Metaverse STANDARDS FORUM™

D2: Gap and Feasibility Analysis for Metaverse Services

Network Requirements and Capabilities Working Group

Scope:

This deliverable provides a gap and feasibility analysis based on networking requirements established in the NRC WG deliverable D1. The existing access networks in scope are listed and synthetically described, highlighting their key features and differentiators. The possible deployment scenarios are defined, highlighting a rich set of scenarios combining the above-mentioned access networks. For all network technology in scope, a feasibility analysis is conducted to check them against the requirements established in D1, the potential gaps or lacks are documented.

Versions:

Version	Authors	Note
1.0	Thibaud BIATEK (Nokia), Omar ELLOUMI (Nokia), Barry FERRIS (CableLabs), Jens JOHANN (DT)	First version

Table of Contents

References

[1] MSF NRC, "D1: Network Requirements for Metaverse Services".

[2] Nokia, "Tackling the network challenges for XR: Understanding, developing and delivering an E2E XR connectivity solution," 2023.

[3] ETSI, "Multi-access Edge Computing," [Online]. Available: https://www.etsi.org/technologies/multi-access-edge-computing.

[4] 3GPP RAN, "TR 38.835: Study on XR enhancements for NR (Release 18)".

[5] CableLabs, "Impacts of WMM on Wi-Fi," [Online]. Available: https://www-res.cablelabs.com/wp-content/uploads/2024/04/10100006/Impacts-of-WMM-on-Wi-Fi.pdf.

[6] Nokia, "Nokia and Vodafone conduct world first trial of L4S technology over and end to end PON network," [Online]. Available: https://www.nokia.com/about-us/news/releases/2024/04/03/nokia-and-vodafone-conduct-worlds-first-trial-of-l4s-technology-over-an-end-to-end-pon-network/.

[7] "CAMARA Project," [Online]. Available: https://github.com/camaraproject.

[8] GSMA, "GSMA Open Gateway," [Online]. Available: https://www.gsma.com/solutions-and-impact/gsma-open-gateway/.

[9] Ookla, "“Speedtest® by Ookla Global Fixed and Mobile Network Performance Maps,”," [Online]. Available: https://registry.opendata.aws/speedtest-global-performance/.

[10] Dgtl Infra, "United States Data Centers: Top 10 Locations in the USA," [Online]. Available: https://dgtlinfra.com/united-states-data-centers/.

[11] IETF, "Comcast L4S Field Trial Update," [Online]. Available: https://datatracker.ietf.org/meeting/118/materials/slides-118-tsvwg-sessa-61-l4s-experience-00.

[12] IEEE, "L4S contributions IEEE 802.11 group," [Online]. Available: https://mentor.ieee.org/802.11/documents?is_dcn=L4S.

Acronyms

API	Application Programming Interface
CAPIF	Common API Framework
FEC	Forward Error Correction
FWA	Fixed Wireless Access
HMD	Head Mounted Display
L4S	Low Latency, Low Loss, and Scalable Throughput
MAC	Medium Access Control
MEC	Multi-access Edge Computing
NEF	Network Exposure Function
OLT	Optical Line Termination
ONT	Optical Network Terminal
PDB	Packet Delay Budget
PDU	Packet Data Unit
PON	Passive Optical Network
QoE	Quality of Experience
QoS	Quality of Service
RAN	Radio Access Network
XR	Extended Reality

1. Introduction

Metaverse services are wide in scope, spanning from industry and enterprise to consumer use-cases. The use-cases have different expectations, technically and functionally, but also involve different technologies, e.g. from extended reality to digital twinning. The involved devices can also be different, from acquisition, haptics gloves, headset, sensors, vehicles to smartphones. Most of the use-cases can be mapped to various processing scenarios, involving different computing and connectivity architecture.

We've identified in [1] that a rich set of requirements need to be fulfilled to enable Metaverse services at scale. While we were focused on requirements and use-cases in our first deliverable, this deliverable looks at how these can be deployed, and assess if the existing networking technologies can support them. The figure below illustrates how deployment scenarios can differ whether you're considering enterprise or consumer services.

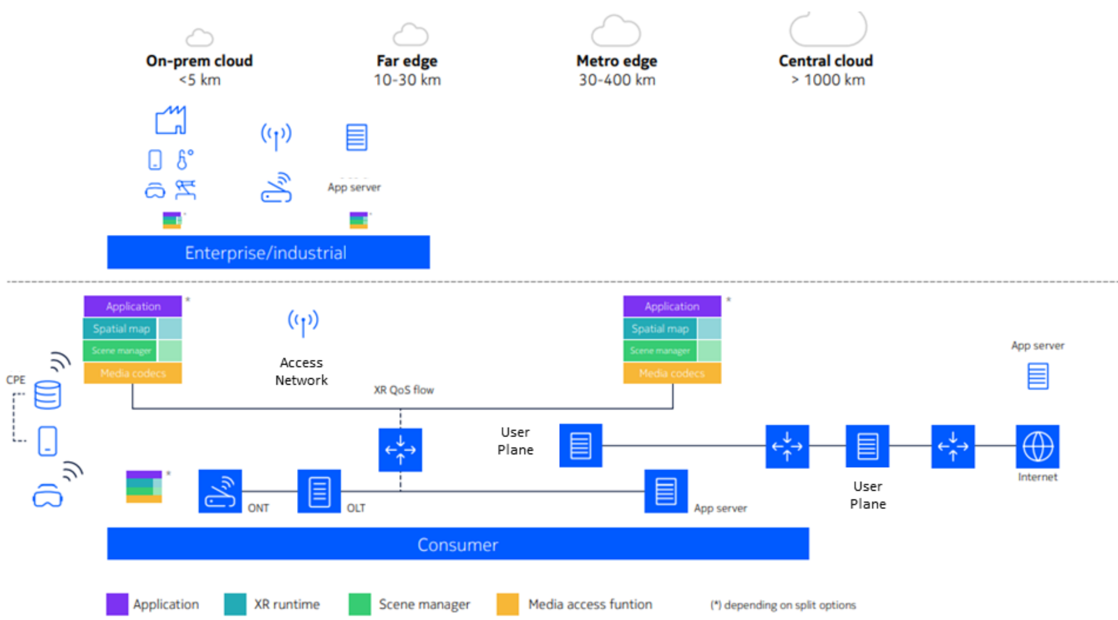

Figure 1: Deployment scenarios to be studied [2]

The deployment scenarios rely on underlying networking technologies that need to be identified. This may include fixed or wireless access, with a potentially large set of possible options. Hence, the relevant networking technologies need first to be identified. Then, it would be possible to identify the addressable set of deployment scenarios to be studied, including a combination of different networking technologies, and to check them against the requirements established in [1].

This deliverable is organized as follows. Section 2 provides a selection of networking technologies in scope. Section 3 establishes the considered deployment scenarios and identifies the possible combination of networking technologies that can be used. Section 4 evaluates the identified deployment scenarios against the requirements established in [1], identifying possible gaps or lacks in existing specifications. Finally, Section 5 concludes this deliverable, providing recommendations on how to address the possible identified gaps.

2. Technologies in scope

This deliverable investigates how the connectivity requirements established in [1] are addressed by state-of-the-art networking and computing technologies. In terms of connectivity, our use-cases may be served by a combination of multiple networking technologies, including local and access networks. The technologies in scope for our study are listed in the Table below.

Table 1: Networking technologies in scope of the gap & feasibility analysis

	Technologies in scope
Cellular network	5G • 3GPP Rel. 15 to Rel. 17 5GA • 3GPP Rel. 18 to Rel. 19 • Note: Rel. 18 is under finalization and Rel. 19 is about to start 6G • From 3GPP Rel. 21 to Rel. ?? • Note: Not started yet.
Local network	Wi-Fi 6/6E • Defined in IEEE 802.11ax Wi-Fi 7 • Defined in IEEE 802.11be • Note: currently rolling out
Fixed network	VDSL2 • ITU-T G.993.2 DOCSIS 3.1 • Defined by CableLabs • Note: deployed everywhere DOCSIS 4.0 • Defined by CableLabs • Note: just starting XGS-PON • ITU-T G.9807.1 25G-PON • tbd

Furthermore, the processing tasks identified in [1] may be deployed in edge or cloud servers located at a variable distance from the device. This deployment flexibility can be achieved by leveraging Multi Access Edge Computing (MEC) [3], enabling to address on-prem, far-edge, metro-edge and central-cloud scenarios.

3. Deployment scenarios

The networking technologies identified in section 2 can be used or combined in various ways, depending on the deployment scenario. Different contexts may require different deployment scenarios, whether it is for industry, enterprise or consumer, for example industry application requires the servers to be hosted on premise for security reason while a consumer service may leverage a public cloud service provider to host the compute tasks. In this report, two main deployment scenarios are described: industry/enterprise and consumer.

3.1. Industry/Enterprise deployment

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.

3.2. Consumer deployment

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.

3.3. Deployment in scope

The abovementioned deployment scenarios are quite similar in terms of high-level architecture, the main difference being in the distance the server is located from the device. Otherwise, the local network is interconnected with an application server through an access network. The different combinations are detailed in the Table below.

Table 2: Deployment scenarios variations in scope

	Local network	Access Network	Distance to App. server
Enterprise/Industry	WiFi		On-prem
	WiFi	PON	Far-Edge
	WiFi	Docsis	Far-Edge

	WiFi	VDSL2	Far-Edge
	WiFi	3GPP (Private Wireless)	On-Prem
		3GPP (Private Wireless)	On-Prem
	WiFi	3GPP (FWA)	Far-Edge
Consumer	WiFi	PON	Far-Edge->Central
	WiFi		On-Prem
	WiFi	Docsis	Far-Edge->Central
	WiFi	VDSL2	Far-Edge->Central
	WiFi	3GPP (FWA)	Far-Edge->Central
		3GPP	Far-Edge->Central

4. Gaps and feasibility analysis

4.1. Cellular networks

3GPP has assessed the performance of 5G for XR traffic [4], and concluded that Rel. 15 5G can support up to 10 simultaneous users with high-quality real-time downlink video. The number of users depends on the application throughput and required packed delay budget (PDB). The study has been done in the context of outdoor Dense Urban (DU) and indoor hotspots (InH), for FR1 and FR2 frequency ranges (100Mhz bandwidth per cell). The figures below provide the evaluation results.

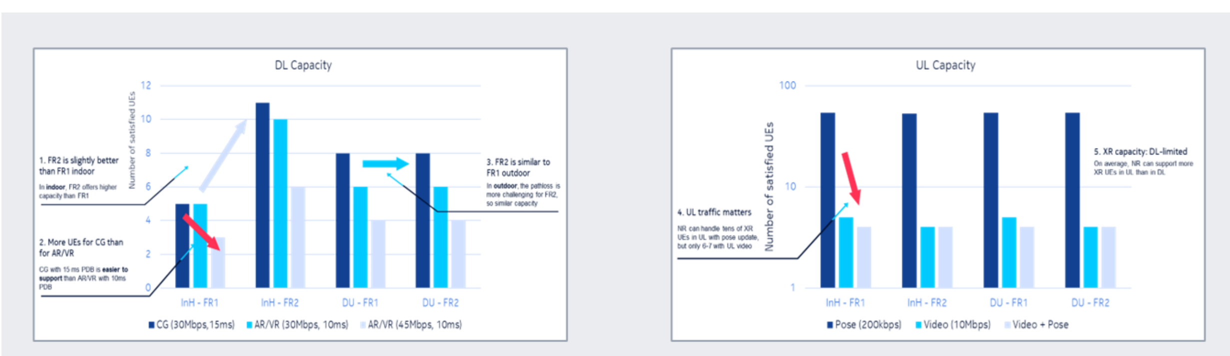

Figure 2: XR capacity evaluation results in FR1 and FR2

While this seems sufficient for most of our requirements established in [1], limitations may be identified when a high number of users are competing for resources. To address this, 5G and 5GA introduce native smart and flexible QoS handling through application aware and PDU set optimization. This increases capacity by enabling flexible scheduling, as illustrated in Figure 3.

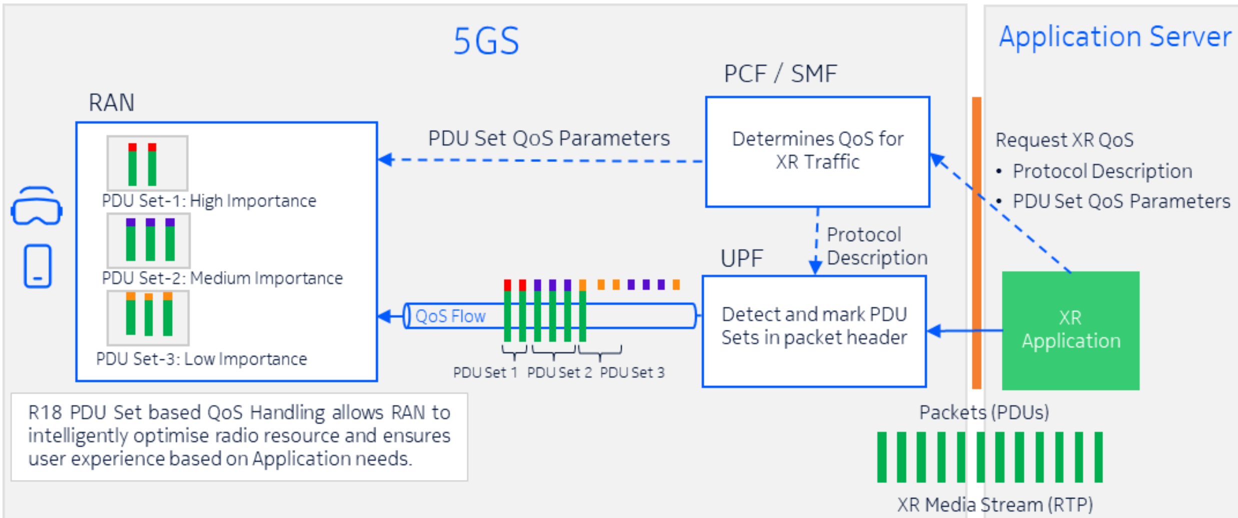

Figure 3: Rel.18 Application Aware QoS

Furthermore, congestion control mechanisms are introduced in Rel.18 to further support rate adaption in XR services. This enables to maximize QoE for the end-users and to avoid service interruption when network conditions are temporarily degraded on the RAN. The congestion control mechanism based on L4S is described in the Figure below.

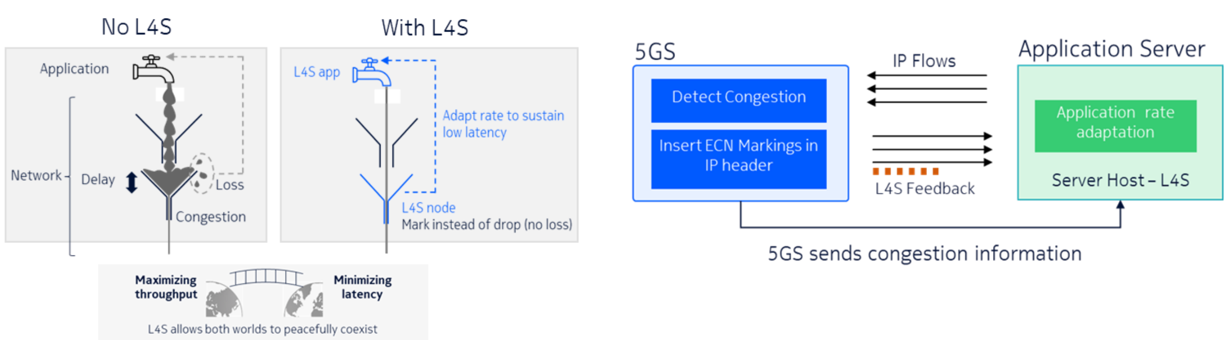

Figure 4: R18 L4S for scalable roll-out of robust low-latency services (UL and DL)

Our established requirements require in some cases important uplink capacity. While this can be addressed in theory, most of the 5G deployments today do not support such inverted uplink/downlink ratio, as highlighted in the Figure below. To circumvent this current limitation, more flexible duplexing may be needed.

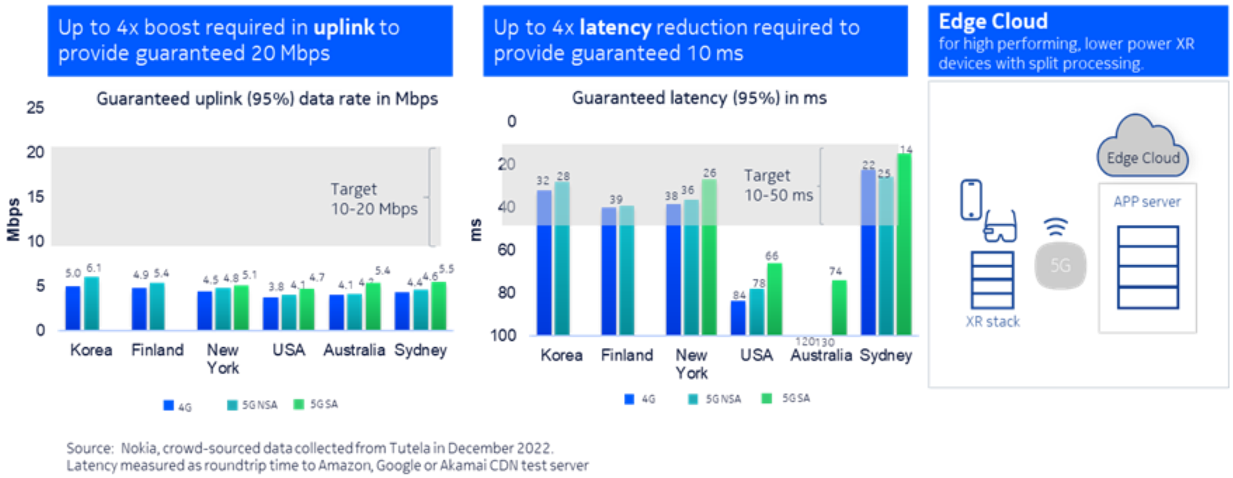

Figure 5: Currently available network throughput and latency

As uplink can be considered as a limiting factor for the most demanding applications, 5G-Advanced introduces enhancements to tackle this challenge, as illustrated in the figure below.

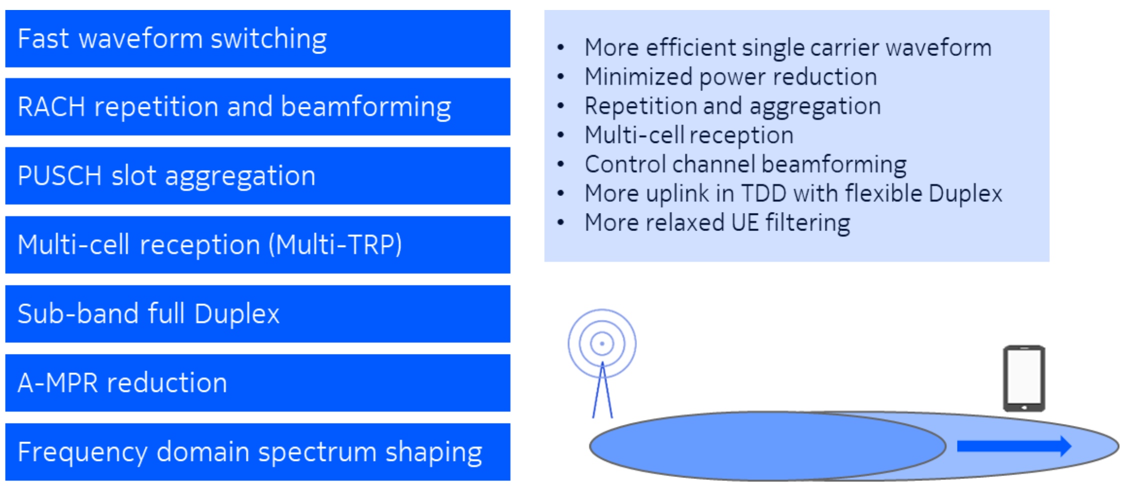

Figure 6: Rel.18 toolset for enhanced uplink capacity

In general, all the requirements established in D1 are achievable by 5G. The possible scaling issues are addressed by 5GA, by introducing a number of enhancements, depicted in the picture below.

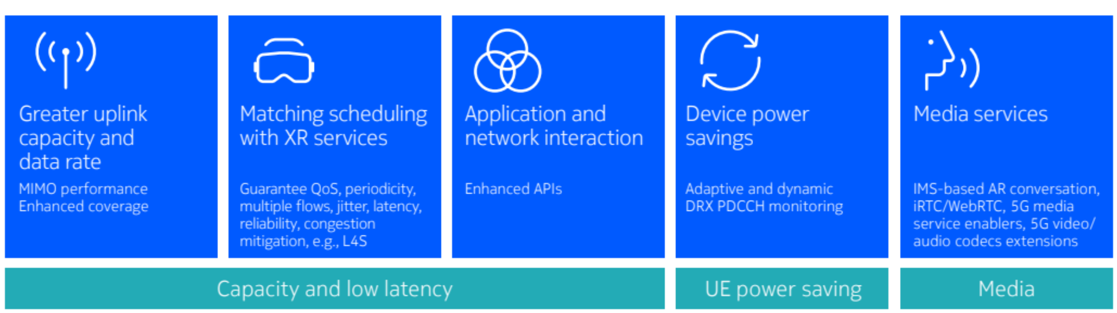

Figure 7: Rel.18 toolset for XR at scale [2]

While all the requirements are addressed by 5GA, both in terms of performance and functionality, some future protocol enhancements might not be compatible with existing features (QUIC vs PDU set handling).

4.2. Local network: WiFi (6/6E/7)

Metaverse requirements require a stringent packet delay budget, as identified in [1]. Achieving such low latency with pre Wi-Fi 6 technology cannot be done under certain load conditions. From Wi-Fi 6, new techniques are added to better manage congestion, such as Target Wake Time (TWT) and basic service set (BSS). However, Wi-Fi 6 is running on the busy 2.4Ghz and 5Ghz bands which limit its scaling potential. Therefore, WI-FI 6E is probably the first generation capable of running metaverse applications, in theory. When it comes to practice, it appears that Wi-Fi 6/6E cannot fulfill the latency requirements we've established as demonstrated in [4] with average latency measurements around 100ms. Enabling such low latency services requires the implementation of advanced congestion control mechanisms, such as L4S, as tested in this trial [5]. L4S can be used coupled with the WiFi router, without being adopted in IEEE standards for Wi-Fi, enabling to lower the latency from hundreds of ms to the MAC latency (~10ms). To summarize, Wi-Fi can be used as an access point for metaverse applications, from Wi-Fi 6E if it is coupled with low-latency congestion control mechanisms such as L4S.

4.3. Docsis 3.1/4.0

In DOCSIS networks, groups of end users - known as Service Groups - share the capacity of a DOCSIS link. Groups typically correspond to discrete geographic areas (e.g. a neighborhood or an MDU) and are sized based on capacity demand. A Service Group could consist of dozens or even hundreds of customers but can be scaled down as needed to increase the per-user capacity.

DOCSIS 3.1 supports data rates up to 10 Gbps in the downlink direction and 1.8 Gbps in the uplink direction, shared by the customers of each Service Group. DOCSIS 4.0 increases the uplink data rate to 6 Gbps. Thus, from a capacity perspective, each DOCSIS 3.1 or DOCSIS 4.0 Service Group could support hundreds of simultaneous XR sessions.

From a latency perspective, DOCSIS 3.1 networks support active queue management, and are software upgradable to support Low Latency DOCSIS features, which include L4S and NQB support as well as optimized low latency scheduling for particularly latency sensitive upstream services. DOCSIS 4.0 networks support both active queue management and Low Latency DOCSIS features as well. Low Latency DOCSIS deployments can achieve even the most stringent XR latency targets at the 99th percentile, with P99 round-trip times in the range of 1-10ms, while traditional DOCSIS deployments can meet many of the less stringent latency targets at the 99th percentile, with P99 round-trip times in the range of 50-100ms.

From a packet loss perspective, DOCSIS 3.1 and 4.0 protocols support powerful physical layer error correction techniques (FEC) which can achieve the PER targets for XR. DOCSIS does not support layer 2 retransmissions, so achieving the most stringent PER targets for XR would require the operator to actively monitor FEC error statistics and adjust FEC parameters accordingly.

No gaps have been identified.

4.4. PON (XGS/25G)

PON covers a wide range of technologies. We assess here the already well installed technologies and future looking ones, including GPON, XGS-PON and 25GS-PON. In the Table below, measurements done under operational conditions from few operators' networks are provided.

Table 3: PON measurement done on few operators networks (source: Nokia)

Technology	Typical throughput (100Mbps busy hour load, 1:64 split ratio)	Typical DS latency (without congestion)	Typical US latency (without congestion)
GPON	2 Gbps DS 900 Mbps US	24us (64B) 130us (jumbo)	500us – 1.5ms
XGS-PON	8 Gbps DS 7.5 Gbps US	25us (64B) 130us (jumbo)	500us – 1.5ms
25G-PON	20 Gbps DS 17 Gbps US	25us (64B) 140us (jumbo)	500us – 1.5ms

The throughput is measured under 100 Mbps busy hour load configuration, in 1:64 split ratio, meaning 1 OLT is serving 64 ONT. For throughput, both the downlink and uplink are measured. Worst-case scenario can be derived from these figures by dividing those by 64 typically. In terms of latency, both downlink and uplink cases are measured in non-congested environments. For downlink, the latency is spanning from 25us for 64B packets to 130us for jumbo packets. Latency is a bit higher for uplink, from 500us to 1.5ms, because of TDMA and more demanding communication with the OLT. From these figures, it is observed that PON technologies met metaverse services requirements in terms of downlink and uplink throughput and latencies.

In case of congestion, these numbers can still be achieved by implementing some traffic engineering mechanisms. L4S can address this and maintain near zero delay by eliminating queuing delay. Recently, this has been demonstrated end-to-end by Vodafone and Nokia Bell Labs [6], measuring 1.2ms latency at local ethernet port running on a fully congested access network, the latency reaches 12.1ms when measured on the WiFi local network termination.

To summarize, PON technologies are capable today of addressing metaverse requirements as a backhaul solution or access network. In an end-to-end solution coupled with another access network, the latter would be the bottleneck. In highly congested networks, PON solutions can support L4S to maintain low latency.

4.5. Ecosystem aspects

Enabling metaverse applications without fully leveraging the network may result in non-optimal user experience. Developers will need to access networking parameters through APIs to enable various use-cases. It has been identified in [1] that capabilities, such as quality on demand, network information / network insights, positioning, cloud resources instantiation and provisioning, or congestion control mechanisms are needed to be exposed from the network. In general, intent-based APIs, such as ones defined by the CAMARA [7] project or 3GPP CAPIF/NEF can be used to address this demand, without necessarily having to deal with the specificities of a particular network type and operator's specific policies and configuration. However, it requires dealing with multiple operators, or connectivity providers, which becomes really complex for more global applications. For example, GSMA Open Gateway [8] has identified multiple ways of distributing and accessing APIs, including API roaming or aggregators.

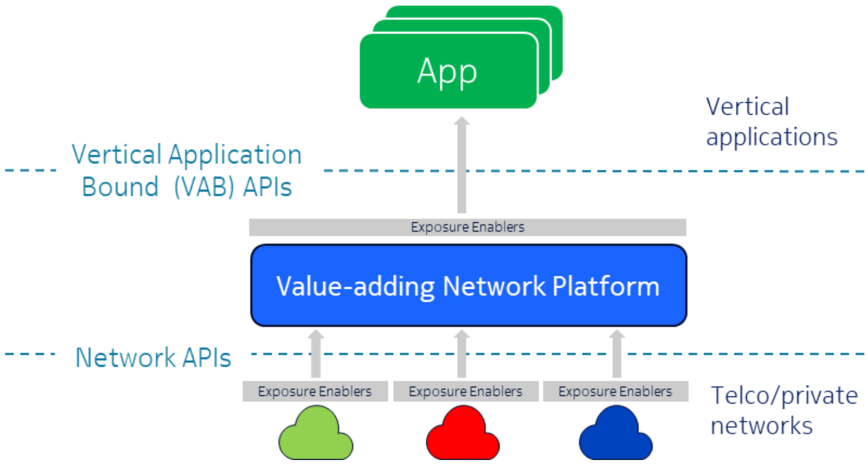

Figure 8: Illustration of Vertical Application Bound (VAB) APIs

In general, there is value in providing intuitive, on-demand, and elastic access to network resources, capabilities and analytics, control and data, while hiding the complexity of the telco capabilities, and opening the network for innovation. By introducing a Value-adding Network Platform, Telcos or Aggregators can provide vertical application bound (VAB) APIs that deliver simple and contextualized services, focusing on the desired vertical industry outcomes, as depicted in the Figure 8.

4.6. State of the Internet in 2024

This section is providing an analysis of the current state of the internet, based on , Ookla (the company behind Speedtest) public open data [9]. The current capacity of the open internet is evaluated against the requirements established in [1].

Ookla is a company specialized in web testing and network diagnostics. Through its flagship testing service "Speedtest.net" or "Speedtest by Ookla", it provides a free tool to test internet access, collecting data rates and latency statistics. Ookla is publicly releasing global internet speedtest data, and provide it on a quarterly basis through its *Ookla For Good*™ program, for fixed and mobile access. This is a public and open dataset made available under CC BY-NC-SA 4.0 terms.

The dataset is provided in both Shapefile and Apache Parquet formats, including per-tile access to the data, a tile being defined as a zoom level 16 web Mercator tile. This enables to access the data in a geographical oriented manner, e.g. per city, or country or continent. The attribute provided in the dataset, for every tile, are detailed in the table in Annex.

In our analysis, we're focusing on three continents: Europe, Americas, and Asia. For each of these regions, we extract the average value for the following attributes:

- avg_d_kbps, avg_u_kbps
- avg_lat_ms
- avg_lat_down_ms

- avg_lat_up_ms

The data are processed per quarter, and the performance of fixed and mobile networks are provided separately.

In the Figures below, we provide the results obtained for Mobile and Fixed access, respectively in Figure 9 and Figure 10. In these results, only the average values are provided, which includes testing on unloaded network, over nights for instance.

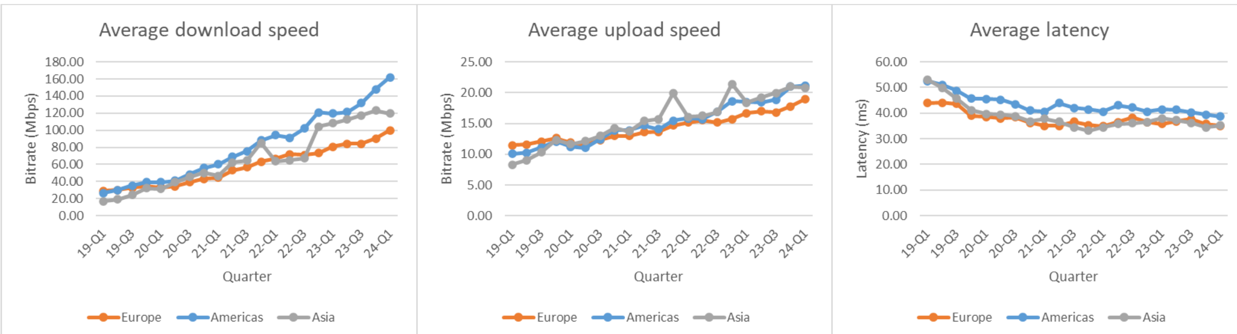

Figure 9: Average performance of internet over mobile network

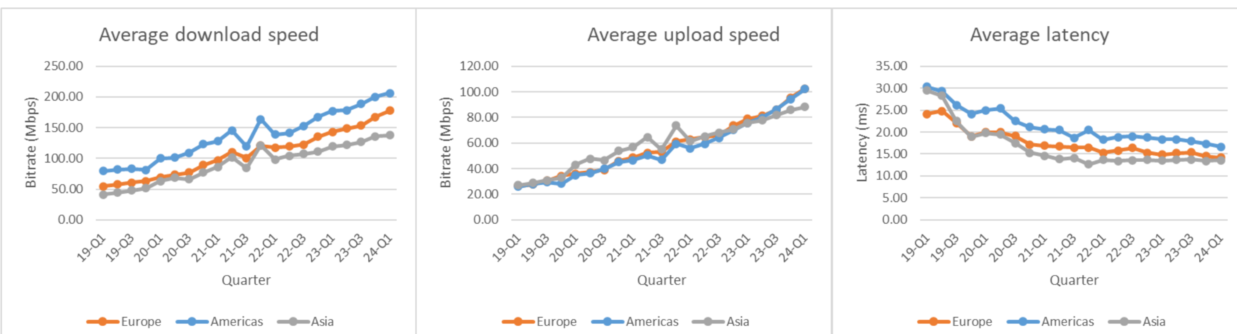

Figure 10: Average performance of internet over fixed network

From these results, we may conclude that today's state of the internet, on average, is sufficient to deliver most of the metaverse applications over the top, on both mobile and fixed infrastructure. In average, mobile networks may not be capable of serving most stringent and extreme case in terms of uplink throughput and latency but would be totally capable of handling most of the use-cases. However, those results do not reflect the state of the network when the connection is under maximum load. In the table below, the average latency measurement during upload and download are provided. These results simulate situations where you're accessing the network while someone is downloading a movie, a game or is in visio session in your home.

Table 4: Average latency under load (from 22-Q4 to 24-Q1)

			22-Q4	23-Q1	23-Q2	23-Q3	23-Q4	24-Q1
Europe	Fixed	lat. DL	276.94	265.91	279.29	263.31	250.92	254.83
		lat. UL	435.62	433.77	410.17	394.09	394.81	393.97
	Mobile	lat. DL	984.99	890.00	849.20	847.26	761.75	808.44

		lat. UL	1315.63	1261.53	1231.52	1272.65	1226.17	1153.40
Americas	Fixed	lat. DL	310.81	314.52	316.38	314.00	306.29	300.86
		lat. UL	401.64	399.21	403.71	405.80	403.51	404.33
	Mobile	lat. DL	848.64	850.34	862.97	851.61	854.79	900.07
		lat. UL	1071.31	1073.63	1096.28	1075.92	1034.39	1016.13
Asia	Fixed	lat. DL	283.41	282.34	279.55	268.64	262.21	263.70
		lat. UL	510.25	487.14	478.33	455.39	446.91	453.18
	Mobile	lat. DL	836.16	873.90	873.55	912.86	849.37	880.64
		lat. UL	1124.29	1162.82	1138.85	1143.64	1082.30	1116.85

When filtering out the data to loaded network conditions, it appears that both fixed and mobile networks provide too high latency to support metaverse applications, respectively around 300ms for fixed and 1s for mobile. This situation advocates for more closest interaction between the application and the network, and possibly for deploying applications closer to the user, at network's edge. Quality on demand, L4S, or 5G advanced features such as PDU set QoS handling can typically help overcome the identified limitations.

The closer you are from the datacentre, the lower the delay is expected to be. To compare the latency difference between low-density and high-density population area, additional results have been generated by state in the USA. Before having a look in the detailed results, the map below is showing the location of primary datacentres in the USA. It is thus expected that measurements in states close to those datacentres show better performance than rural states such as Wyoming, Colorado, Nebraska, Montana or North Dakota for instance.

Figure 11: Location of primary datacenters in the USA [10]

In terms of fixed reception, the results in Figure 12 highlight the disparity between higher and lower population density states. The distance from metropole areas where primary datacenters are located is visible in latency curves. While Wyoming is far away from dense urban area, South Dakota is in and intermediate location with most of the population located at the extreme east, so closer to dense urban areas where most of the big cloud provider datacenters are located. In terms of matching metaverse requirements, it is unlikely that agriculture or mining industry in Wyoming would get sufficiently lower latency to enable metaverse services deployment over the open internet. Although South Dakota might be seen as a bit remote location, it is close enough to datacenter to provide near sufficient latency performance.

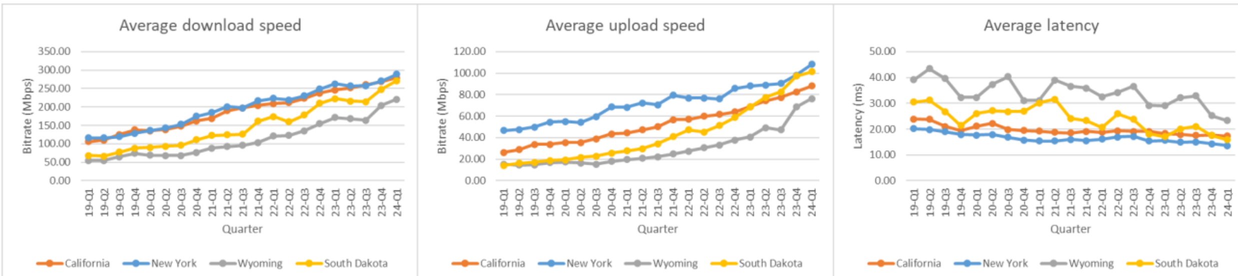

Figure 12: Average performance of internet over fixed network in US states

In terms of mobile network performance, the same disparity is observed, as depicted in Figure 13. However, the performance is low-density state such as Wyoming, with 60ms latency, 100Mbs and 15Mbps downlink and uplink speed respectively prevent any metaverse services to be deployed.

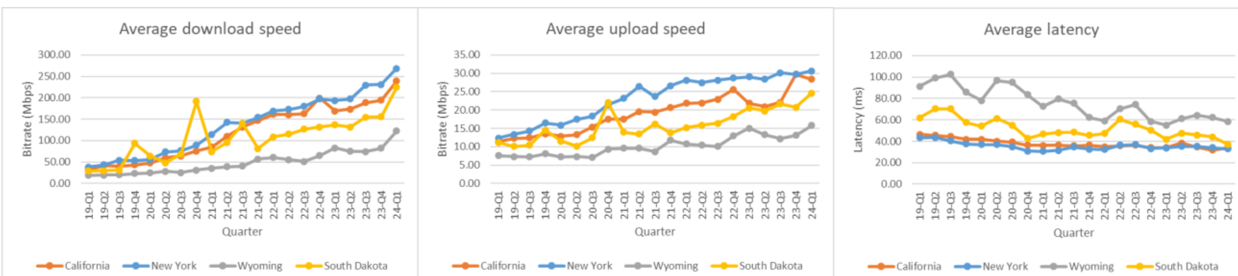

Figure 13: Average performance of internet over mobile network in US states

The disparity between dense urban areas and rural areas is highlighting that remote locations are not compatible today with next generation metaverse services. In terms of fixed networks, the download and upload performance are sufficient although latency is not low enough, advocating for edge deployments, traffic engineering and congestion control. Regarding mobile networks, high density states such as New-York or California provide sufficient performance in terms of bandwidth, but still without a low-enough latency to support most demanding metaverse applications. Rural and low-density states such as Wyoming would clearly require more 5G deployment to increase their performance while introducing edge server to lower the latency to an acceptable threshold.

Table 5: Example of European countries ready for most demanding metaverse applications (2024-Q1 mobile)

	DL	UL	Latency
North Macedonia	152.557	33.7201	21.0721
Denmark	239.484	30.7403	22.2969
Bulgaria	214.214	30.2402	22.9872
Switzerland	165.952	32.53	23.9475

Croatia	246.31	35.3198	24.0005

In Europe, although average latency measurements do not address most demanding metaverse applications, several countries are already capable of delivering expected network performance, providing low enough latency and high enough upload speed. Those are small countries with few locations concentrating most of the population.

To summarize, the following observations can be extracted from this study on the open Internet network capacity:

- Round trip latency for wireline access has dramatically dropped over the years to reach an average acceptable level for the identified applications needs (20ms)
- Uplink capacity for mobile access has improved but could be a limiting factor for some of most demanding applications. 3GPP has been working on features to enhance uplink capacity in Rel. 18, once they get deployed the uplink throughput should improve.
- Round trip latency for mobile networks is not consistent across countries. In some cases it may not offer the acceptable latency for some applications. Traffic management, e.g. L4S, would be super beneficial to master the round trip latencies and offer levels acceptable to applications. The deployment of edge cloud should also help in bringing the latency to acceptable level.
- During peak hours and under load conditions, neither wireline nor wireless access provide the sufficient performance needed by applications. Without traffic differentiation (e.g. PDU set features in 3GPP-Rel.18) or a mechanism to master latency (such as L4S), the current Internet cannot support metaverse applications under load conditions.
- 6G networks must have a significant emphasis on device to application low latency (not ultra low latency) and uplink to bring the performance level on par with wireline

5. Conclusion and recommendations

From a connectivity perspective, recent technologies can support deployment of isolated metaverse services. Fixed networks such as PON and Cable deliver sufficient throughput and latency to support such applications. Cellular networks defined by 3GPP, from 5G can cope with connectivity requirements, while Wi-Fi does not provide low enough latency. However, these networking technologies need additional features to support deployment at-scale. Traffic engineering is required, including flexible QoS handling and congestion control mechanisms. We've identified that 5GA includes a flexible set of features incorporating PDU set level QoS handling and prioritization, as well as L4S, enabling metaverse deployment at scale. It is noted that although Wi-Fi 6 under certain load conditions seems insufficient from a latency perspective, coupling it with L4S can drastically reduce the latency to around 10ms, making it a viable option for most of the metaverse services. Cable and PON can also benefit from L4S [6] [11], making them a strong option for at-scale deployment of metaverse services.

Furthermore, a study on Ookla Speedtest Intelligence® data highlights that current internet would support most of the metaverse applications over the top, during low-load traffic hours. However, it appears that deploying such services under loaded network is not possible advocating for quality on demand, traffic engineering and congestion control and deployment of applications at network's edge, especially in low density and rural areas.

Although the connectivity technologies seem quite mature and ready for future metaverse applications, we've identified potential gaps and limitations. First, future media delivery formats and protocols such as QUIC/HTTP3.0 introduce full encryption, making traffic engineering and prioritization potentially more complex. Second, we've seen that congestion control mechanisms such as L4S are not natively supported by Wi-Fi access networks, which as of today forces a tight and separated integration of L4S with the router. Some work is currently being done in the WiFi Alliance and IEEE 802.11 to get it integrated [12]. Last, but not least, we've highlighted that all applications require interaction with the network for achieving an acceptable user experience, e.g through exposure mechanisms. While plenty of APIs exist, this ecosystem might be complex to manage for the application developers that want to reach a global audience. The latter is where Aggregators, making APIs available across multiple telcos, can play an increasingly important role.