# Response to: Delivering the Broadband Universal Service

## Telzed reply to the Ofcom consultation

**February 2019**

File reference: Telzed response to Ofcom Broadband USO Consultation

12022019.docx

Details: Telzed response

Author: R Steele ( ✂ )



Telzed Limited UK. Consulting, Advisory and Interim Management Services
All correspondence: Lilac Cottage, Westerfield Road, Westerfield, Ipswich, IP6 9AJ, UK. Tel: +44 (0) 777 178 7607
Registered Office: Orchard House Park Lane, Reigate, RH3 8JX, UK

# Table of contents

# Executive summary

This paper has a limited focus: the definition of the broadband USO in the Ofcom consultation document “Delivering the Broadband Universal Service” 5 December 2018. This technical definition is referred to several times in the Ofcom document and implies that the Service Providers (SPs) must respond with proposals that meet this specification (service definition or definition). A central problem is that the definition is not sufficient. Much of it is sensible but it has some items that are missing, additionally the inclusion of a contention ratio definition has a number of problems: primarily, such definitions have little practical use and the values used contradict the need to provide downloads of 100Gbyte per month over a 10Mbit/s minimum speed service.

The definition ought to be clarified by the responding SPs and hopefully they will provide a new definition that exceeds the Ofcom definition and is fully practical. A danger is that a service could be provided that can be demonstrated to meet the definition, yet it will *not* be suitable in practice. Potentially Ofcom would have to evaluate the service offer based only on its own definition and could have problems rejecting or modifying the proposal, even if it is not *really* suitable, because the Ofcom definition has deficiencies.

This paper provides an enhanced service definition and explains why the Ofcom definition is insufficient. The issues with statistical traffic variances are discussed in a brief tutorial to help understand the problem and therefore why the definition needs to be enhanced.

The key problems with the definition are:

- The contention ratio. This is not a sensible factor to use, though it is often mentioned in network dimensioning discussions. Furthermore, the value of 50:1 is totally unsuitable. It should be removed
- The minimum speed of 10Mbit/s has no qualifications on its availability and likelihood of being attained. Real network services have statistical factors that must be considered
- The target values have no time related factors. This is a serious omission given that speeds and downloads change so rapidly each year. Average downloads increase tenfold approximately every ~7 years (Ofcom market data). Without corrections to the definition it will be insufficient in just a few years.

The inclusion of a download target per month of 100Gbyte/month is sensible. This value needs to be related to the network busy hour performance, and the value must change over time.

This paper is supplied to assist Ofcom and other parties. It has not been sponsored by any other party.

# 1 The Broadband USO technical definition

## 1.1 Ofcom USO definition

The required service definition is stated to be:

a) Download sync speed of at least 10 megabits per second

b) An upload sync speed of at least 1 megabit per second

c) A contention ratio of no higher than 50:1

d) Latency which is capable of allowing the end-user to make and receive voice calls over the connection effectively

e) The capability to allow data usage of at least 100 gigabytes per month.

This is mostly acceptable, as the values are minimum targets. The 100Gbyte value is not hugely less than the current UK fixed line average (240Gbyte in 2018). Hopefully SPs will provide a solution that exceeds the definition and includes technical design factors that are long term sustainable and meet consumer needs over time.

The absolute values of speed and download usage are open to considerable debate and it is expected that other consultation respondents will argue that the values (10Mbit/s and 100Gbyte) are too low. Given the average speeds already seen (~20+Mbit/s even over mobile and c46Mbit/s[1] on fixed) and the known huge growth rates in downloads (see Ofcom Connected Nations and market reports), the values are probably too low. Notice is drawn to the upload speed – something that is rarely discussed. Although downloads matter most for the majority of consumers, uploads matter for a number of applications and in particular for home workers. This value should be examined further and increased significantly. Interactive two-way communications requires close to symmetrical speeds. Home working requires uploads of large files.

## 1.2 Telzed USO definition

This paper does not provide precise final values for the USO. That is left to others and of course any value can be argued as to high or too low, depending on the party (customer or SP).

A revised USO definition is supplied below. This is discussed in the following section and the deficiencies in the Ofcom definition are explored there.

a) Download sync speed of at least 10 megabits per second

b) An upload sync speed of at least **5** megabits per second

[1] https://www.ispreview.co.uk/index.php/2018/05/ofcom-2018-study-average-home-broadband-speeds-hit-46-2mbps.html

c) Latency which is capable of allowing the end-user to make and receive voice calls over the connection effectively

d) The capability to allow data usage of at least 100 gigabytes per month

e) The download and upload speeds shall be obtained at least 9x.y% of the time based on random measurements during the Network Busy Hour

f) The service shall be Available 99.z% of the year (Service Availability)

g) The minimum download and upload speed [a) & b)] shall be updated every three years [TBC] by Ofcom/UK Government

h) The SP shall ensure that the service design is capable of upgrades to double the minimum download speed as a minimum, after three years

i) The SP shall ensure and demonstrate that the network design is capable of 10x increase in the download target [d)] every 7 years. The actual target download value shall be defined as per g)

j) The SP shall ensure the service design is capable of meeting reasonable demand growth in customer numbers. The assumptions for this must be provided and be passed by an Ofcom review. Such demand growth shall not degrade the other performance criteria of this definition

k) The download volumes [d)] shall be delivered in a manner that allows the speed requirement of 10Mbit/s [a)] to be met in the Network Busy Hour as per definition e) above

l) The design must be demonstrably capable of meeting the time of day demand variations and resulting Network Busy Hour usage as indicated by average fixed line broadband customers not obtaining the USO service, without failing to meet the above criteria. Therefore the download in d) is assumed to be delivered with the same time of day distribution of traffic of other non-USO customers.

The values (10M, 100G, x, y, z etc.) can be adjusted after further discussions with SP and network experts.

The changes are not major but provide a more solid basis.

## The following notes are relevant

Service Availability. This is the percentage time that a service is functionally usable. This is commonly used in service definitions. It is recommended that the value is the same as, or better than, the availability of PSTN fixed line rentals. There is no logical reason for new technology to work worse than traditional copper/PSTN. Data on this performance should be readily available. In some countries this may be a regulated specification.

A figure of 99.8% (for example) allows over 17 hour a year on average without service. This is unacceptable. So a higher value must be used. This does not mean that this outage happens every year. In reality there should be no failures in most years. This availability is an important factor.

Service penalties paid to customers when services fail can also be defined. These are not included above, but are recommended to be included based on the same commercial terms

used for non USO fixed line customers. Whether these penalties/targets are sufficient is not within the scope of this paper.

Network Busy Hour. This is the period (usually measured over an hour) when the weekly traffic is at its peak. The performance over this period is the critical value, not average measurements over a day or month – speed is normally excellent after midnight. The download usage in Network Busy Hour must be part of the specification. Speeds when downloading traffic must be satisfactory during the Network Busy Hour. A monthly average download speed target is not sufficient as most measurements would be outside the Busy Hour.

The changes ensure the service performance moves in line with the actual needs of average broadband customers. With the known fast changes in both speed and downloads, the existing Ofcom definition means it is not suitable after only a few years. Time-changing definitions must be introduced. Failure to do this will soon result in customer complaints and a repeat debate on why a percentage of customers cannot obtain a reasonable service. Their inclusion avoids the SPs complaining about the changes and SPs objecting to Ofcom "changing the rules" after a few years. By flagging that the targets will change, the SPs should not deploy a technology that is short term and unable to migrate to better services over time.

Ofcom must verify that the solutions are able to migrate to meet rising demands (speed and usage) as well as the customer-number changes. Therefore the SP proposed solutions must include technical and market demand assumptions to show that the solution is long term sustainable. Ofcom needs to be able to evaluate such proposals.

Clearly a problem exists as these additional requirements, even if sensible, are not in the current definition – SP responses might only meet the current definition. Ofcom and UK Government will need to sort out a final solution with the SPs. It is recommended that this is based on a revised USO definition, using these new proposals.

# 2 Discussion of the Ofcom and Telzed service definitions

## 2.1 The changes provide a better basis, but still build upon the existing Ofcom definitions

The first major change to the definition is the deletion of the contention ratio requirement. This is not necessary as other factors like: the speed; how often the speed is obtained in busy hour; the download per month and inclusion of time of day factors, cover the consumer needs. Contention ratios are defined in many places and are easy to understand. However they are often *not* used in many network designs. “Even” Wikipedia notes:

- “*…a stated contention ratio is that it is not, on its own, adequate for comparing services. There is a huge difference between 1000 users each on a 2Mbit/s service sharing a 40 Mbit/s pipe, and 50 users each on a 2 Mbit/s service sharing a 2 Mbit/s pipe. In the latter case two users trying to download at the same time means each get 50% of the speed. When there are a [sic] 1000 users, it would take 20 users maxing out their 2Mbit/s link at the same time to show any congestion. However both of these would be quoted as 50:1 contention.”*
- “*With the advent of ADSL2+… BT no longer work[s] on "contention ratio" as a planning rule*.”

As written, the Ofcom definition implies that a 10Mbit/s resource (say a mobile mast providing Fixed Wireless Access - FWA - broadband) could have 50 customers. This works and gives 10Mbit/*s so long as it is rare for more than one customer to use the service*. In reality this will not happen. Even a 20Mbit/s or 50Mbit/s mast can still have problems with just 50 customers and even more so with 100 or 250 customers as implied by a 50:1 contention ratio and 10Mbit/s target speeds. The slowing down of services is a fundamental aspect of all services that share a finite resource that has less than the capacity of all customers’ simultaneous use. The service is statistical. This has to be included in the definitions (see definition e) above).

The second major change is to ensure that the performance targets are related to actual changes seen elsewhere in the fixed broadband arena. The values used by Ofcom do not suggest they were defined with any serious headroom to account for the well-known growth in traffic downloads and in average speed. The values must rise with time. The Telzed definition provides a basis to enable the values to be updated. The values would typically be related to a known percentage of the average speeds and traffic usage levels per month seen by other fixed line services. This means the USO customers remain behind the average customer, but do not fall significantly behind. Given the fast changes, this linkage to trends is vital.

The upload speed has been revised upward. This could be the subject of wider Ofcom investigations as the asymmetric services in both fixed and mobile restrict a wide range of consumer behaviour. The reasons may not always be technical barriers that cannot be overcome with limited effort or investment. By restricting a wide class of symmetric services, this protects SPs from demand increases and so reduces investments, but at the expense of limiting the significant minority of customers who need a more symmetric service. Consider:

gaming; video interaction; healthcare meetings with a doctor; multi-users editing shared data/documents; database uploads; home publishing; home web sites; system back-up[2], etc.

Availability has been added. This perhaps is an almost trivial addition, but should be included. This can be enhanced to have Mean Time Between Failures or Mean Time To Repair and service penalties. Such details are left for Ofcom and the SPs to develop.

The addition of statistical measures of speed is important. No service speed can be absolutely specified if there are systems that have a finite capacity and the sum of all customers usage could in theory exceed that capacity. So fibre to the home might have 200Mbit/s direct point to point from the exchange site. But in the core network, routers and IP transit services have finite capacity. If every customer were to use services simultaneously it could slow down. As customer numbers rise to thousands or tens of thousands sharing the systems, the probability of this happening becomes small, so core networks can be reliably dimensioned with low chances of overload. Yet it is not an *absolutely* guaranteed speed.

In access to residential premises, there may well be systems (such as a radio mast or a cabinet system) that has backhaul to the exchange site with less capacity than the sum of: 10Mbit/s x number of customers. The systems could be overloaded if every customer uses the network simultaneously. This is not very different from PSTN and telephone calls. A variation of erlang type statistics and definitions can be used: there is a probability of a PSTN call being blocked (say <1%) in the busy hour. Broadband speed can slow down, but the network can, and should be dimensioned so that this is unlikely. The chances of slowing down can be defined accurately based on traffic statistics. Unlike phone calls, broadband is unlikely to be blocked at busy hour traffic peak times. The users each get a smaller share of the total capacity available – the familiar slow down of a service seen in mobile networks or when other users share a single broadband router's capacity in the same household.

## 2.2 Evaluations of SP responses need to appreciate network statistics

If SPs propose fibre to the customer in USO areas then both speed, monthly usage and busy hour performance are very likely to easily far exceed the Ofcom and Telzed values. As there are financial issue to deal with (otherwise there would not be a USO debate in the first place), this probably will not happen to every USO customer. Fibre to a cabinet or fibre to a mast with a mobile or FWA (Fixed Wireless Access) delivery solution, are likely. This is where there is now a system that has a finite capacity contended by customers. It is this *point of traffic concentration* that limits service.

As noted above, the contention ratio is insufficient to deal with this. The introduction of a target download per month (100Gbyte) as well as a speed value goes much of the way to ensuring adequate service. Additional Telzed factors above improve the service definition. Without these factors 10Mbit/s, or far more, may exist but at the busy hour it will slow down and not be practical to routinely download 100Gbyte/month, yet the service still meets the

[2] How do you back up or restore a typical PC/laptop/tablet with 200Gbyte of data over 10Mbit/s? Optimistically at least ~2 days or 20 days to upload over 1Mbit/s. The fastest way to transfer data is still by courier

Ofcom specification because the downloads could be forced to be done off-peak. This is unacceptable.

It is emphasised that the key cost driver that forces the mast or street systems' capacity up, is not the service speed, but the download. So, if a mast has 50Mbit/s capability then customers can get 50Mbit/s as the service speed. This can slow down if many customers use the same mast simultaneously. The probability of this happening is set by the download per customer, number of customers and traffic profile by time of day. These determine the traffic in the busy hour and if the total of this traffic exceeds the mast limit – the service must then slow down. Note that the same calculation is relevant to other concentration points such as a cabinet that could have many customers sharing a finite capacity.

The busy hour usage demand (1) is:

The average busy hour usage (Mbit/s) =

Telzed Factor x Download per month (Gbyte/month)

Where the Telzed Factor (TF) is ~0.01-0.035. This is simply defined by the time of day profile of average customers, and this shows how "peaked" the traffic is. This is country dependent. It differs by fixed line and mobile customers. It is probable that the TF value may increase over time with increased video usage in busy periods dominating the usage. The value can be defined using SP time of day profile data and can be updated.

This busy hour usage demand (1) defines the customer numbers (2):

The number of customers per mast =

mast capacity (Mbit/s)/ (The average busy hour usage per subscriber)

Formula (1) means 100Gbyte per month needs a few Mbit/s per subscriber on average. **This does not mean that 10Mbit/s is hugely over specified.** A user downloads, then uses the data, then downloads more. This physical download speed must be a good deal more than the average speed to assist with short term random traffic peaks.

So, if each customer on average needs 2Mbit/s, then a mast with 50Mbit/s can only have about 25 customers (TF = 0.02). This provides a simple evaluation method for submissions. In practice some additional headroom may be needed (fewer customers) to ensure good performance. This follows from:

- Lower customer numbers results in a higher variance in traffic levels. Masts or systems with 1000+ customers tend to have more predictable daily traffic peaks (from basic traffic statistics)
- Any system will need room for growth. If run at the design limit, a new mast or system must be ordered almost immediately
- The reality that mobile masts are rarely run at 100% of capacity in busy hour (perhaps 70-80% on average). If this is not done, then service slow-downs start to happen too frequently and customers may churn and/or complain. Values can be modified/confirmed by SPs.

Note how a 50Mbit/s mast actually gives close to 50Mbit/s service as the *physical service speed*. This is far above the 10Mbit/s target. Yet the customer only makes use of 2-3Mbit/s on average. With 25 customers, there is a fair chance more than 5 use the service simultaneously and hence do not get 10Mbit/s.

The mast could also be “split” to give 5 separate services each of 10Mbit/s. This *seems* to allow five customers (using 2Mbit/s average) per service or per 10Mbit/s “slice of capacity.” However it only needs two of the five to make simultaneous use to slow the service below 10Mbit/s. This traffic management is inefficient and therefore unlikely to be used[3]: it doubles the chance of slowing down[4]. It is better for all customers share the full 50Mbit/s, but there still remains a reasonable chance of less than 10Mbit/s.

These calculations mean that the headline service speed (50Mbit/s) may seem hugely better than the USO target, yet the service statistics still need to be considered. As traffic exceeds the mast capacity, users’ downloads *must* slow down, and they do not experience the 2-3Mbit/s average usage that they need to meet the 100GByte download per month target, when that target is adjusted for busy hour usage. The service definition and SP evaluation therefore need to consider the monthly download, the busy hour effects and how often traffic slows below the target. This is critical as the “50Mbit/s” service seems to far exceed the USO demand. However, *without* inclusion of requirements to not slow down *too often*, the definition is insufficient.

Furthermore, a blanket specification to always exceed 10Mbit/s (implied in the Ofcom definition) is also impossible: all shared systems are subject to random traffic factors and a SP ought to demand some statistical factors be included (unless fibre is used and there is no concentration point in the access link).

## 2.3 Setting the target speed or download usage

Meeting the download target of 100Gbyte reliably without frequent busy hour slow-downs is arguably more important than whether it is delivered at 10, 20 or 50Mbit/s physical peak speed. Any of these speeds can deliver the 2-3Mbit/s average needed. That said, as speeds >10Mbit/s are easy to deliver over FWA or over fixed lines, a 10Mbit/s minimum speed is still probably too slow a definition. It is far below actual current average speeds. This is supported by Ofcom market data that shows customers whose average fixed line speeds are less than ~10Mbit/s download considerably less per month than those on >50Mbit/s. This is due to:

- Frequent interruptions to streaming services
- The time taken to download large files of data. Users simply do not frequently attempt such activity
- Once users have >30Mbit/s the downloads do not rise significantly even with 100Mbit/s. Clearly the speed, once adequate, is not the key factor affecting customer behaviour
- Clearly a service speed below 10Mbit/s does significantly reduce the traffic per month and so it is restricting consumer behaviour.

[3] It is noted that many masts have three sectors due to spectrum management reasons, but this is ignored here for simplicity

[4] 2 in 5 versus 5 in 25 (=1 in 5) when sharing 50Mbit/s.

The 10Mbit/s definition implies that Ofcom is starting with a speed that is already known to limit existing customer behaviours. If this cannot be increased, at least the new definitions that ensure the values increases over time, will help to ensure USO customers are not further disadvantaged. The inclusion of time of day and statistical factors linked to the monthly download, will help to ensure a reasonable speed that is usable.

## 2.4 Summary messages on the probability of a slow-down of service speed

A recap on key points on the effect of overloading a mast or concentration point:

- A service speed of say 50Mbit/s vastly exceeds the 10Mbit/s target
- 100Gbyte per month usage implies customers on average use ~2Mbit/s at busy hour. But this is used in bursts, of often up to 50Mbit/s
- If there are many more than 25 customers or more than 100Gbyte/month is downloaded by each, then the service *must* slow down in the busy hour if they each behave normally
- With this additional demand, the customers will not see the average 2-3Mbit/s nor very often the 50Mbit/s peak speed. They will see or measure *average downloads* of just a Mbit/s or so (it will be less than 2Mbit/s) and yet they may still have peak speeds of anything up to 50Mbit/s – hence the need to make a selection of measurements using realistic traffic amounts, if verifying the speed.

The last point emphasises how consumer measurements of speed may be a problem. They can see up to 50Mbit/s most of the time. Only busy hour measurements really matter and values then could be below 2Mbit/s or up to 50Mbit/s. This makes analysis of values more complicated. What is the true average? It is possible that SPs will not directly do this measurement and will more simply define that that average usage across the busy hour, made by all subscribers on the mast, will not exceed say 70% of the maximum. This is easy to measure, but not by customers. Ofcom could agree to such measures or related guarantees of a low frequency of slow-downs. This might be included in the technical USO service definition.

It is emphasised that basic traffic theory and statistics should be appreciated and the service definitions modified to consider this. The SP proposals must be evaluated using such knowledge.

## 2.5 Conclusions

The Ofcom definition provides a reasonable basis for USO definitions. It can, and should, be improved to take account of statistical factors and the known trends to faster speed and greater downloads per month. Without such changes, the service definition will rapidly become inadequate.

The absolute speed value and other parameters in the new service definition are open to adjustment. In general: speed, downloads per month and quality factors such as probability of service speed slow downs can be improved.

The Telzed service definition is technology agnostic. However, any fixed line (i.e. fibre-based) solution ought to far exceed the specification, unless there are poorly selected cabinet-based traffic concentration systems. Therefore it becomes most relevant to FWA type solutions.

Broadband may also include a PSTN service, but this is not discussed here, as not strictly part of the Broadband USO. It is noted that PSTN line and Broadband are frequently bundled and their delivery has many common costs, so further study of bundled solutions may be needed.

It is not thought that modifying the service definition will be onerous and most SPs ought to be able to meet the revised requirements without significant effort or cost. The adjustments should reduce arguments and poor service to customers

# Appendix A Additional information and references

## About Telzed and the Author.

Roger Steele is an experienced management consultant in telecoms who has worked extensively on service level agreements and telco service definitions. This has included regulated reference interconnection and access offers. He is a chartered engineer.

He has recent experience of costing and dimensioning of mobile network models for data services. He has also seen how contention-ratio-derived values of traffic do not closely relate to actual traffic levels.

He has been an expert witness in legal cases relating to regulation, costing, pricing, interconnection, USO and access.

More details can be obtained on request or from the Telzed web site.

## Telzed Factor

The Telzed Factor is stated to be ~0.01-0.035. The values are approximate and depend on the busy hour profile. It is expected that a low volume traffic per month is likely to be more evenly distributed (low TF) than high volume traffic (higher TF). This needs confirmation by SPs. So a TF of 0.01 might be suitable for mobile data volumes of <10Gbyte/month, but fixed line usage >200Gbyte/month may need a TF of closer to ~0.035.

## References

The basics of this analysis are in the following Telzed papers, along with extensive wider discussions:

- A guide to understanding broadband usage
- Mobile cell site numbers with growing demand and higher capacity per site
- The need for speed
- Strategic issues for fixed and mobile broadband.

These are available on the Telzed web site www.telzed.com. These also have references to relevant source data from Ofcom.

## The analysis in this paper has wider relevance

An additional basic take-away from the analysis and the above formulae is that even 1Gbit/s masts do not enable many customers to have even the current UK fixed line downloads (240Gbyte/month in 2018). So mobile substitution of fixed lines, even with 5G, is unlikely – this is totally in line with the observations made by DCMS in its 5G strategy. This is not directly relevant to the USO discussion but there seems to be a belief in some circles that 5G can meet many demands. These might include coverage of rural areas (USO met by cheap/fast 5G FWA) and even substitution of many existing fixed line broadband lines. The latter is unlikely (from formulae 1 and 2), without many more masts than exist today.