Three's response to Ofcom's call for evidence on Net Neutrality

Non-Confidential

2 November 2021

Executive Summary.

People have significantly changed the way they have used telecoms networks over time as we have moved from basic communications and voice calls to gaming and streaming very high-definition content.

As a result, the internet is no longer just a meeting point for end-to-end communication: it has transformed into a marketplace for the world's largest content providers.

To enable access to these providers, mobile operators have invested heavily in access networks, now providing 5G coverage and high-capacity backhaul to cope with ever-increasing data usage of these providers' services.

Further investment in new 5G technologies, and different network features such as edge computing and network slicing that allows for industry innovation, depends on there being a more permissive regulatory regime.

These new products will be of interest to content providers and enterprises, and mobile operators should be able to allocate them based on efficient means, which could mean commercial pricing.

This would not affect the quality of access to the general internet. If anything, improved traffic management measures will allow customers to have more reliable access to various types of more basic services as more data-heavy services are managed more appropriately.

Contents.

1. Introduction.

1. Introduction

1.1. This document sets out Three's response to Ofcom's call for evidence on net neutrality.

1.2. This response is structured as follows:

- First, we set out how the mobile internet market is changing and how content providers have changed their balance of control within the industry in the period since the original open internet debates took place, and how this may change further
- Secondly, we set out how the current rules are not fit for purpose today, or in the future as 5G-era technologies are being developed

1.3. We conclude on the timeliness of this review, and how we should seek to amend the regulations today to provide certainty for upcoming business cases.

2. The balance of control in the internet has shifted towards content providers

2. The original net neutrality debate is outdated in today's world

2.1. The original concerns leading to the creation of open internet regulations came out of a fear of the techniques that ISPs can exercise on restricting its customers' choice of the content they can access. In the earlier days of the debate, the telecoms industry was still managing voice and SMS traffic on switched networks, internet data use had only just started to emerge, but VoIP applications threatened the investments ISPs had to make into their physical networks.

2.2. The debate continued as video-on-demand services started to develop, and the type of ISP under the spotlight was the dual-play type who had both telecoms and broadcast services. Here it was assumed that ISPs would block or restrict the traffic of video-on-demand services to protect their broadcast business.

2.3. In both cases, such fears were broadly unfounded in the UK. The competitive nature of the ISP market has meant operators cannot go ahead and discriminate to favour their own business.

2.4. On the other hand, the content, application and services market has become more concentrated over time as large players have continued to acquire smaller innovative players and own a wide range of services. This must lead to a question as to whether open internet regulations – designed to regulate ISPs so that content providers can gain unfettered access to consumers – are suitable in today's age.

Data traffic continues to grow strongly, driven by a few large content providers

2.5. Customers are mainly using their mobile services for data in today's world, and this is reflected in the ever-increasing data usage patterns that we see.

Figure 1: Mobile data traffic (PB) by technology (monthly consumption)

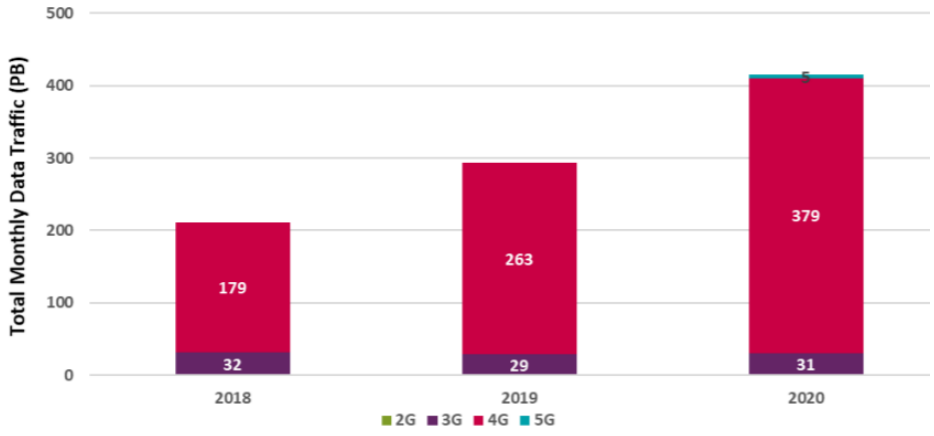

Source: Ofcom Connected Nations 2020

2.6. As reported in Ofcom's Connected Nations (2020), UK mobile data traffic continues to grow by about 40% each year. 70% of that traffic is estimated to be generated by a small group of providers.

Figure 2: Mobile Application Traffic Share, Europe

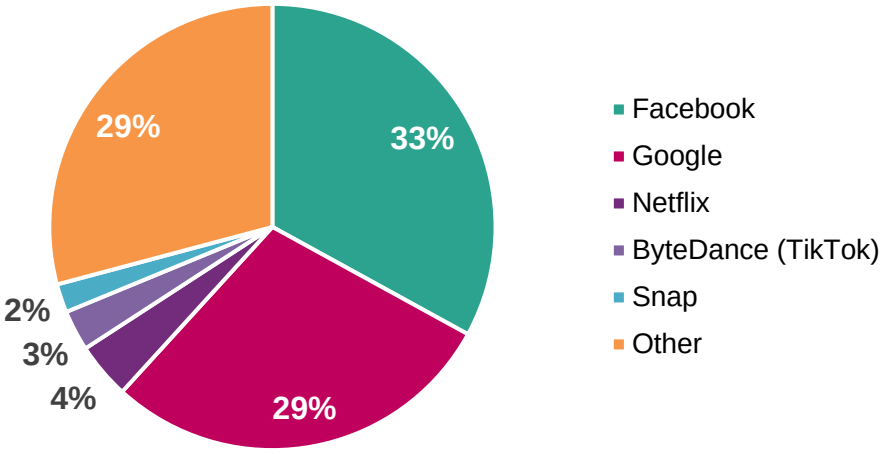

Source: Three analysis of data provided in Mobile Phenomena Report 1H 2020, Sandvine. Facebook includes: Facebook, Facebook Video, Instagram, Whatsapp; Google includes YouTube, Google Play

2.7. Such growth in mobile data needs considerable investment. Three has embarked on a £2 billion network enhancement programme across its entire network to cater for these future traffic trends, including an improved core network with 20 data centres, fibre-based transmission and RAN upgrades to LTE Advanced and 5G.

2.8. This investment is taking place at a time when customer prices continue to drop across the sector.

Figure 3: Weighted average monthly prices for average mobile use (excluding handset cost)

Source: Ofcom, Pricing Trends for Communication Services, 2021

2.9. Weighted average monthly prices have dropped from £14.53 in 2016 to under £11 in 2020. In the face of increasing data requirements and associated investments, the mobile market business model will become unsustainable.

2.10. In the meantime, the two content providers that have benefitted from mobile operators' network investment continue to record significant growth in revenues.

Figure 4: Google and Facebook revenues

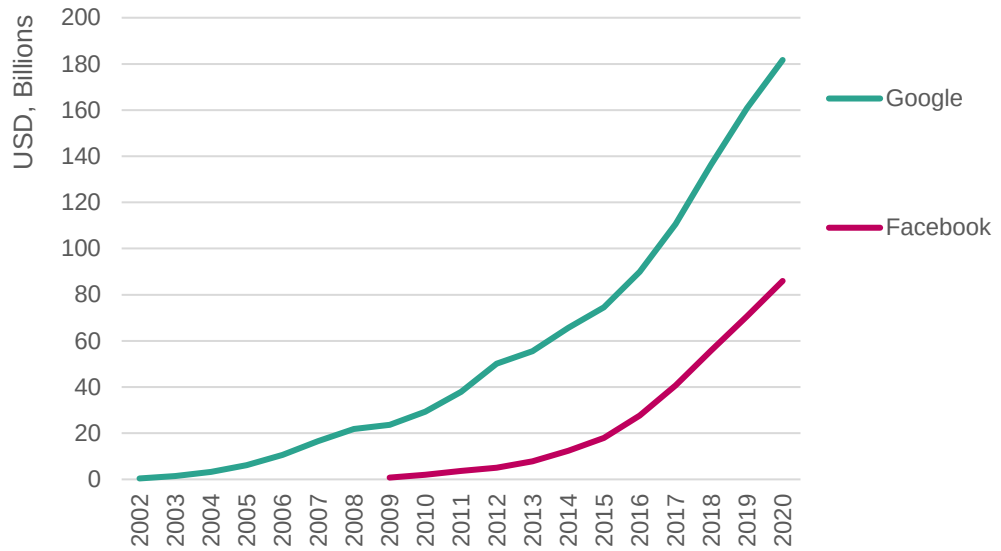

Source: Statista

2.11. To continue funding further network investments in the future, with continued pressure on customer pricing and revenues, operators must be able to explore other business models.

The business model of the mobile internet has already changed

2.12. The traditional assumption of customers' network access has already changed. For the most popular services, customers are no longer simply buying access to an ISP's network and being handed over via the internet in order to access content. Acknowledging that customer content is becoming larger, and needs to access customers more quickly on-demand, many content providers are engaging in different forms of network interconnect to position themselves deeper into the ISP network.

2.13. The traditional route of delivering content via transit providers is reducing in popularity due to the costs involved. Most internet traffic today instead involves some form of peering.

2.14. Today, with content being concentrated to a few providers as discussed above, the majority of internet traffic is delivered by direct peering to an ISP's data centre, and a large proportion of this traffic is also cached locally via Content Distribution Networks. There is a mix of content distribution being provided by third parties (e.g. Akamai), but also the development of in-house distribution networks (e.g. Netflix, Google).

2.15. Direct Peering and CDNs give an advantage over other content providers using internet exchanges or transit to reach ISP customers. Direct peering allows the content provider to better manage of traffic flows, and CDNs provide a lower latency, higher reliability and more efficient delivery of content (i.e. by avoiding duplicate transit of the same content to different customers).

2.16. However, direct peering and CDNs are only achievable at scale, with high traffic volumes. Unsurprisingly, it is the largest content providers who have been able to invest in these arrangements and therefore improve the quality of the content delivery to customers. Smaller content providers with lower traffic volumes must rely on public peering or transit to reach their customers.

2.17. We therefore already see a difference in the quality of access to the internet which, as it involves private interconnection where the onward connection is

managed by the content provider (and not the ISP), is therefore outside of the scope of net neutrality guidelines.

The 5G era will bring more changes to the telecoms business model

2.18. With content now being introduced to the edge of ISP networks, the next step in the evolution of the telecom business model is the passing of control of a share of the access network to enterprises and content providers.

2.19. This will predominantly be enabled by two technologies: network slicing which allows for differentiated services and layers of the network to be segmented and delivered over the same shared infrastructure; and edge computing which allows control of different parts of the network to be handled virtually and move applications closer to the end user or placed at enterprise premises. The combination of the two is compelling: if different parts of a network can be configured virtually, and a dedicated slice of the network can be segregated, then operators may be able to offer entire virtual networks to enterprises and content providers.

2.20. The potential for operators is how a business model can be created to offer these services. The current customer-oriented business model cannot suffice. It is difficult to envisage individual customers controlling or wanting to control network elements to deliver a service that they need. It would be more likely that a service provider will offer various packages that offer some level of customisation for the customer, but in a marketable way.

2.21. The second approach would be differentiated pricing to enterprises and content providers, with different levels of QoS set by the operator for categories of enterprise or application type (e.g. eHealth, automated transport, AR/VR applications etc.). This allows the operator to offer services upstream, without necessarily offering the highest level of customisation to the enterprise or content provider.

2.22. A third approach is for the operator to offer its entire network services via a platform, and then allow others to build a fully-customised network service according to the needs of the enterprises. Under this approach, the operator no longer controls the network in a traditional sense, and is simply offering underlying infrastructure and capacity to allow others to run their services.

2.23. In all cases, there is a clear trend towards a high level of customisation, a larger set of interactions that go beyond the end-user customer, and a retreat from network operators maintaining control over the whole network. This is the context that we recommend the UK's open internet regulations should be reviewed within, with any changes reflecting not just the state of the market today, but the potential for innovation that has yet to be developed.

3. Current rules are not fit for a 5G world

3. The current rules completely stifle innovation

3.1. The open internet regulations are designed for an era where ISPs control the entire infrastructure connecting people to the internet. As discussed in section 2 above, the infrastructure to connect to digital services is an increasingly complex environment, with a mixture of internet access, private cloud and public cloud applications. One thing is clear however: the ISPs control far less of this new environment.

3.2. In order to reflect the ever-increasing need for diverse and customised access to digital services, MNOs should be able to innovate in the 5G era to develop networks that are tailored to business' needs as well as for consumers' individual needs.

3.3. Unfortunately, such tailoring is difficult within the current open internet regulation and framework. Some might be able to interpret the rules to tolerate certain types of services to operate (e.g. network slicing). But we ultimately see it as just that – a tolerance under the current rules, under an interpretation that may not stand the test of time.

3.4. The ECJ's recent opinion[1] on zero-rating stands as the best example of how chilling an environment the framework can prove to be for investment: six years after the rules were put in place, and after the issuance of two full sets of BEREC guidelines, zero-rating has suddenly become effectively banned in the European Union on the basis of one judgment.

3.5. We think that this makes the clearest case for permissive regulation across the board: the rules must clearly state the conditions for which an activity is not permitted and therefore assume under all other scenarios that the ISP can carry on investing and developing propositions that can bring a new range of services for consumers and businesses.

Ofcom has managed zero-rating assessments well, but a permissive framework can reduce regulatory burden

3.6. Ofcom has not appeared to have blocked any zero-rating offer on the basis of the ability of the offer to have influenced end-users' exercise of rights. In some cases (e.g. Vodafone/VOXI Passes), it was the traffic management element of the offer (i.e. restricting the bandwidth available to the applications that Ofcom had concerns with).

3.7. The issue in this case is not to do with principles that Ofcom applies to the regulation, but the factors that operators must consider when developing zero-rating offers. Zero-rating offers are a commercial proposition designed to give customers a wider range of choices. For example, if a customer is particularly keen about social media, then zero-rating tariffs give the customer an opportunity to combine this with a lower volume data bundle for all other types of traffic. Alternatively, a customer who wants to keep their usage flexible between traffic

[1] https://curia.europa.eu/jcms/upload/docs/application/pdf/2021-09/cp210145en.pdf

types will continue to use a higher volume data bundle. We note that all MNOs in the UK market offer an unlimited data tariff and many others with very high data volume bundles.

3.8. This type of proposition is therefore beneficial to customers. However, the framework asks many questions of the offer to assess whether it influences the end-users' exercise of rights.

Figure 5: Ofcom framework for assessing Zero-rating offers

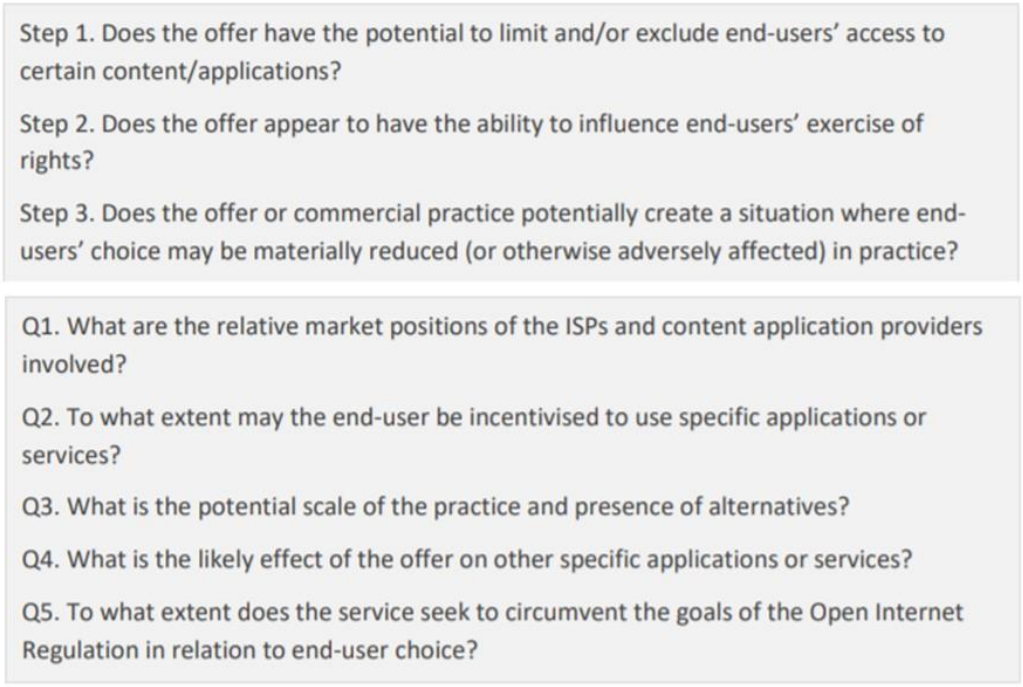

3.9. Many of the questions focus on the state of competition in the content services market. We consider that this is an unnecessary assessment in understanding whether end-users' exercise of rights have been limited. Open internet regulations are fundamentally about end-users' access to internet access services. To the extent that this access is unencumbered by zero-rating offers, then such offers should always be assumed to be permissible.

3.10. On the latter point, it is important to note that the understanding the zero-rating offers should cease to be active if a customer has completed their monthly data allowance is a flawed understanding of both the regulations and how customers use their services. Clearly, customers are able to choose a data allowance that is suitable for them, can top-up if they are a pay-as-you-go customer, or move on to using volume-based rates if they are pay-monthly. Customers are never blocked from using internet access services when their allowances finish – they simply pay a new/different tariff on the extra usage.

Regulations on device usage restricts choice for consumers in the market

3.11. Unlike for zero rating, Ofcom's investigations into traffic management practices has resulted in the closing of various congestion and other traffic management programmes by mobile ISPs.

3.12. Three was investigated in 2018 for traffic management practices on the following:

- Tethering restrictions: customers on Essential plans were unable to use their handset to connect to another device on the internet

- Traffic management whilst roaming: Three managed some categories of traffic such as video traffic, P2P and VPN when customers were roaming in the EU

3.13. Subsequent to this investigation, Three dropped the tethering restrictions on Essential plans and removed traffic management whilst roaming in the EU. However, the consequence of this decision was that the Essential Plan was no longer a commercially viable option and was subsequently discontinued.

3.14. That Three was forced by net neutrality regulations to remove the Essential Plan showed how incompatible the regulations are with enabling an innovative, diverse range of services for customers. The Essential Plan enabled customers with limited requirements from their mobile network (as only a minority of our customers were using tethering) to access a service that was priced accordingly to meet their needs.

3.15. By forcing tethering into the package by default, an Essential Plan with raised prices would have effectively seen these customers subsidising high-use customers (such as those using tethering domestically and high-volume services whilst roaming). We argue that this is not fair on those customers and crucially reduces the choice of tariffs available and restricted end-user choice in the market.

3.16. Tethering is not a mainstream service. According to Ofcom's Technology Tracker[2], approximately only 20% of households use tethering to access the internet at home (note that tethering is not the only option to access the internet for these respondents). This suggests that there is a majority of mobile customers who do not use the service and that there is, therefore, a great deal of cross-subsidisation by those who do not use tethering for those who do use tethering.

3.17. Furthermore, according to Ofcom's Technology Tracker, tethering tends to be used by younger people and those from the ABC social demographics. Older people and those in DE social demographics are significantly lower users of tethering and yet are the ones affected the most when tethering was removed as an additional option (for premium tariffs only).

3.18. Equally, different devices have different requirements from the network, and there is a need to tailor network access by device type. The most obvious example is for IoT devices, which generally require low bandwidth, high latency and a need to maintain the energy usage of the device. Likewise home routers need high bandwidth with high data usage, without the need to switch cells frequently. In both cases, the operator would be well-placed to provide the right SIM for these services to prevent the wrong SIM from providing a sub-standard service.

We cannot continue relying on goodwill approaches from content providers to manage our networks

3.19. Congestion management is required for operators to efficiently manage scare network resources in order to provide a good service to their customers. The current regulations set some building blocks to allow for this, but are ultimately not clear enough to allow for full traffic management to be put in place to allow for a smooth-running network.

3.20. The ISPs worked closely with the content providers during the COVID-19 lockdowns to manage network usage times, and the largest content providers agreed to voluntarily reduce video quality to manage the traffic load as people stayed at home.

[2] https://www.ofcom.org.uk/__data/assets/pdf_file/0022/216733/technology-tracker-2021-cati-omnibus-survey-digital-exclusion-data-tables.pdf

3.21. Whilst the voluntary action is appreciated, it highlights the imbalance in the power to manage network usage between content providers and ISPs. ISPs in this scenario would have been best-placed to manage the network between various providers. With effective traffic management, ISPs would be able to impose transparent measures for all content providers rather than relying on the goodwill of the largest firms.

The traffic management framework does not allow for managing congestion constraints

3.22. The key issue with the congestion management rules is the requirement for the policy to only be in place temporarily for temporary incidences of congestion.

3.23. This is because, on the one hand, congestion management can be based on *objectively different technical quality of service requirements for specific categories of traffic*, which we interpret as allowing mobile ISPs to manage traffic so that, for example, traffic requiring low latency is prioritised at peak time over those that do not.

3.24. But crucially, all traffic management must only be applied *only for as long as necessary* to prevent impending network congestion or mitigate effects of exceptional or temporary network congestion. This creates uncertainty on how long such a traffic management rule can be applied for.

3.25. The first question is on whether the mechanism to trigger the rule can be in place more than temporarily? Investing in traffic management programmes and rolling it across the network is a large exercise. Once developed, the ISP must be free to keep the policy in place permanently.

3.26. The second question is how many days can the trigger point be reached (and the management become active) before there is an interpretation that the traffic management is not being imposed *temporarily*? Traffic forecasts are consistently accelerating beyond the ability of ISPs to build out new networks. This means in the busiest hours, there will always be a breach of capacity and a requirement to manage traffic efficiently so that users are able to get a good service.

Figure 6: Monthly data usage per mobile user, UK

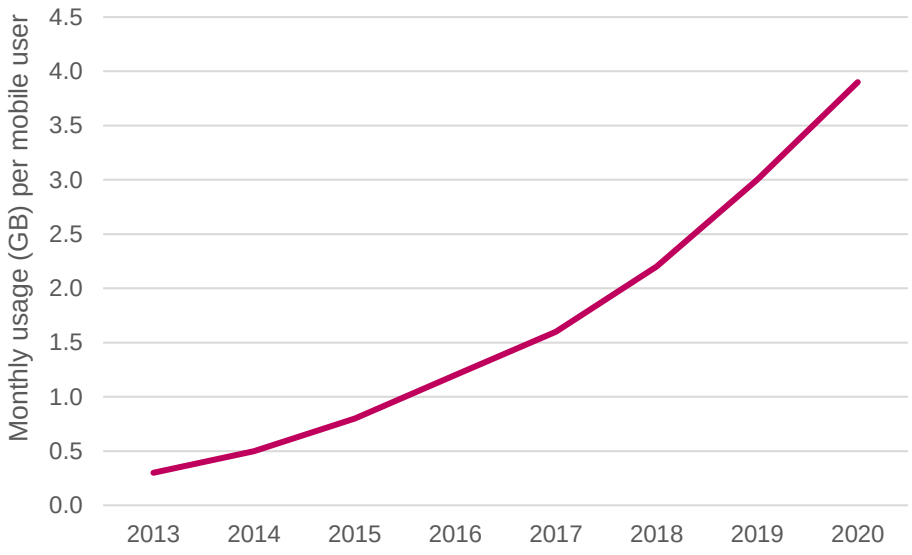

Source: Ofcom Pricing Trends for Communications Services, based on Ofcom Connected Nations data

3.27. It was suggested under the BEREC guidance that consistent traffic management should not be applied and instead operators should be investing to increase

capacity. As discussed above, Three has been investing heavily in its network as part of its £2 billion expansion programme. However, this will not be enough to avoid congestion in the busiest parts of the network during the busiest hours. This is because it is uneconomical to invest for such events. No other network industry is forced to manage capacity in such a way:

3.27.1. The transport sector manages its peak time capacity with higher prices during the times that most people commute.

3.27.2. The energy sector promoted off-peak tariffs to encourage usage outside of the busiest times.

3.28. We believe that there should be no time limit imposed on the ability to conduct traffic management. As long as there are rules which encourage transparency, there should be no reason to limit the policy.

We cannot rely on favourable specialised services decisions to invest in future technologies

3.29. The success of 5G will rely on the implementation of technologies that seek to allow mobile networks to be customised to the needs of consumers and businesses. These include a variety of services that can be enabled by Software Defined Networking (SDN) such as network slicing and Multi-Access Edge Computing (MEC).

3.30. There are two issues with how these developments can be safely interpreted within the current net neutrality regulations.

3.31. The first is technical: there has been ample discussion both by Ofcom and at the BEREC level on whether network slicing can be classed as a specialised service. The summary interpretations are positive (i.e. they suggest that it could be classed as a specialised service), but also make reference to the need to comply with the open internet regulations where network slices deliver to internet access services (IAS).

3.32. The rules on specialised services ask that there is no detrimental impact on IAS. In some respects, if network slicing reserves capacity for some providers or services, then IAS will always have a lower capacity than it previously did. However, that need not mean that the quality will be lower: in fact particularly demanding applications will likely be served in a separate network slice, away from IAS. In fact, enabling regulations to permit fully-fledged network slicing could help alleviate the congestion issues we discuss above.

3.33. The second is commercial: network slicing is unlikely to exclusively be a consumer end-user oriented model. There are applications where consumers might be interested in buying network access to a particular network slice (e.g. AR/VR gaming), but on the whole, demand for such services are likely to come from enterprise and content providers themselves. It is upstream businesses who will best benefit from the ability to access sliced offerings. In this scenario, therefore, operators must be able to charge businesses for access to the network slices.

3.34. We see that the rules currently offer two ways to interpret the ability to deploy network slicing:

3.34.1. **Specialised service route:** comes with the ability to commercialise the specialised service. But will be hard to prove that it is not detrimental to general internet access. Also subject to periodic review to determine whether services previously granted specialised service designation are no longer required because they can be provided by IAS instead. This latter point makes it

particularly hard to invest given the uncertainty of future regulatory interpretations.

3.34.2. **Traffic management route:** where there is no specific requirement to maintain same quality of access to IAS. However, all commercial relationships are banned and QoS differentiation is only permitted on broad categories of traffic. This would make it very hard to sell network slices fully to businesses who would buy a network slice to impose their own QoS arrangements over it (e.g. through the Network-as-a-service model).

3.35. We would propose amending the traffic management route to allow for network slicing and other future developments. The specialised service route creates too much uncertainty and does not allow future technologies to become widespread. The latter is due to the point that specialised services must not replace IAS at any point. What we envisage from future technologies is that they will enable the replacement of services delivered over IAS today (e.g. if there are provider-specific slices, gaming, video applications, etc.).

3.36. One good example of how the specialised service framework creates odd outcomes is the different treatment of IPTV and Video-on-demand services by European NRAs. IPTV, delivered using dedicated QoS characteristics over an ISP's network, is seen as an acceptable specialised service[3]. On the other hand, Video-on-demand is not seen as an acceptable specialised service as it is offered by many providers over IAS[4].

3.37. The outcome is unsatisfactory because the customer only sees the difference between the services as 'live' video and 'catch-up' video. Both type of video are often provided on the same platform: Amazon's predominantly VoD service has live sports channels included; BT's IPTV platform also has catch-up services in built. These take the form of full-length content, and news/sports bulletins of a much shorter length. There will therefore be inconsistent delivery across the platforms and within platforms because of net neutrality rules.

3.38. But the biggest problem with the decision is the binary nature of whether a service is an IAS-based service or a specialised service. It is clear that live TV services have higher demands for a specific quality of QoS compared to VoD. However, it is misleading to then conclude that VoD has no specific QoS requirements, and that it must only be accessed best-efforts over an ISP's network. There is clearly another level of QoS that would deliver good outcomes for consumers that is not being allowed by the current interpretation of the regulations. We therefore need rules that allow different levels of QoS for many different services.

3.39. However, the current traffic management framework is too restrictive to allow for this customisation. We believe that the framework and regulation should return to a simpler, permissive set of rules that protects access to the open internet, but allows for innovation in how services are delivered

3.40. Crucially, the regulations should allow for:

3.40.1. **An ability to commercialise the differentiation of QoS**: as set out above, all future technologies are based on the differentiation of QoS. A mobile operator will need to both understand how to allocate network resources and also earn fair returns on their investment in this technology. Therefore, there must be an acceptance that differentiated QoS can be commercialised so that the burden of investment does not simply fall to the end-user tariff.[5]

[3] See BEREC guidance, BEREC Guidelines on the Implementation of Open Internet Regulations, 2020, Paragraph 113
[4] See TKK's decision on A1's VoD service within A1TV, R 3/16 - A1 Telekom Austria AG | RTR
[5] This would, for example, require the removal of the phrase "and shall not be based on commercial considerations" in the traffic management clauses of the regulations.

3.40.2. **Permanent differentiation of QoS**: since the differentiation of QoS is being built into the business model, this must be allowed permanently, unlike in the traffic management rules where such differentiation is only allowed temporarily, or the specialised services rules where the differentiation can be reviewed over time. This is required for the certainty of both the operator and the businesses that rely on differentiated services[6]

3.40.3. **High degree of customisation of QoS**: some concepts of network slicing can be based on broad categories of industry verticals or applications, and therefore associated broad categories of QoS. However, to future-proof the regulations as business models develop, we think that they should allow for a high degree of customisation of QoS, perhaps at even the provider-application level[7].

The core of any revised approach to open internet should be based on transparency and the ability to switch between operators

3.41. We believe that the rules and guidance on traffic management should be adjusted to consider the core fundamentals.

3.42. The first key fundamental is transparency. We believe that transparency regulations are important in ensuring that customers are fully engaged with how their services are being managed. Three had previously set out its traffic management position in a key facts indicator template[8] which clearly set out traffic management rules and what services are affected by them, along with limits and hours of usage. We also set out our traffic management measures in our price guides and in the relevant pages on our website.

3.43. Armed with this information, customers should be able to choose between network providers to find one that is suited to their needs. The ability to do this has been improved through switching reforms (such as text-to-switch), the production of end-of-contract notifications and annual best tariff notices and the mandatory unlocking of handsets.

3.44. Finally, the ability to switch is also determined by the availability of other operators who are likely to differentiate their services to entice new customers. The UK retail mobile market is extremely competitive, with, in addition to the four MNOs, a large number of competing MVNOs.

[6] This would, for example, require the removal of the phrase “and shall not be maintained for longer than necessary” and “and only for as long as necessary” from the traffic management clauses of the regulations
[7] This would, for example, require the removal of the phrases “provided that equivalent categories of traffic are treated equally”
[8] Designed in collaboration across the Broadband Stakeholder Group

4. Conclusion

4. The time to enable innovation is now

4.1. Operators have been seeking certainty on future propositions and investments for some time. These investments are being made on real commercial business cases and require certain regulation to be implemented successfully.

4.2. It is therefore important that Ofcom is able to act quickly on amending the rules to clear the way for future investments to take place. In some cases, we recognise that Ofcom has a large amount of control (e.g. the zero-rating framework). On the other hand, we appreciate that Ofcom's ability to amend the traffic management framework is restricted somewhat by the regulations in place.

4.3. It might be tempting to use parts of the framework as a temporary solution in order to delay any decision making on changing the regulations. For instance, a clear decision on how network slicing and edge computing can deliver specialised services would help, but would need regulatory certainty that such decisions would not be put under review again during the investment lifecycle for these services.

4.4. We would therefore be happy to work closely with Ofcom to provide evidence to the Government on how to amend the traffic management regulations in order to foster a pro-investment and -innovation environment fit for the 5G era.