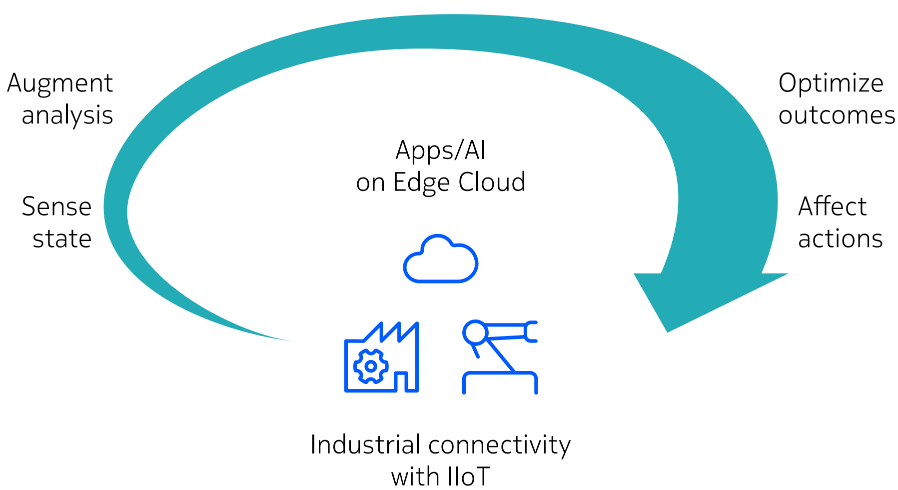

Source: Gartner Edge Computing in support of the Internet of Things

In these industries, large amounts of data sent by various devices and sensors needs to be handled in real time for process automation; at the same time humans monitor, instruct, and execute several tasks by using diverse types of devices, tablets for 2D operations as well as HMD for 3D immersive operational experiences.

The Nokia MX Industrial Edge (MXIE) [12] is part of the E2E Nokia Digital Automation Cloud (NDAC) ecosystem [13] for industry transformation and automation. It is a future-ready edge solution that accelerates the digital transformation of industrial operational technology (OT) environments. OT provides the technology to monitor and control physical processes, systems, devices and infrastructure in a production environment. It combines the agility and simplicity of an edge-as-a-service model with a high performance, resilient and secure edge architecture designed to meet the mission-critical needs of asset-intensive industrial environments.

MXIE runs virtualized functions including a local connectivity network and the wireless connectivity (5G, 4G, and Wi-Fi). The platform can be used to enable Nokia and third-party applications such as high-accuracy object positioning, push-to-talk and scene analytics, some of which have ultra-low latency requirements. To utilize the full ecosystem including webscalers Nokia MXIE has expanded to onboard third-party software suppliers. Taqtile is such a third party with a footprint in the XR market already.

Part of the application catalogue running on MXIE, Taqtile's 3D AR platform, Manifest, gives deskless workers access to step-by-step guidance and remote assistance from remote experienced technicians and trainers (Figure 18).

Figure 18. Worker operating in 3D

By leveraging augmented reality, Taqtile' s Manifest improves operational work flows so that jobs get done more accurately and consistently. Using Manifest, deskless workers document maintenance and operational procedures, follow step-by-step instructions, remotely collaborate with experts, and analyze job performance and history. Figures 19 and 20 show the on-premises deployment of MXIE, which highly improves the user experience by significantly lowering the latency. It also maintains data sovereignty by ensuring that all data stays on premises. Figure 19 shows MXIE deploying Taqtile Manifest from a high level and Figure 20, the MXIE architecture.

Figure 19 - MXIE Taqtile private wireless

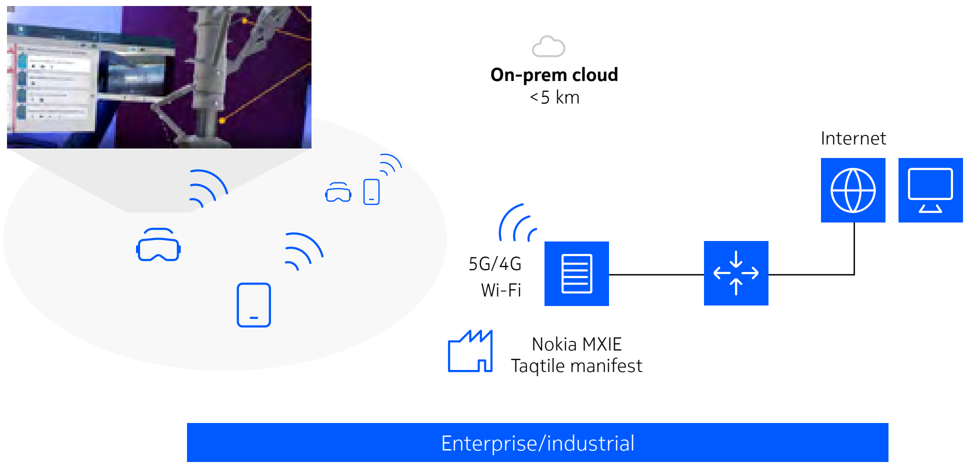

Figure 20 - MXIE private wireless architecture

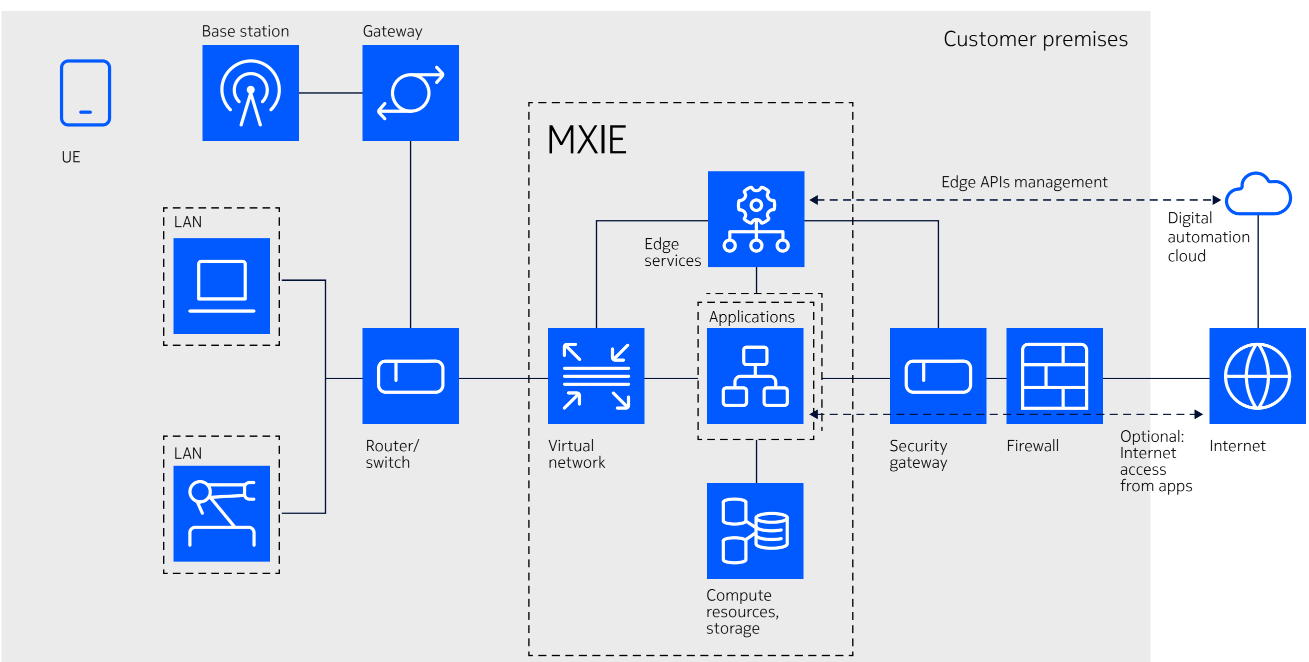

Learning Space

The learning business is rapidly expanding from classical face-to-face training into digital and remote learning. This is an ideal use case for XR technology. Nokia's Learning Space application is a real-life example of this. Multiple teams have already been trained and certified using it. Learners enjoy new learning experiences and interact with the simulated world using the Nokia Learn Mobile app and Nokia Learning Space VR platform.

Engineers can, for instance, augment a work procedure by using their phone or tablet and follow a step-by-step work instruction. Alternately, they can be immersed in a VR experience to replicate the installation procedure of Nokia equipment in a safe environment. The experience is impactful as the learner is fully immersed in the hands-on activity using the HMD. The view simulates an actual experience and includes spatial audio. User actions are directly reflected in the simulation and subjects not normally visible to humans can be visualized, e.g., simulation of radio signal propagation, traffic in a fiber cable or bytes flowing within a network element.

One example is the procedure to install and commission Nokia equipment, where the engineer can perform the end-to-end installation tasks and get certified with a 3D digital replica of it. The XR classroom is the place where people can meet with instructors and peers to learn, attend workshops, do simulations, and connect as if they were in the same physical space. The Learning Space allows them to more explorative as the user is at the center of the experience and can move into either real or imaginary places or situations without needing to travel.

Figures 21, 22 and 23 show examples of VR Learning Space experiences.

Figure 21. VR-based learning for hands-on practice

Figure 22. VR certification on installation and commissioning

Figure 23. Learning classroom with instructors and peers

Conclusion

As an umbrella term for VR and AR as well as cloud gaming, extended reality (XR) will be a crucial enabler for many use cases in the industrial, enterprise and consumer metaverses. XR devices will be essential for connecting the virtual and physical worlds. The paper introduces Nokia's view of XR use cases, network challenges, different access technologies and deployments.

Nokia believes that an E2E connectivity solution between XR devices and the application server is critical to meet the application requirements of low latency and high throughput. The paper demonstrates that today's networks can, depending on the use case, support certain types of XR.

We believe in a future where XR will be deployed at scale. This will require a set of network enhancements to increase the network capacity and minimize UE battery consumption. In addition to connectivity, technologies for XR use cases involve multiple domains and system components spanning applications and devices along with optimization technologies for improving the end user experience.

While developing key technologies for meeting the use case requirements of future XR applications, Nokia is also working with key global players to experiment, develop and integrate networks, devices, and applications to ensure the best user QoE for XR and other services. Our collaborative work shows that to achieve interoperability among key XR components some standardization is needed; Nokia is driving this standardization effort by contributing to significant work items in the relevant standardization bodies.

Through ecosystem collaborations, standardization leadership, and industry initiatives, Nokia is enabling a compelling XR E2E solution that harnesses Nokia's strengths: providing products, services, and technologies to enable end-to-end XR solutions.

Abbreviations

25GS-PON	25-Gigabit Symmetrical PON
3GPP	3rd Generation Partnership Project
5GS	5G System
AI	Artificial Intelligence
AR	Augmented Reality
BNG	Broadband Network Gateway
BSS	Basic Service Set
CAGR	Compound Annual Growth Rate
CG	Cloud Gaming
CSP	Communications Service Provider
D2D	Device to Device
D2S	Device to Server
DL	Downlink
DOCSIS	Data Over Cable Service Interface Specification
DSL	Digital Subscriber Line
DU	Dense Urban
E2E	End-To-End
F2F	Face-to-Face
FR	Frequency Range
FOV	Field of View
FTTH	Fiber to the Home
FWA	Fixed Wireless Access
GBR	Guaranteed Bit Rate
GPU	Graphics Processing Unit
HMD	Head Mounted Display
IEEE	Institute of Electrical and Electronics Engineers
IETF	Internet Engineering Task Force
InH	Indoor Hotspot
IoT	Internet of Things
LAN	Local Area Network
L4S	Low Latency, Low Loss, and Scalable Throughput
MXIE	Nokia MX Industrial Edge
NaC	Network as Code
NDAC	Nokia Digital Automation Cloud
OLT	Optical Line Termination
ONT	Optical Network Termination
OT	Operational Technology
P2P	Peer to Peer
PC	Personal Computer
PCF	Policy Control Function
PDB	Packet Delay Budget
PON	Passive Optical network
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
RTT	Round Trip Time
RXRM	Nokia Real-time eXtended Reality Multimedia
SDO	Standards Development Organization
SLA	Service Level Agreement
TDD	Time Division Multiplexing
TWT	Target Wake Time
UE	User Equipment
UL	Uplink
ULCL	Uplink Classifier
UPF	User Plane Function
URSP	User Equipment Route Selection Policy
USB	Universal Serial Bus
VDSL	Very High-speed Digital Subscriber Line
VR	Virtual Reality
Wi-Fi	Wireless Fidelity
Wi-Fi 6E	Wi-Fi Sixth Generation Extended
XGS-PON	10 Gigabit Symmetrical Passive Optical Network
XR	eXtended Reality

References

[1] PR Newswire, “Global Extended Reality Market Report 2022: Rising Demand for Immersive Entertainment Bolsters Growth,” CISION, Feb 2023. https://www.prnewswire.com/news-releases/global-extended-reality-market-report-2022-rising-demand-for-immersive-entertainment-bolsters-growth-301751085.html

[2] Markets and Markets, “Metaverse Market Overview, Industry Share and Forecast,” May 2023. https://www.marketsandmarkets.com/Market-Reports/metaverse-market-166893905.html#:~:text=%5B273%20Pages%20Report%5D%20The%20Metaverse,at%20a%20CAGR%20of%2047.2%25.

[3] Nokia, “Metaverse explained.” Viewed 20 Jun 2023 at https://www.nokia.com/about-us/newsroom/articles/metaverse-explained/

[4] Nokia, “Six trailblazing use cases for the metaverse in business.” Viewed on 20 Jun 2023 at https://www.nokia.com/metaverse/six-metaverse-use-cases-for-businesses/

[5] Nokia, “Discover 2030: Our technology vision for a future where realities merge.” Viewed on 20 Jun 2023 at https://www.nokia.com/innovation/technology-vision-2030/

[6] LR Mobile News, “Docomo Counts Cost of Signaling Storm,” LightReading, Feb 2012. https://www.lightreading.com/mobile/packet-core/docomo-counts-cost-of-signaling-storm/d/d-id/693779

[7] 3rd Generation Partnership Project, Technical Specification Group Radio Access Network, NR, “Study on XR enhancements for NR (Release 18)”, 3GPP TR 38.835, Apr 2023. https://www.3gpp.org/ftp/Specs/archive/38_series/38.835/38835-i01.zip

[8] Petrov, P., Gapeyenko, M., Paris, S. Marcano, A. and Pederson, K., “Extended Reality (XR) over 5G and 5G-Advanced New Radio: Standardization, Applications, and Trends,” ResearchGate, Mar 2022. https://www.researchgate.net/publication/359055025_Extended_Reality_XR_over_5G_and_5G-Advanced_New_Radio_Standardization_Applications_and_Trends

[9] Nokia, “Nokia enhances fiber access portfolio to accelerate 5G,” Mar 2020. Viewed on 20 Jun 2023 at https://www.nokia.com/about-us/news/releases/2020/03/03/nokia-enhances-fiber-access-portfolio-to-accelerate-5g/

[10] Nokia, “Nokia introduces new Wi-Fi 6 mesh router and technology innovations to deliver ultimate gigabit experience in a 5G world,” Mar 2020. Viewed on 20 Jun 2023 at https://www.nokia.com/about-us/news/releases/2020/03/19/nokia-introduces-new-wi-fi-6-mesh-router-and-technology-innovations-to-deliver-ultimate-gigabit-experience-in-a-5g-world/

[11] Nokia, “Real-time eXtended Reality Multimedia.” Viewed on 20 Jun 2023 at https://www.nokia.com/networks/industries/rxrm/

[12] Nokia “MX Industrial Edge.” Viewed on 20 Jun 2023 at https://www.dac.nokia.com/mx-industrial-edge/

[13] “Nokia Digital Automation for industrial campus.” Viewed on 20 Jun 2023 at https://www.dac.nokia.com/

About Nokia

At Nokia, we create technology that helps the world act together.

As a B2B technology innovation leader, we are pioneering networks that sense, think and act by leveraging our work across mobile, fixed and cloud networks. In addition, we create value with intellectual property and long-term research, led by the award-winning Nokia Bell Labs.

Service providers, enterprises and partners worldwide trust Nokia to deliver secure, reliable and sustainable networks today – and work with us to create the digital services and applications of the future.

Nokia OYJ
Karakaari 7
02610 Espoo
Finland
Tel. +358 (0) 10 44 88 000

598285 (July) CID213401