# Openreach analysis of additional factors impacting service costs in very high performance scenarios

Issue date: **14/11/13** Status: **For Information**

**The following information has been provided to EY by Openreach – as an input to the Discrete Event Simulation (DES) service cost modelling work.**

**This report addresses two specific factors which impact unit costs for the service delivery engineering workforce when addressing very high levels of required service performance:**

- **Factors affecting a minority of jobs in which successful completion within SLA lead-time is either physically impossible or extremely challenging (impacted by the "glass ceiling" limitations on Openreach's ability to successfully complete some types of engineering tasks on the day of issue to the field engineer);**
- **Factors causing increased task times due to local peaks in work volumes.**

## 1. Openreach analysis of the 'Glass Ceiling' for repair and provision service performance against SLA

Even with sufficient funding and a very much improved industry and Openreach ability to anticipate demand from unforeseeable events, the widely varying nature and environment of the UK network and the range of random incidents that can impact, there is still an upper threshold for the proportion of engineering tasks that Openreach can successfully complete on the day. Even if Openreach had limitless access to field engineering resource, a "glass ceiling" will always exist and will limit the percentage of tasks that can successfully be completed on the day the Openreach engineer is dispatched. There are a number of factors, some outside and others within Openreach's control, that prevent 100% task completion on the day[1]. The Openreach response to Ofcom's Fixed Access Market Review (FAMR) consultation describes this "glass ceiling effect" and establishes the level of this practical limitation at approximately 79.5% of repair tasks and 83.7% for provision.

The practical factors causing these limitations are not explicitly modelled in the discrete event simulation model. However, given the facility within the model to effectively limit the maximum achieved performance by GM patch by week, it is possible to introduce a factor to reflect the overall impact on modelled average performance. To do this, it is necessary to consider the impact of the glass ceiling limitation on the ability of Openreach to complete jobs within SLA lead times. Not all task failures will lead inevitably to an SLA failure.

[1] In this context, a "task" is the activity undertaken by an engineer on a specific visit. A "job", refers to either the fixing of a fault or the provision of a new product. Each job may require more than one individual engineering task to be issued, in order to successfully complete the job.

For most MPF repair tasks, where there is a Care Level 2 requirement to fix within the day following receipt of the fault report, a failure to complete the task on the day generally represents an immediate SLA failure. For WLR Care Level 1 repair tasks, in some cases where the first task fails to complete on the day it is issued to the field engineer, there may be one further day available, within which the job might be re-issued to a field engineer, and may still be successfully completed within the SLA (i.e. next working day +1). Also, some task failures are caused by factors within the domain of the Openreach customer – and hence may cause the SLA clock to be stopped for a period.

Openreach has carried out a further analysis of operational outcomes for field jobs. In this analysis, Openreach analysed a sample of jobs from 2012, considering those jobs which actually failed the lead-time SLA (as opposed to failures of individual tasks not successfully completed on the day of issue). For these jobs, Openreach identified a number of categories of the causes of such failures – and quantified each category as a proportion of total jobs performed.

The Discrete Event Simulation model, developed for Openreach by Ernst & Young, models data at a "job" level – i.e. the jobs that are modelled have total task times and lead-times which factor in a range of jobs, some of which are executed by multiple engineer visits. In order to factor into that model an appropriate level for the limitations imposed by the glass ceiling effects, the results of the study of the sample of jobs failing SLA and the nature of the different failure categories, were used to estimate the practical upper limit on Openreach patch performance.

**Methodology**

In this analysis, Openreach analysed all repair jobs, excluding SFI2[2] and BBB[3],that were issued to Service Delivery engineers in two separate weeks during 2012. These two weeks were chosen to reflect firstly a period of broadly "normal" levels of performance (week ending 12th July), and secondly the week ending 5th October where performance levels were under greater pressure from events such as poor weather. A total of over 130,000 jobs were covered by these samples.

The analysis identified whether each job met the relevant SLA repair lead-time, and if not, why it was not completed.

Openreach identified all the various categories that prevented successful completion within the SLA lead time and then reviewed which issues were within Openreach's current ability to address or improve, based on the information known to the engineer prior to attending the end user premises.

The table below details the types of activities attributed to each category.

**Detailed Definition of Repair Task Failed Establishing the "Glass Ceiling"**

| Repair Activity | Description |
|---|---|
| Completed | Job completed successfully. |
| Resource | Job failed due to lack of field resource available on time with right skills. |
| Engineer ran out of time | Engineer did not have sufficient time to complete the job. |

[2] Special Fault Investigation
[3] Broadband Boost

| | |
|---|---|
| Common fault found | Tasks that were allocated to an engineer but were subsequently identified as part of a common fault scenario that had not been identified when the fault was raised. |
| System | Openreach Operational Support System failure |
| MSO | Major Service Outage |
| Different skilled engineer | The fault was initially diagnosed as likely to be in a specific part of the network but was subsequently proved to lie elsewhere, such that the initially dispatched engineer did not have the necessary skills to complete and a different engineer had to be dispatched. |
| CP Issues | No Access - faults where the engineer has no access to the customer premises or where there is access to the premises but the customer or their representative are not available.<br>The customer is not ready for the work to be carried out.<br>There is no customer at the premises indicated.<br>Tasks that cannot be progressed as customer is querying service status. |
| Needs civil engineering | No line plant is available to resolve the fault and subsequently the task is passed to the planning department to design the solution.<br>Job requires a Planning Survey to assess work required to resolve an issue.<br>Second stage repair activity is required – for example duct work or maintenance dig type work.<br>Temporary service has been provided but additional work or stores are required to enable the fault to be cleared permanently.<br>Payphone faults needing contract work to resolve.<br>Includes a range of issues not covered by other codes – for example the need for traffic lights, water pumps, generators, gully suckers, tree cutting, roof ladders, or long ladders, etc. |
| MBORC[4] | MBORC declared and job failed target as a result |
| Job stopped as hoist required | Faults in the overhead network that cannot be progressed as additional access equipment and/or assistance is needed |
| Frames engineer required | Reports proved to the frame in the exchange where the engineer was not skilled for frames work.<br>Access to the frame was restricted and required a frames specific skilled engineer to be allocated. |
| Hazard | All faults that involve a safety related hazard – for example low wires, asbestos, medical reasons, hazardous environment, dangerous dog, abusive/violent customer, etc. |
| Specialist tools/stores required | Tasks that require specialist tools or stores to restore service that could not be identified prior to engineering visit. |

Each category was quantified – and classed either as a category which could be addressed hypothetically by increasing Openreach resource applied to the issue, or one which could not be addressed. Based on this analysis, a line was drawn – indicating the estimated proportion of jobs which are expected to fail the SLA lead time due to the underlying "glass ceiling" limiting the successful completion of the underlying tasks, due to factors not addressable by application of additional Openreach resource.

[4] Matters Beyond Our Reasonable Control

## Results

The following two waterfall diagrams show the results for each of the two sample weeks. Categories marked in blue reflect factors not addressable through an increase of Openreach resource. The green columns indicate the proportion of successful jobs within SLA, and categories marked in red represent factors where additional Openreach resource, if it had been available in the right location with the right tools and skills, may have been expected to succeed – but in fact failed.

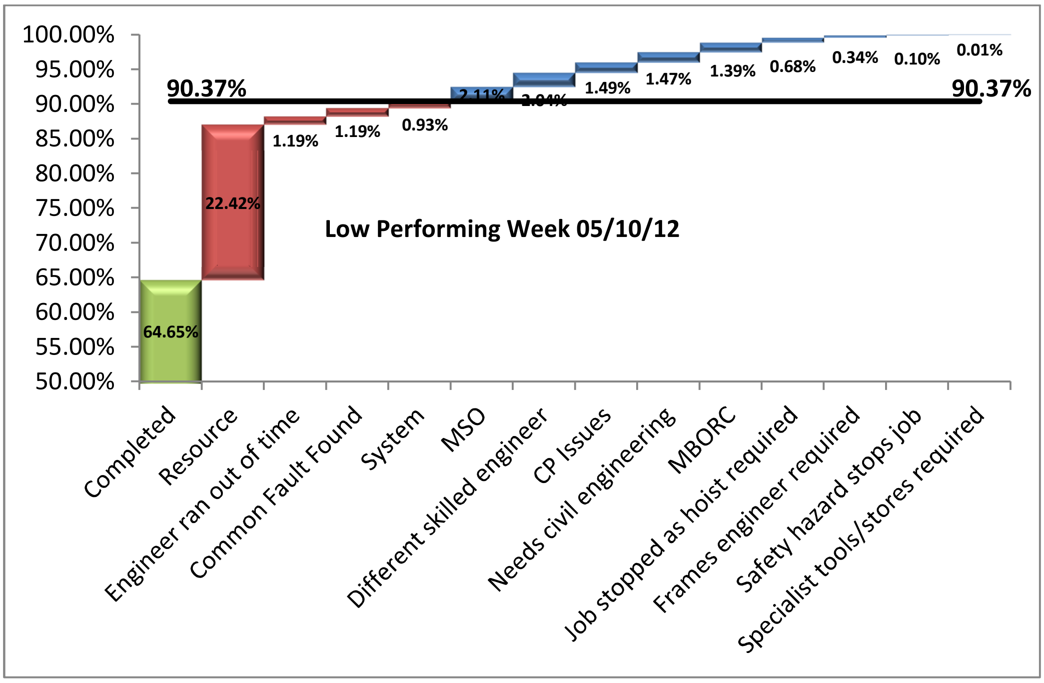


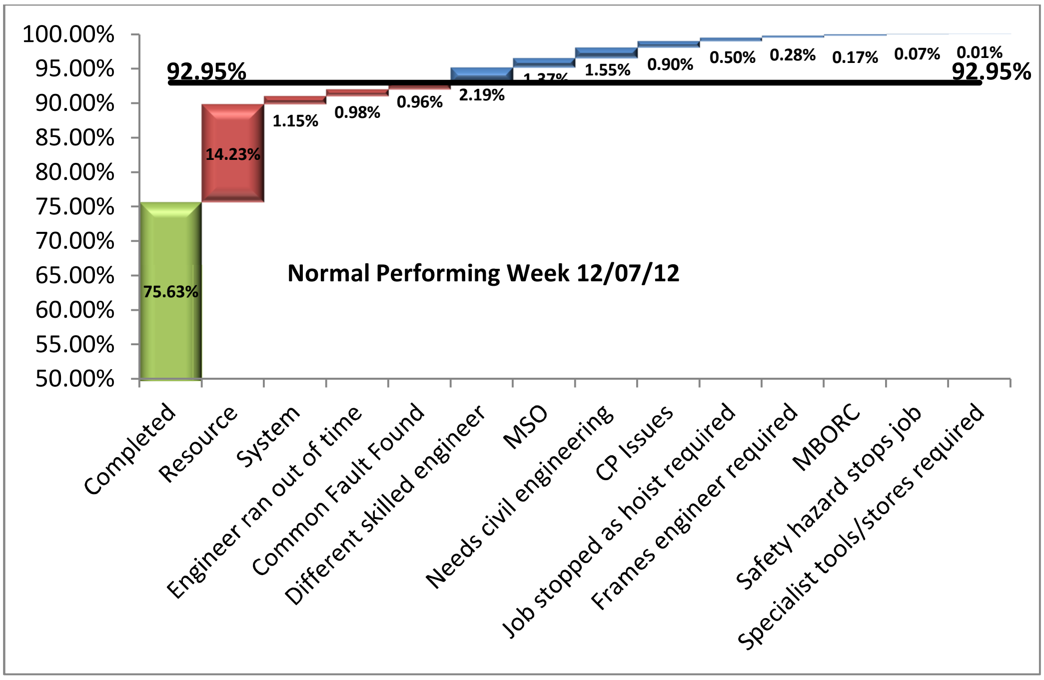


## Conclusions

Based on these samples, the waterfall diagrams above show estimates of the impacts that the glass ceiling may have on ability to achieve SLA success. These indicate that Openreach achievement of on-time repair during 2012 was effectively limited to a level of between approximately 90% and 93%. The various factors that determine this upper limit are shown in the charts above and explained more fully in the table.

It should be noted that this analysis was derived from actual performance during the year 2012 – when the mix of copper lines subject to Care Level 2 (rather than Care Level 1) was approximately 40%.

As Openreach has set out in the Service section of its FAMR response, there is a clear and increasing trend towards greater proportions of jobs being subject to Care Level 2, rather than Care Level 1 – as MPF line volumes grow at the expense of WLR. Between 2012 and the final year of Ofcom's new control period (2016/17) using Ofcom's own forecast data, the proportion of Care Level 2 compared to Care Level 1 lines will shift from roughly 40% to more than 50%. This will clearly impact on the level of the glass ceiling as it impacts Openreach ability to achieve the repair SLA. In broad terms, as the proportion of Care Level 2 jobs increases, it is reasonable to expect the impact of the glass ceiling to increase – so that the practical limitation on achievement of SLA success would reduce from the c. 90%-93% established above, towards the 79.5%-83.7% glass ceiling limitation on individual task completion (set out in Openreach's glass ceiling analysis submitted as part of its FAMR response).

Therefore, Openreach concludes from the above analysis that a figure of 90% for the glass ceiling factor in the DES model – representing the practical limitation on achievable performance against SLA for any patch over a given time period – is an appropriate level for this parameter, and represents a conservative forward-looking view of this limitation, factoring in a potential reduction over time in some of the volumes in the failure categories. A less conservative (and arguably more appropriate) figure to use for a forward-looking view of service costs in 2016/17 would be of the order of five percentage points lower than this – taking into account the greater forecast proportion of jobs requiring fix by close the following day.

## 2. Analysis of increased task time (including travel time) caused by local work volume peaks

This analysis covers two specific issues:

- Task time premium for unfamiliar engineers: an estimation of how much longer a task takes to execute (including travel time) when it is performed by an engineer working outside his/her normal working area, and

- Increased Minimum Service Level impact on “unfamiliar” job volumes: an estimation of how much the volume of jobs which need to be performed by an engineer working outside their normal working area would be expected to increase for a given increase in the overall average service performance level required.

Other relevant issues are also described and overall conclusions are drawn and informed by the analysis.

**Task time premium for unfamiliar engineers**

Openreach’s Service Delivery field force is organised into 9 General Manager (GM) areas. These areas are further split into a total of 58 Senior Operations Manager (SOM) patch areas. Each Openreach field engineer is typically assigned to a “preferred working area” of which there are approximately 430 across the UK – and in the first instance the engineer would typically be assigned jobs within that area in order to minimise travel time. This structure has evolved over many years to strike an efficient balance between the number of field engineers and the area they serve. However, on occasions, due to variations in local activity, and in particular as a result of high levels of demand, engineers may be required to travel outside their normal preferred working area – and occasionally outside their home SOM patch, and the evidence suggests that this can adversely affect their productivity.

Using historic Openreach data, an estimate was derived of how much longer the average task time has been for jobs executed by engineers working outside their normal preferred working area, compared with the overall average task times. In the first instance, jobs executed outside the engineer’s home SOM patch were analysed.

Data from Openreach’s work tracker system was analysed for the period from 01/09/2012 to 31/08/2013, using the following assumptions and simplifications:

- Only home and visiting direct labour was considered.

- “Home direct labour” are engineers in their home “Senior Operations Manager” (SOM) patch, “visiting direct labour” are engineers in a different SOM patch.

- Only field repair work was considered.

- Only non-assist, completed tasks were considered.

- Idle tasks, quick task returns and unusually short and long tasks were excluded.

- Analysis focused on main repair task types that account for more than 95% of all repair tasks; frames-related tasks were excluded.

- The most frequent repair task “type R1” was divided into sub-categories based on main Work Manager skill required.

- Task time was taken to be the time between issuing and completing a task (with adjustments made for lunch, etc.).

- Average task times for home and visiting direct labour are based on the same repair task type mix.

- End of Day travel was based on data as declared in engineer timesheets. The analysis compared average End of Day travel for field engineers working exclusively in their home SOM patch on a day with field engineers working exclusively in a foreign SOM patch.

- Increased travel between tasks was based on period 01/04/2013 to 31/08/2013, with data set identical to the one used for Task Time (above). This analysis considered the travel time differential between Care Level 3/4 repair tasks and Care Level 1/2 repair tasks as an indicator of the additional travel time required to react quickly to urgent tasks.

Openreach's assessment of the impact this may have on task times and the analysis of job-specific data shows that a job performed by an Openreach engineer outside their own SOM patch has a total task time of at least 16% greater than the average – including additional time "on the job" plus additional travel time specific to each job, plus additional start/end of day travel.

This is comprised, on average, of an additional 9.5 minutes total task time "on-site" (due to increased unfamiliarity with the location), an additional 2 minutes travel time between tasks, and an additional 9.1 minutes for start and end of day travel. Note that this figure for start and end of day travel indicates the typical or average time for the engineer to reach the "unfamiliar" location of around 4.5 minutes – from which we can reasonably deduce that most of the jobs performed by unfamiliar engineers are in adjacent patches and that historically, and on average, the travelling engineer is generally based close to the border of the unfamiliar patch. It would be reasonable to expect the overall premium on task times to increase the greater the volume of jobs that are required to be executed by unfamiliar engineers due to higher required minimum service levels. Therefore, we would expect task times for jobs executed by unfamiliar engineers to increase markedly above the 16% estimate derived from historic data under sustained periods of stress and high levels of high performance. One might also expect the frequency of such inter-SOM travel to increase under higher performance targets, particular as performance approaches the glass ceiling.

A similar factor would be expected to be encountered for jobs performed by engineers from the "home" SOM patch, but working outside the engineer's individual normal preferred working area. As noted above, typically an engineer's normal working area is significantly smaller than a whole SOM patch. This factor may be expected to be a lower premium than the inter-SOM unfamiliarity.

Additionally, as the volume of additional jobs required to be executed by an unfamiliar engineer increases, one would expect the additional jobs to be executed by engineers travelling greater and greater distances – and hence these jobs may be expected to have greater increases in travel times compared to the average durations quoted here.

**Increased Minimum Service Level impact on "unfamiliar" job volumes**

To estimate the impact of increased minimum service levels on the proportion of jobs which would be required to be completed by engineers working outside their normal working area (and hence expected to have longer task times including additional travel as described above) Openreach performed a simple analysis of job volumes by patch.

Firstly, for a given SOM patch, historical data was obtained for the daily repair job arrival volumes for each of Care Level 1 and Care Level 2 types. By applying a simple indicative distribution of execution lead times for these jobs, a daily profile of jobs executed was produced. This distribution was applied as follows:

- For Care Level 2 jobs, for X% performance requirement, X% of jobs received on day n were assumed to be completed on day n+1, with the remainder completed on day n+2;

- For Care Level 1 jobs, for X% performance requirement, X/2% of jobs received on day n were completed on each of days n+1 and n+2, with the remainder completed on day n+3.

This distribution was applied to the base daily profile of job arrivals, and a daily execution profile was produced – firstly for the base performance level (for 2012/13) of 63%, and another for one of the higher performance scenarios suggested by Ofcom at 85%.

For each of these scenarios, the volume of jobs executed on days over and above the mean jobs per day for that SOM patch was calculated as a percentage of total job volumes for the patch. On the simplifying assumption that each SOM patch has (on average) a number of engineers located within it equal to the required number to execute the mean number of jobs expected, this percentage indicates the proportion of jobs executed (on average) by engineers based outside the patch. For the base year 2012/13, the extended task times (including travel) are included within the base data that the DES model uses to calculate resource requirements. However, the impact of the additional proportion of such jobs required to meet the higher minimum service level (85% say) is not included in the task time assumptions for the DES model – and should be factored in explicitly.

From this analysis, the difference between the proportion of "unfamiliar" jobs (i.e. those performed by engineers outside their normal working areas) between the base case and the 85% case, is approximately 1% of total job volumes. Therefore, Openreach estimates that this additional volume of jobs would have to be executed by engineers from outside the home SOM patch compared to 12/13 baseline to meet an 85% performance scenario.

It is also important to consider the much higher volatility of such jobs when viewed at the level of each individual preferred work area. Openreach's FAMR Service response quantifies the volatility of work volumes for specific skill types at individual preferred work area level as at least 3 or 4 times the volatility of all work at total SOM patch level. Therefore Openreach considers that this same effect has a much greater impact on the volume of jobs performed outside the normal engineer working area but within the same SOM patch. Openreach therefore estimates that this would mean at least an extra 5% of all jobs would need to be performed by unfamiliar engineers for an 85% service performance scenario compared to the baseline.

**Other relevant factors**

In addition to the adverse impact on average task times caused by an increased requirement for engineers to execute jobs further away from normal working areas, there is likely to be another factor causing task times to increase as required service performance rises towards the upper theoretical and practical levels. This is caused by an increased tendency, as pressure increases, to divert resources towards greater numbers of individual jobs in jeopardy of SLA failure. In these circumstances there will be an increasing need to halt individual tasks part-way through so that the engineer can be diverted to another job which is in jeopardy. Effectively, as required performance levels rise, overall average task times will be likely to extend, and more jobs will eventually be completed by multiple field tasks. There is therefore an increased probability of difficulty in completing these marginal jobs.

Further, average task times increase the more the job has failed the SLA (i.e. typically it is a complex job requiring multiple visits which may fail the SLA). Therefore at higher performance levels more of these jobs will need to be completed, requiring more resource than the average and this needs to be reflected in higher task times applied in the model. The chart below illustrates that for jobs that have failed the SLA then task times are higher, and this increases (as would be expected) the further the task completion date moves from the original target. Hence to properly take account of task times at higher levels of performance there is a need to reflect in the analysis the fact that the average task time would be higher as the service organisation tries to deliver a greater proportion of the jobs in delay.

As an indication of this effect, Openreach considered the relationship between average task time, and the job completion time as compared to the SLA target date. The results are shown in the graph below:

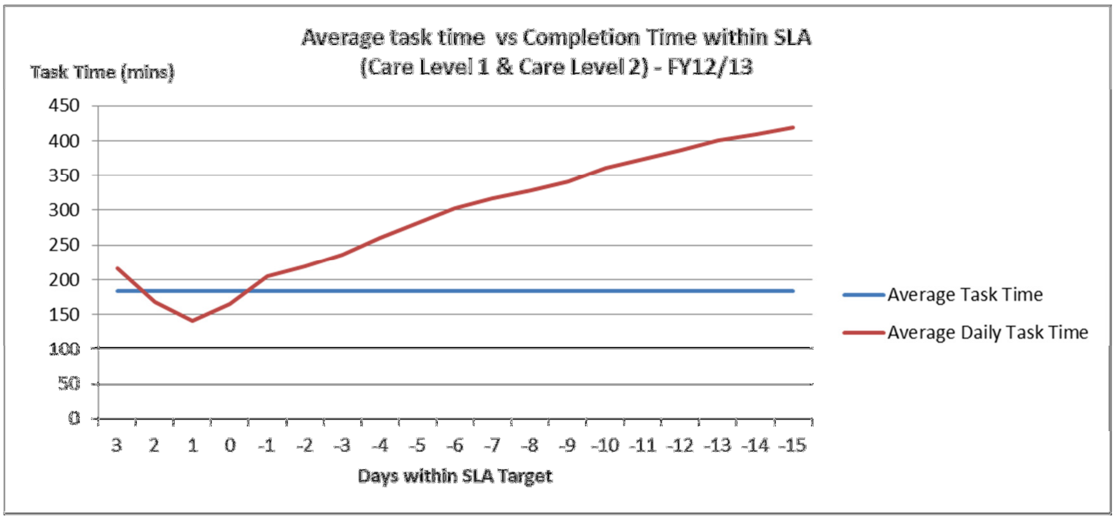


The graph illustrates that jobs which are more complex and take more time to complete, are inevitably more likely to fail the SLA target. Under a scenario where greater proportions of jobs are required to be executed within the SLA target, this will necessarily cause a greater number of these more complex jobs (i.e. those represented by the right hand part of the graph) to need to be executed within SLA - i.e. to move towards the left hand side of the graph. This effect will inevitably add to the size of the peaks of work volume to be executed – as there will be less freedom to schedule the more complex work over a longer period.

## Conclusions

The above analyses describe various different aspects of the tendency towards increasing average task times, as higher overall levels of performance are considered. The increased impact of these issues, as higher levels of performance are contemplated, are not reflected in the results of the DES model if task times are unadjusted (though baseline levels of the various issues are implicit within the model baseline data).

Considering all these issues together, Openreach concludes that under very high service performance scenarios (such as the 85% scenario suggested by Ofcom) the combined effect of all these issues would be that total task times would extend significantly, compared to those used implicitly in the DES model, especially for jobs executed on days where work volumes in specific areas are above average.

The amount by which task times for such tasks would be expected to increase is in the order of 20% or greater – due to the factors described in this paper. As this effect may not be expected to inflate task times across all tasks during the year, the question arises of what proportion of this increase should be reflected in relevant DES model runs.

Openreach estimates that if repair task times in the DES model run are inflated by between 5% and 10% for the 85% service performance scenarios, this will likely provide the best estimate of overall resource implications – including the issues described in this paper.

The same sets of issues would be expected in the case of provision tasks – as the impact on the Openreach Service Delivery workforce is likely in practice to be managed in terms of the total workload for the workforce (rather than individually and separately for each of repair and provision work). However, given the fact that provision work is largely appointed some days ahead, it may be appropriate to inflate provision task times for 85% performance scenarios by a lesser amount. So a range of between 0% and 5% uplift for provision would be most appropriate for these scenarios.