Ethernet Tails QoS: The need for change

Submission to the Telecoms Access Review

03/10/2025

1. Summary

- Openreach is making this additional submission to Ofcom's Telecoms Access Review following recent market changes that have made the current Ethernet Upper Percentile Limit (UPL) Quality of Service (QoS) Standard (also known as Tails standard) no longer suitable. The changes mean the Standard now goes beyond Ofcom's stated objectives and has become unachievable for the 2026-2031 period.[1]
    - The UPL has changed a number of times since 2016. The current UPL was set in WFTMR 2021 based on Openreach's historic performance, market conditions at the time and Ofcom's judgement. The maximum permitted level of tails orders was linked to the size of the workstack and therefore was susceptible to changes in market conditions and Openreach processes.
    - Specifically, the volume of tails orders has not varied in the same proportion as the workstack. In recent periods, we have observed:
- The size of the open workstack has fallen because:
    - [✂]
    - We've increased the speed with which we can deliver orders [✂]
- The absolute volumes of Tails has remained relatively stable:
    - We continue to close Tails orders and work to improve our processes, however our ability to complete tails orders is restricted by the prevalence of civils, wayleaves and traffic management.

[1] Based on our current forecasts we also expect that we will fail the current UPL for 2025/26 and we will make further submissions on this in due course.

    - Tails which reach deadlock and are ultimately cancelled under the force majeure process are counted within the end-of-month UPL measures even if they are subsequently cancelled

- Accordingly, [✂] this will continue as a new long-term equilibrium, meaning the current measure is no longer appropriate.
- We request that Ofcom revise the measure such that it fits with current and future market conditions.
- We consider that the simplest change would be to amend the current standard by adding a safeguard for low levels of the workstack. We propose that the standard is set as the higher of **4.5%** of the open workstack or **600** open tail orders. This is equivalent to the current measure applying unless the open workstack is below 13,333 orders, in which case the limit is an absolute level of 600 orders.
- This will meet Ofcom obligations as set out in the Section 49 of the Communications Act 2003.

Contents

2. The current UPL QoS Standard and recent changes

Ofcom has set QoS Standards on Openreach in order to ensure it maintains a sufficient level of service for customers. One of these measures is the Upper Percentile Limit (UPL) standard, which limits the proportion of long duration orders.

There are two components to the UPL standard – the level of the workstack (the denominator) and the volume of open tails orders (the numerator). The level of the workstack has changed significantly recently whereas the volume of open tail orders has not. Tails have not reduced to the same extent due to a number of factors that limit Openreach's ability to make material improvements to their provision times, which are explored in the next two chapters.

Ofcom's objectives for QoS

2.1. Ofcom believes that the best means of delivering appropriate levels of Quality of Service (QoS) is through network competition.[2] Ofcom has set its WFTMR and proposed TAR regulatory frameworks to facilitate the development of such competition, where it is feasible. In Openreach's view the level of network competition for leased lines has grown substantially since 2021. This growing competitive pressure is borne out in a recent fall in demand. However, Ofcom has determined that while such competition develops, there is a need to impose regulatory standards for QoS to protect customers.

1.1. In relation to long duration orders Ofcom has stated its objectives as being *"[the UPL] is intended to protect customers with complex orders from excessively long lead times, by attempting to limit the number of orders experiencing such lead times."*[3] Further, in the 2021 WFTMR Ofcom recognised the improvements in service that Openreach had made in previous years and set out that its objective for the measure was now "*away from a push to drastically improve Openreach performance and towards maintaining an appropriate level of service for those customers with complex orders*".[4] Essentially Ofcom was stating that if Openreach performance continued at prevailing levels or better, then this was sufficient to protect consumers.

[2] TAR, Volume 5: Quality of Service, para 2.4.
[3] WFTMR, Volume 5: Quality of Service, para 4.13.
[4] WFTMR, Volume 5: Quality of Service, para 4.16.

2.2. Alongside this, Ofcom itself has acknowledged that there is a risk of regulatory failure in setting QoS Standards and a risk of setting QoS standards too high.[5] While this submission covers a broader concern about one of the standards, we encourage Ofcom to be mindful of the risk of regulatory failure if it was to maintain a standard that was not achievable (hence clearly disproportionate) or provided perverse incentives.

The current UPL measure

2.3. In-line with its objectives, Ofcom set QoS Standards on leased line access (LLA) and inter-exchange connectivity (IEC) services in order to ensure that Openreach delivers the QoS that Ofcom has determined customers need.[6] They apply to 'relevant ethernet services' which include our EAD, EBD, cablelink and dark fibre products. In this report, references to 'ethernet' mean to these 'relevant ethernet services'. The geographic areas covered are Leased Lines Access Area 2 (noting dark fibre is not available in the area), Leased Lines Access Area 3 but not in the High Network Reach area or Central London Area. IEC services are also included within QoS, except for circuits between BT+2 exchanges.

2.4. The standards cover a range of aspects of provision and repair. One of the provision standards is the UPL, designed to protect customers with the most complex orders. The current measure is defined as:

> *The Dominant Provider must ensure that the mean Monthly Upper Percentile Open Orders is no more than 4.5% in each Relevant Year.*[7]
>
> where:
>
> *'Monthly Upper Percentile Open Orders' means in relation to all Orders that were Accepted Orders but not Completed Orders by the end of the relevant month, the percentage of orders that had they become Completed Orders on the last day of the relevant month, would have had a Time to Provide of more than 133 Working Days.*[8]

[5] TAR, Vol 5, para 2.22.
[6] Wholesale Fixed Telecoms Market Review (WFMTR), Volumes 5 and 7.
[7] WFTMR, Volume 7, page 221.
[8] WFTMR, Volume 7, page 224.

2.5. The QoS standard deals with a small percentage of a what is already a relatively (e.g. compared to the volumes covered by the QoS standards for GEA-FTTC and MPF) small number of circuits in the first place.

2.6. The circuits covered by the Upper Percentile QoS standard tend to be those with the most complex delivery attributes, including contributing factors that are not fully within Openreach's ability to control such as civils, wayleaves and traffic management. [✂] of Tails orders have one of these factors affecting their provision.

2.7. Thus the measure is very susceptible to market fluctuations such as demand volatility or changes in the incidence of root causes of delay.

2.8. We are at a particular point of flux in the market with changes in both the competitive landscape and increasingly the types of technology that customers demand.

2.9. Therefore, as these market and internal conditions change it is necessary to review the measure.

Recent changes

2.10. Openreach has met the UPL Standard in recent years. However, in approximately the last 12 months, it has experienced changes that have put compliance with the measure at risk. Indeed, our forecasts show that we will miss the Standard for 2025/26.

1.2. Mathematically, the measure is made up of the denominator (the size of the workstack) and numerator (the number of tails). The changes we have experienced have led to an unprecedented and persistent sizeable reduction in the workstack but little/no impact on the number of tails, as shown in Figure 1.

Figure 1: Workstack and tails

[✂]

2.11. We now look at the root causes each of these changes in more detail.

3. Changes in the workstack

The workstack is a measure of all open orders that we have received and are in the process of delivering to customers.

The level of the workstack has historically been very stable. The workstack level is essentially determined by the rate at which we receive orders (intake) and the speed at which we make completions.

[✂]

[✂], we have continued to use our resource to work through the workstack to maintain completions. Indeed, in response to the competition underpinning the [✂], we have looked at ways to improve service further. We have undertaken process improvements to increase the speed at which we can complete certain orders. We have managed to improve service speed significantly under this process which has markedly improved our Mean Time To Provide (MTTP).

The result of [✂] combined with our enhanced completion speed has led to a dramatic fall in the workstack away from its historical norms.

This is not a temporary change. [✂] we have modelled a number of scenarios, which show that the workstack will remain lower than historical levels in future. Even if [✂], the faster completion rate, combined with this period where the workstack is already lower, means that the workstack will not return to previous levels.

3.1. The workstack is a measure of all open orders that we have received and are in the process of delivering to customers. In this report references to the workstack are to the

QoS applicable workstack, which is lower than the total workstack, since it excludes products and geographies where regulatory QoS standards do not apply.[9]

3.2. In this chapter we review the recent changes in our workstack, the reasons for these changes and our forecasts of future scenarios for the workstack.

3.3. The workstack itself can be described mathematically as:

$$Workstack_t = Workstack_{t-1} + Net\ Demand_t - Completions_t$$

3.4. Accordingly, the main drivers of the workstack are the level of demand that we receive and the completions that we make.

[✂] demand

3.5. [✂]

Figure 2: [✂]

[✂]

3.6. [✂]

Change in complexity mix of demand

3.7. [✂]

3.8. [✂]

3.9. [✂]

3.10. [✂]

[9] Geographically, orders for Area 2 and Area 3 access services and BT+1 and BT-only IEC services are within the QoS workstack, whereas HNR, CLA and BT+2 orders are not. Product-wise, EAD, EBD, dark fibre and cablelink are within the QoS workstack, but products like OSA, OSEA are not.

[✂]

3.11. [✂]

3.12. [✂]

Figure 3: [✂]
[✂]

3.13. [✂]

[✂] Area 2 orders

3.14. We have looked at differences between orders placed in Area 2 and Area 3. Orders placed in Area 2 will on average be closer to our network and will be easier to deliver, potentially because they require less infrastructure build.

3.15. The MTTP of delivering services within Area 2 and Area 3 is shown in Figure 4 below. This shows that there is a persistent difference between the two areas, with Area 2 being on average c. [✂] days to deliver faster than Area 3. These are both broad geographic areas containing a high number of volumes, so a difference of this magnitude across all these orders is indicative of a persistent difference in the complexity of these orders.

Figure 4: [✂]

[✂]

3.16. Further evidence that orders in Area 3 are generally more complex than orders in Area 2 comes from looking at Excess Construction Charges (ECCs). ECCs arise where the provision of a circuit requires additional work and cost to connect a customer site to our network. We recover these costs through ECCs.[10] Figure 5 below shows the average ECC (before accounting for the ECC threshold) attributable for orders in each of the

[10] If ECCs are above the £2,800 threshold then they are recovered as additional charges levied on top of our standard rental and connection charges, approved and accepted by the customer. If ECCs are below £2,800, then the amounts are not charged to customers for that individual circuit, but instead are summed up and added to a balancing charge which is recovered from connection charges.

Areas. It shows, that on average Area 2 orders incur around [✂] of ECCs, whereas Area 3 orders tend to average around [✂] of ECCs, albeit this value is more volatile.

Figure 5: Average ECCs in Area 2 and Area 3

[✂]

3.17. The implication of higher ECCs in Area 3 than Area 2 is that the nature of work in Area 3 is on average more complex, involving more build, than in Area 2.

3.18. Having examined the MTTP and ECCs within Area 2 and Area 3 and found them to indicate different levels of average complexity, we have next looked at how the demand changes split across the two Areas. Figure 6 shows that [✂], the area with comparatively less complex orders.

Figure 6: [✂]

[✂]

Completions

3.19. [✂]. As a response to the competition [✂] we have sought to improve customer experience, particularly our provision lead times, to remain attractive to customers. As a result, we have continued to complete new and existing orders but at a faster rate than before.

3.20. Those process improvements have allowed us to push down MTTP through delivering many orders more quickly. The key features of these service improvements include:

3.20.1. Leveraging the Full Fibre network to deliver service quickly;

3.20.2. Planners using new advanced planning tools;

3.20.3. Engineering expertise being applied earlier in the order journey

3.21. As a result of our ongoing completions and enhanced processes we have managed to push down the overall MTTP for completed orders to a level below 30 days (see Figure 7). This is a substantial improvement beyond already strong levels and is materially below the QoS Standard of 38 days.

Figure 7: Trend in MTTP

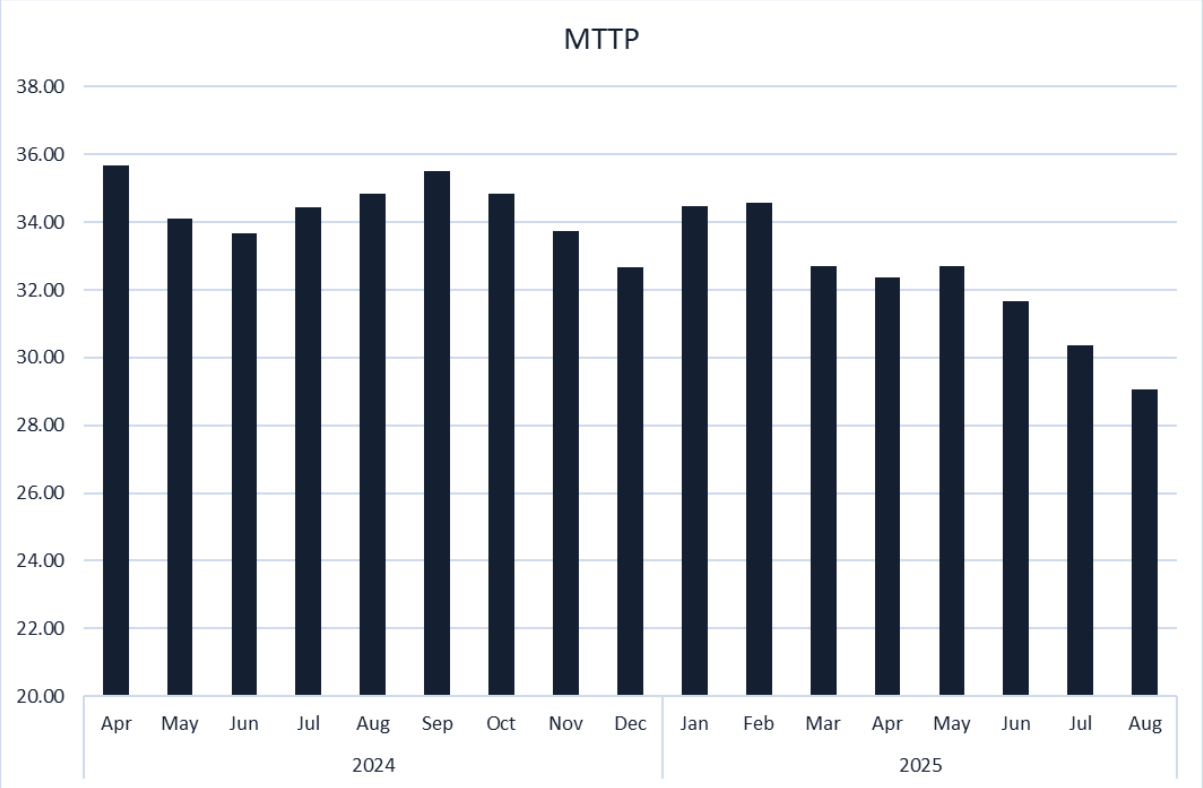

3.22. However, these particular improvements do not affect all orders equally. Indeed, measures such as involving engineering resource in a provision process at an earlier stage are more suited to improving provision lead times for the least complex orders.

3.23. We have not neglected seeking other types of improvements for more complex orders and we outline these in Section 4. Whilst this has improved the MTTP of these orders it hasn't resulted in a reduction in the absolute number of tail orders.

3.24. The mathematical consequence of increasing provision speed is to reduce the size of the open workstack, for any given level of intake. Provision speed also determines the relationship between the open workstack and closed orders. Another way of looking at the issue, is to recognise that where we are able to deliver orders in under a month, some of these orders will never be counted in the open workstack measure for the purposes of the UPL calculation. For example, an order received on the 2nd of the month completed in 10 working days, will never appear in the end of month open workstack because it is a closed order at that point.

3.25. One way to abstract from the increased rate of completions is to look at the metric on the basis of closed orders rather than the open workstack. Openreach reports this monthly to Ofcom[11] as KPI (e):

> KPI (e) - Time to provide (upper percentile) In relation to all Orders that became Completed Orders in the relevant month, the percentage of Completed Orders in respect of which the Time to Provide was more than 133 Working Days in each Relevant Year.

3.26. For example, in the latest available month (August 2025), Openreach's performance on the closed orders KPI was [✂][12] compared to [✂][13] on open orders. Accordingly, if both measures were compared to a threshold of 4.5% [✂].

3.27. To be clear, Openreach is not necessarily advocating for a return to the closed order measure. Rather, we are highlighting performance on this closed order KPI because it illustrates the impact of measuring tails using the open workstack during a period where the workstack is historically low. We recognise that there are other issues with a closed order metric, as there are with all individual QoS metrics. We comment further on these in section 5 below.

Impact on workstack

3.28. The factors set out above have had a material impact on the level of the open workstack. Figure 8 shows the QoS workstack since mid-2019. It was around 15,000 orders at the start of this period, the impact of COVID-19 lead to a decline to a previous low of 13,000, before the workstack picked up again remained steady between 14,000-16,000 for around 3 years until mid-2024, since then it has declined steadily to historic lows, reaching 11,000 by Sept 2025.

Figure 8: QoS workstacks

[✂]

[11] As part of Openreach's regulatory reporting and transparency obligations, Openreach provides Ofcom with a monthly set of KPIs as specified within the WFTMR. This includes KPI(e). See WFTMR, Vol7, page 233.
[12] As measured by KPI (e) for Relevant Ethernet Products within QoS applicable areas.
[13] As measured by KPI (h1) for Relevant Ethernet Products within QoS applicable areas, i.e. equivalent to the spot monthly value of the UPL QoS Standard.

Workstack forecast

3.29. The dramatic reduction in the workstack is unprecedented, and even exceeds the declines seen during COVID. [✂] we recognise that demand is always somewhat volatile and future demand is inherently uncertain, we do not believe that the workstack will recover to previous levels. This also reflects the success of the WFTMR in encouraging market entry and with the assets now literally in the ground, we consider this competition is here to stay and will only strengthen further.

3.30. Figure 9 below sets out our forecast for the workstack under different scenarios for future demand, which are themselves set out in

3.31. Figure 10. The four scenarios are as follows:

- As is – This was our base case forecast. [✂].
- MTO (Medium Term Outlook) – This forecast takes into account our revised demand forecast as of September 2025. Our forecasting team produces annual forecasts that feed into our Medium Term Plan (MTP). The regular cycle to refresh the MTP starts in late summer / autumn each year where performance against the current MTP and refreshed market insight are incorporated into an updated 2 year volume forecast. These refreshed forecasts, the Medium Term Outlook (MTO), are agreed and then processed into a fully-fledged set of volume, revenue and cost forecasts and a new MTP, reflecting more up to date market intelligence, which is generally authorised in following spring. The MTO is submitted to BT Group and forms the basis of their target setting for the Group for the year ahead. In September of this year we have produced a MTO containing a new set of forecasts for a two-year outlook, [✂]
- Uptick – [✂]
- Uptick (untrended) -This scenario is an artificial scenario created to show the trend required if the workstack is to return to the previous levels. It has been set to end at 13,500, which was the QoS workstack level at the start of the WFTMR period.

Figure 9: Workstack forecast under different scenarios

[✂]

Figure 10: Demand profile of different scenarios

[✂]

3.32. This scenario forecasting exercise demonstrates that in the most likely scenarios the workstack will settle at a new equilibrium somewhere in the region of [✂] orders. [✂]

3.33. Accordingly, we consider that it is necessary for a QoS standard that is linked to the workstack to be calibrated on workstack levels that are lower than they were at the start of the WFTMR period.

4. Stability of complex orders

The volume of orders in the tail has remained very stable.

[✂]

Our ability to use process improvements to speed up the delivery of tails orders is restricted because of the greater role of external factors that cause delays to tails orders.

The three most common causes of delay are civils, wayleaves and traffic management. Our ability to influence these measures is more restricted than for less complex orders.

We have worked to improve the delivery times of these orders and have made a number of improvements, particularly to the most extreme delivery times. While we have improved the delivery times of tails orders, for many of them it is not feasible to improve their delivery time to such an extent that they are not classified as tails orders. Accordingly, we have not seen a drop in the absolute volume of tails orders.

We continue to seek to improve processes and expect to make further reductions in tails delivery times. However, we don't think this will materially affect the volume of tails themselves.

Volumes of tail orders

4.1. The volume of open tail orders has remained very stable as shown in Figure 11 below. The monthly volume of open tails orders has been consistently between 500-600 orders.

Figure 11: Volume of orders in tails

[✂]

4.2. Unlike [✂], we have not seen any clear drop in the volume of tails orders. We recognise that due to the age of tails orders there is a degree of lag [✂] the volume of tails orders. Specifically, since a tail is a minimum of 133 working days old, changes in intake only feed through to changes in tails after at least six months. However, we do not believe

this lag effect is the primary driver of the stability of the tails volumes. Instead, we believe this is because [✂] is focused on the less complex orders (see previous section).

Composition of tail orders

4.3. The nature of delays on tails is such that they the majority of tails orders have delays associated with civils, wayleaves or traffic management. We set this out regularly within our bi-annual formal tails report submissions to Ofcom as part of our regulatory reporting obligations. Table 1 below shows the high proportion of open tails that have these types of delay. Around [✂] of orders have wayleaves delay.

Table 1: Analysis of delay on open tail orders as of 31 March 2025

[✂]

4.4. The high propensity for tails to be affected by these types of delay demonstrates the challenges for Openreach in making improvements to the delivery time of these services. We acknowledge that Ofcom has determined that these types of delay should be attributed to the 'Openreach clock' for the purposes of measuring QoS, because Ofcom is of the view that Openreach is best placed to manage these delay types. Nevertheless, the presence of these third-party factors inevitably makes it more challenging to make improvements to delivery times. We try to make such improvements but are limited by these factors.

4.5. We also note that Tails which reach deadlock and are ultimately cancelled under the force majeure process are counted within the end-of-month UPL measures even if they are subsequently cancelled

Improvements in tail order management

4.6. Where we can, we have worked to improve service performance for tails.

Outlier sprints

4.7. Since July 2024 Openreach undertook a targeted workstream that was focused on improving outliers.[14] We communicated this to Ofcom, OTA2 and CPs (see, for example, Figure 12 below presented to the OTA exec in March 2025). We've processed these orders differently, hand holding them and removing all barriers to completion. We put in place a specific team of individuals to proactively review orders as they approached CCD or became a failure risk and to also review orders earlier in the process so as to eradicate failure altogether. The team problem managed individual orders until completion or until they completed planning. They also undertook a root cause analysis on these orders, identifying the root causes (also see Figure 12) and allowing us to address some of the root causes. We sought to both reduce the number of days for these orders in the tail but also to reduce the need for customers to escalate.

Figure 12: [✂]

[✂]

4.8. As a result of our work on this Outlier sprints we brought down the total number of days of failure in the workstack by [✂] and saw a reduction in high level escalation volumes of [✂].

Impact within tails

4.9. The impact of our work on tails is borne out in the statistics about the tails themselves. Examining the age and distribution of tails orders, provides a greater degree of insight than simply looking at a binary measure of whether an order is greater than 133 days old or not (and hence whether or not it is classified as a tail).

4.10. Figure 13 shows the full distribution of tail durations. It shows a reduction in the mean duration of tails, from [✂] days at the start of the period to [✂] days in the latest period. This reduction is equivalent to the MTTP of tails being reduced. Further, the duration of the maximum length tail has also reduced from [✂] days to [✂] days.

Figure 13: Distribution of duration of tail orders

[14] We shared an overview of this at the Openreach/Ofcom/OTA2 quarterly QoS update in February 2025, slide 5.

[✂]

Future actions on tails

4.11. We continue to work with our dedicated team focusing on outliers.

4.12. [✂]

Table 2: Actions being undertaken in relation to Wayleaves

[✂]

Inability to materially change the volumes

4.13. Despite the improvements within the tail that we have outlined above, we have not seen a change in the volume of tails that is proportionate to changes in the workstack.

4.14. Fundamentally, the issues that affect wayleaves are long duration issues, and improvements that reduce the duration of tails (e.g. [✂] (see Figure 13), does not typically prevent a tail order from becoming a tail order.

4.15. Further, where these orders are delayed, additional resource will not usually expediate the process. Where an order is dependent on a third-party or CP next step (e.g. stuck in a wayleave deadlock or awaiting fee approvals), then additional resource is not able to speed-up the process.

Future trends will also tend to maintain tails volumes

4.16. In its TAR consultation, Ofcom proposed a cost-based charge control for low bandwidth services within Area 3. It proposed a glide-path down to cost, which would result in price reductions from the current levels. If Ofcom was to proceed with its proposed charge control, the likely consequence is that it will encourage volumes within Area 3, and especially in the more complex/harder to serve parts of Area 3, hence the balance between easier to provide and harder to provide orders in Area 3 will shift towards the former. By increasing the volumes of these types of provisions, there is a greater risk of tails orders being retained or even increased.

4.17. Ofcom's TAR consultation also proposed a tighter charge control on the Dark Fibre Access (DFA) services. These services are also only available within Area 3. One of the

most common uses of DFA is to connect mobile sites. Mobiles sites have a higher propensity to be in difficult to reach sites. Hence an increase in DFA use to mobile sites would also likely increase the volume of tails.

4.18. We consider that some of these most complex order types are network extensions (see volume 3 of our response to Ofcom's March 2025 TAR consultation). Accordingly, we consider that we could legitimately treat these as services outside of the scope of our regulatory obligations. Without a change to the tails measure, we will review how we treat orders that we consider to be network extensions, for the purposes of QoS.

5. A simple alternative

Given the factors outlined above the UPL is no longer suitable for the 2026-2031 market review period.

We consider that an alternative is needed and that relevant criteria for it are:

- it adequately protects customers by maintaining service standards in relation to ethernet tails;
- it is achievable and robust to changes in market conditions for the entire TAR market review period;
- it is as simple as possible to measure and implement.

Accordingly, we consider the best alternative would be to amend the current UPL into a hybrid approach, where the measure is set as the higher:

4.5% of the open workstack; or

600 open tail orders

Since this hybrid approach relies on two different measures, it could be monitored throughout a compliance year, but with the final end-of-year metric calculated on the basis of annual averages and whichever part of the hybrid measure that was appropriate.

This measure would ensure that customers of tails orders would be protected with service maintained at its current level. It would also ensure that the Standard meets the requirements of Section 49 of the Comms Act 2003.

Alternatively other plausible options would be:

- To estimate the relationship between workstack size and tails and adjust the percentage with the workstack size
- Set only an absolute level
- Turn to a KPI with ongoing monitoring

A track-record of different UPLs

5.1. The current UPL Standard has only been the relevant metric since 2021. Prior to this a Standard set on the basis of closed orders was in place. This has also changed between market reviews in respect of the threshold at which orders were classified as tails.

5.2. This illustrates that there are a range of options for alternative measures that have been used to achieve Ofcom's objectives. We note that a reversion to the closed orders standard would address the issues outlined above, particularly those stemming from faster throughput. However, the issues with the measure, including the perverse incentives it causes to not close tails, that were the basis for moving away from that metric, still remain. This reflects that there is no perfect measure for this QoS Standard, rather Ofcom has been trying to implement a practical measure that will protect customers and work in conjunction with the other QoS Standards.

5.3. Recognising this difficulty in setting a QoS measure, we are putting forward a proportionate and pragmatic measure that we think is suitable on the basis of the current market circumstances. Nevertheless, it is always possible that unexpected or material changes in market circumstances (for example, some of the factors set out in paragraph 4.18) may lead to a need for a further change, but our proposal is based on our best view of the current situation.

Criteria for assessment

5.4. We have assessed alternatives against the following criteria. The measure must:

- meet Ofcom's stated objectives in protecting against customers of complex orders;
- be achievable across the TAR 26-31 timeframe and as robust as possible to market changes;
- be simple to measure and implement; and
- Meet the criteria of the Section 49 test.

Preferred alternative

5.5. We have considered a number of alternative measures and have concluded that an appropriate approach would be a simple revision to the current UPL by adding a safeguard mechanism that applies when the workstack is low.[15]

5.6. Currently the UPL is defined exclusively in relation to the size of the open workstack. For the reasons set out above, we consider that this is not a viable approach when the open workstack is lower than historical levels, because the permitted percentage of tails falls to an unachievably low level.

5.7. Our proposal is to revise the UPL such that it is defined as the higher of 4.5% of the open workstack or an absolute number of orders. This would ensure that there is still a limit on the volume of tails orders and the customers remain protected. Customers would be no worse off than today with these protections.

5.8. Specifically, we would propose that the UPL is defined as:

> the volume of open orders which are older than 133 days of age must be no more than the higher of:
>
> - 4.5% of open orders; or
> - An absolute number of open orders

5.9. Figure 14 illustrates how this alternative approach to UPL would work. The bold line illustrates how the volume of permitted tails (y axis) would vary with the level of the workstack (x axis). At lower levels of the workstack the maximum permitted tails is set as 600, but once the workstack is at typical levels the 4.5% UPL applies, and the volume of permitted tails increases linearly.[16] The dashed line in the Figure illustrates the current level of permitted tails is 4.5% is applied to all levels of the workstack.

[15] We note that our proposal for a safeguard triggered at a certain threshold is in principle similar to our proposed safeguard for QoS Standards for copper-based WLA services. We have proposed that if copper repair or provision volumes fall below 0.25 million per quarter for two consecutive quarters that Ofcom review the standard. See, Ofcom, 'Telecoms Access Review 2026-31, Volume 5: Quality of Service, para 3.55', March 2025.
[16] The inflection point is at 13,333 workstack, since 13,333 x 4.5% = 600.

Figure 14: Alternative approach to UPL

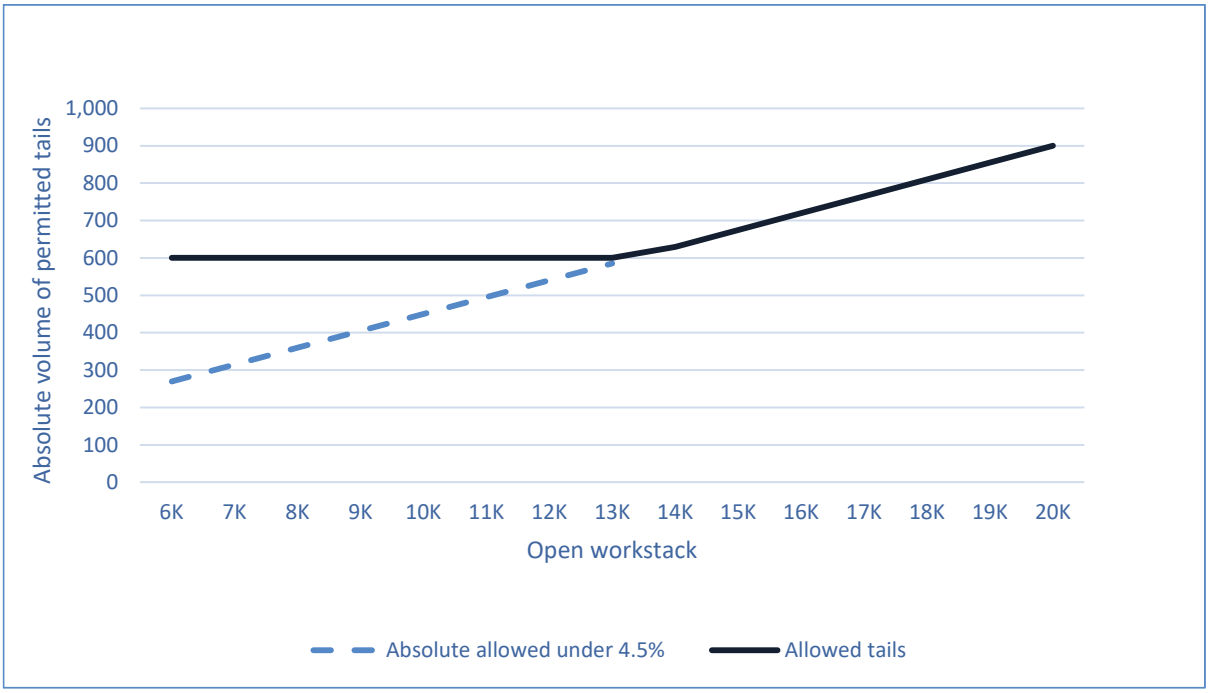

5.10. The absolute level of 600 has been chosen because:

5.10.1. It is consistent with levels of the workstack (c. 13,000) that were common at the start of the WFTMR period when the current UPL was set (see Figure 8). As set it is consistent with maintaining an equivalent level of absolute service, in-line with Ofcom's objectives.

5.10.2. We believe it is achievable over the course of the TAR period, recognising that tails have been consistently within the 500-600 range over the last few years (see Figure 11). While we would still be susceptible to other changes in the complexity of our intake, we believe we would be able to manage most of these.

5.11. Since this hybrid approach relies on two different measures, it could be monitored throughout a compliance year, but with the final end-of-year metric calculated on the basis of annual averages and whichever part of the hybrid measure that was appropriate.

5.12. This proposed approach will not lead to any softening of the QoS Standard. Figure 15 shows the historical levels of actual tails and permitted tails. At the start of the WFTMR period (April 2021) the level of permitted tails was around 600 – i.e. consistent with our proposal. As the workstack increased slightly the absolute level of permitted tails also grow and has been around [✂] for most of the WFTMR period. It is only since the recent declines in the workstack since mid-2024 that the permitted volume of tails has started to fall and now dropped below [✂].

Figure 15: [✂]

[✂]

5.13. Accordingly a hybrid approach with a safeguard at an absolute level of 600 tails would maintain QoS Standards at a level consistent with the WFTMR approach.

5.14. This measure would enable Ofcom to meet the requirements of its legal tests, as required by Section 49 of the Communications Act 2003. The tests require that measures that Ofcom sets by Direction are:

- objectively justifiable: the proposed measure will maintain protections for customers at a level similar to that as required by Ofcom when the current standard was set in 2021 and that Ofcom is seeking to maintain in the TAR consultation;
- not unduly discriminatory: it will ensure that all customers are equally protected.
- proportionate: it will protect customers at the level sought by Ofcom, without imposing on Openreach requirements that go beyond Ofcom's objective (see above under objectively justifiable) and is unachievable for Openreach (in light of the permanent changes explained in this submission).
- transparent: the measure is clear and is closely linked to the current measure.

Other options

5.15. While we consider our above proposal is both pragmatic and consistent with Ofcom's objectives, we note there are several other alternatives that could be implemented to make the measure suitable for the TAR period.

5.16. Alternatively other plausible options would be:

- A dynamic UPL percentage - estimate the relationship between workstack size and tails and adjust the UPL percentage with the workstack size
- Absolute level - set an absolute level for all values of the workstack. This would be similar to our primary proposal but without the hybrid approach. However, it is not a flexible measure and we consider it would be inappropriate, since it would not cater for large increases in the workstack
- KPI only - Turn to a KPI with ongoing monitoring. We consider there is merit in this option, as per our proposal in our WFTMR submission.[17]

5.17. While there are these alternative options, we consider that the simplest approach is the hybrid approach. It keeps the measure largely consistent and is easily understandable for industry.

[17] Openreach, Promoting competition and investment in fibre networks, 15 May 202, para 8.259