analysys mason

Report for Ofcom

Overview of the Quality-of-Service Model and its outputs for WLR/LLU Charge Control 2017

9 March 2017

James Allen, Mark Colville, Charles Khe

Ref: 2008936-102b

www.analysysmason.com

Contents

Analysys Mason Limited
St Giles Court
24 Castle Street
Cambridge CB3 0AJ
UK
Tel: +44 (0)1223 460600
cambridge@analysysmason.com
www.analysysmason.com
Registered in England No. 5177472

1 Executive summary

Ofcom is interested in better understanding the relationship between Openreach's performance in the completion of repair and provision jobs against specific quality of service (QoS) targets and the number of resources (engineers) that Openreach requires to carry out these jobs. For this purpose Ofcom has developed a high-level internal model to help quantify the relationship between QoS targets, performance levels against those targets, and the resource levels needed to achieve them, for various wholesale local access (WLA) products.[1]

Ofcom has commissioned Analysys Mason to review this in-house model and develop its level of usability and functionality in a number of areas, including incorporating real Openreach data on the volume and timing of historic repair and provision jobs (obtained by Ofcom under a Section 135 request). The scope of our assignment has also included producing outputs from the model for a range of scenarios of interest to Ofcom.

We understand that Ofcom is likely to use the outputs of the model in developing its upcoming consultation on the fixed access market. The model inevitably relies on some simplification of the real-world situation, and could potentially be refined further. However, in our view it is currently likely to provide a useful approach to addressing the problem at hand.

1.1 Model design

The model aims to calculate percentage deltas (i.e. percentage changes) in the number of full-time equivalent (FTE) resources needed to complete different percentages of repair and provision jobs within specified target times. These targets are set as follows:

- For *repair jobs*, target times for job completion are specified by service level agreements (SLA), which depend on the care level assigned to the product. Products with care level 1 (CL1) and care level 2 (CL2)[2] have been modelled.
- For appointed *provision jobs*,[3] two targets are specified:
  - The percentage of jobs for which the first available date (FAD) at which an appointment can be offered falls within a target number of days. The target can be set to either 12 days (current target) or 10 days (Ofcom's proposed target); and

1 These products include Wholesale Line Rental (WLR), Local Loop Unbundling (LLU) and Generic Ethernet Access for Fibre To The Cabinet (GEA-FTTC) products

2 To meet the SLA, CL2 repair jobs must be completed by the end of the next working day, while CL1 jobs must be completed by the end of the next working day plus one. Both MPF and WLR products can theoretically be designated CL1 or CL2, depending on the arrangement between the retailer and Openreach.

3 The model considers only appointed provision jobs and not un-appointed ones (those which do not require a customer appointment).

— The percentage of provision jobs completed by a contractual delivery date (CDD). Our modelling assumption is that the customer always accepts the offered appointment slot and so the modelled CDD is the same as the FAD, although in practice we recognise that it could be agreed to be on different date to the offered FAD. The CDD target is always met once the FAD target has been met for any job that does not fail at the first attempt. Conversely, where a job does fail[4], the CDD target is not met and thus the percentage of provision jobs that fail sets a 'glass ceiling' on the performance level that can be achieved against this target in the model.

In simple terms, the model reads-in data relating to job volumes and the time at which each job was logged, along with resource numbers and target definitions (in relation to how quickly jobs need to be appointed and completed). It then performs a range of calculations to schedule jobs and record the time between each job being logged and it being completed for repair jobs, and whether an appointment was available within a target number of days for provision jobs. The outputs are performance levels against targets (i.e. a percentage of jobs for which the relevant target is met).

As such, the model does not calculate the number of resources required to achieve a given performance level, but rather a performance level that can be achieved with a set number of resources. The number of resources required for a particular level of performance therefore has to be calculated by interpolating between pairs of scenarios generated using different numbers of resources.

Although the model runs the field-force simulation at a "Senior Operations Manager (SOM)" level, it aggregates the results up to a General Manager (GM) level to measure the performance level against targets in each GM area. Ofcom requires that a specified performance level is achieved within each GM area as well as nationally. As such, the performance level in the model is actually calculated as being the worst-performance level achieved in any of the 9 GM areas modelled[5] against each target.[6]

A summary of the model process flow is shown in Figure 1.1 below.

4 A small percentage of jobs fails because the customer or the RSP does not show up for the appointment. Such jobs are modelled, but are excluded from the measures above.

5 Due to a lack of suitable input data, we exclude the GM area for Northern Ireland

6 The performance level against each target is equal to the worst-performing GM area in relation to that target. It is possible that the worst-performing GM area may be different for the different targets.

Figure 1.1: Model flowchart [Source: Analysys Mason, 2017]

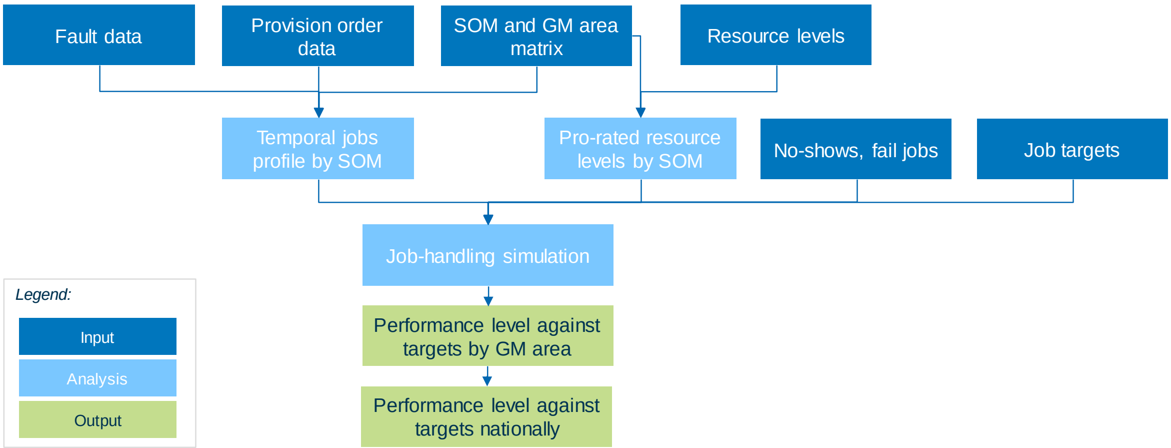

1.2 Limitations of the model

The functionality of the model has been kept relatively simple in order to produce results within the available time. In our view, the model provides a good high level approximation to quantifying the relationship between QoS and Openreach's resource requirements, whilst inevitably making some simplifying assumptions. We consider the model to be sufficiently robust for the purposes of Ofcom's consultation, with scope to be further refined in light of stakeholder feedback arising during the consultation process. However, because of the relatively simple nature of the model, it is subject to certain limitations. We identify some of the more significant limitations below and discuss them in more detail in Section 4.2.

- constant job duration (i.e. no variation in the time to complete a job) is assumed
- a constant volume of resources is assumed for each modelled scenario, with no day by day resource fluctuations (e.g. due to sickness, holidays, part-time working) other than reduced staffing on Saturdays (and zero on Sundays)
- resources are allocated across SOM regions according to historical fault rates and measures of working system size
- the allocation of resources between repair and provision jobs uses simple allocation rules rather than the type of complex real-time system Openreach is likely to employ when scheduling real jobs (e.g. there is no explicit consideration of travel time within the modelled SOM areas)
- the impact of engineers with different skill sets is not considered (all engineers are assumed capable of carrying out all jobs, notwithstanding that some jobs still fail at the first attempt)
- the assumed number of jobs which are "minor fails" has been modified (i.e. it is assumed that in future fewer jobs will fail at the first attempt)
- the model makes no allowance for jobs designated by Openreach as 'matters beyond our reasonable control' (MBORC).

1.3 Key results and conclusions

The model has been developed in order to run the scenarios requested by Ofcom, most notably in relation to varying both the minimum standards (i.e. performance level against targets) and the care-level volume mix for repair jobs.

We have tested the impact of these variations on resource requirements for Openreach in both 2014/15 and 2015/16 against a base case for each year, seeking a minimum standard of 80% provision FAD (for a 12-day FAD) and 80% repair within the SLA target.

- For 2014/15, the base case requires 7276 resources,
- while in 2015/16, a less benign year (in terms of weather, though also with different levels of demand for provisioning)[7], the corresponding number is 8182.

The results of our scenario runs for these base cases are shown for 2015/16 and 2014/15 in Figure 1.2 and Figure 1.3, respectively. For example, Figure 1.2 shows that relative to a base case with a 12-day FAD, a 0.6% increase in resources is required to achieve the same performance against a 10-day FAD. Similarly a 2.4% increase is required if all jobs are also assumed to be CL2 (and a 10 day FAD), for 80% performance.

[7] When the weather is calmer it is expected that less faults in the network arise and consequently there are fewer repair jobs, and vice versa in bad weather.

Figure 1.2: Percentage resource deltas by minimum standard option and care-level volume mix for 2015/16, relative to 12-day FAD, 80% attainment [Source: Analysys Mason, 2017]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
1 (80/80)	0.6%	0.0%	0.4%	0.4%	0.4%	2.4%
2 (90/93)	4.7%	2.2%	4.8%	5.8%	5.9%	8.0%

Figure 1.3: Percentage resource deltas by minimum standard option and care level volume mix for 2014/15 [Source: Analysys Mason, 2017]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
1 (80/80)	0.8%	-0.3%	0.2%	0.5%	0.9%	1.8%
2 (90/93)	4.6%	1.7%	6.0%	6.4%	6.5%	9.3%

Although the relationship is non-linear, there is a general trend that the required number of resources increases both as the proportion of repair jobs at care level 2 (CL2) increases and as the minimum standard for provision and repair is raised. The increase in minimum standard has slightly more impact on the required number of resources (and hence resource deltas) than the increases required to address higher proportions of CL2 jobs.

The observed trends are broadly similar between 2015/16 and 2014/15, though varying the care-level volume mix from 100% CL1 to 100% CL2 requires a somewhat higher resource delta for minimum standard option 2 in 2014/15.

A line of best fit could be used to help determine a 'gradient' to establish changes in the level of resources needed for a (small) unit change in care-level volume mix; the 100% CL1 case may be an outlier and might be excluded from this calculation. We note that some further investigation may be required to understand the behaviour in relation to the care-level volume mix more fully.

We have also tested (see Annex A) the variation in performance across GM areas and have observed that for high minimum standards (e.g. 90/93%) there is little variation, but for lower minimum standards (e.g. 80/80%) the worst-performing GM area may exhibit significantly lower-than-average performance. The volume of additional resource that Openreach will require to meet a set minimum standard in every GM area rather than just on the basis of a national average, will vary according to the level of the minimum standard.

2 Introduction

Ofcom has a duty to ensure effective communication services provision for UK consumers and citizens. As part of its work to fulfil this duty, Ofcom is currently reviewing:

- the remedies relating to Quality of Service (QoS) for Wholesale Line Rental (WLR) and Metallic Path Facilities (MPF) that were introduced for the first time in the last market review in June 2014;
- whether charge controls should be applied; and
- where relevant the appropriate levels of these charge controls for the period from April 2018

This review covers various wholesale local access (WLA) products, including WLR, Local Loop Unbundling (LLU) and Generic Ethernet Access for Fibre To The Cabinet (GEA-FTTC).

The charge controls exist as a remedy designed to set wholesale charges at an appropriate level where no realistic prospect of competition exists. To understand the cost associated with providing these products and propose appropriate QoS remedies, Ofcom needs to understand how changes to the QoS provided affect the efficient costs to Openreach of providing the relevant products.

For this purpose Ofcom has developed a high-level internal model to help quantify the relationship between QoS targets, performance levels against those targets, and the resource levels needed to achieve them. Ofcom has commissioned Analysys Mason to review this in-house model and develop its level of usability and functionality in a number of areas, including incorporating real Openreach data on the volume and timing of historic repair and provision jobs (obtained by Ofcom under a Section 135 request). The scope of our assignment has also included producing outputs from the model for a range of scenarios of interest to Ofcom.

We understand that Ofcom is likely to use the outputs of the model in developing its upcoming consultation on the fixed access market. The model inevitably relies on some simplification of the real-world situation, and could potentially be refined further. However, in our view it is currently likely to provide a useful approach to addressing the problem at hand.

The remainder of this document is laid out as follows:

- Section 3 provides an overview of the model design, discussing both the functionality of the model at the point we inherited it from Ofcom and the refinements we have made to it
- Section 4 summarises the limitations of the model and highlights areas of uncertainty, which may impact the results
- Section 5 introduces the scenarios we have been asked to consider and presents their outputs
- Section 6 discusses the conclusions suggested by the model outputs.

The report includes three annexes containing supplementary material:

- Annex A provides a further breakdown of some of the results presented in Section 5
- Annex B presents the results of some sensitivity tests performed on the model
- 0 includes a glossary of terms used throughout this report.

3 Overview of the model design

The model aims to calculate percentage deltas (i.e. percentage changes) in the number of full-time equivalent (FTE) resources needed to complete different percentages of repair and provision jobs within specified target times. These targets are set as follows:

- For *repair jobs*, target times for job completion are specified by service level agreements (SLA), which depend on the care level assigned to the product. Products with care level 1 (CL1) and care level 2 (CL2)[8] have been modelled.
- For appointed *provision jobs*,[9] two targets are specified:
  - — The percentage of times for which the first available date (FAD) at which an appointment can be offered falls within a target number of days. The target can be set to either 12 days (current target) or 10 days (Ofcom's proposed target); and
  - — The percentage of provision jobs completed by a contractual delivery date (CDD). Our modelling assumption is that the customer always accepts the offered appointment slot and so the modelled CDD is the same as the FAD, although in practice we recognise that it could be agreed to be on different date to the offered FAD. The CDD target is always met once the FAD target has been met, for any job that does not fail at the first attempt. Conversely, where a job does fail[10], the CDD target is not met and thus the percentage of provision jobs that fail sets a 'glass ceiling' on the performance level that can be achieved against this target in the model.

In this section we provide an overview of the model design and the nature of the calculations undertaken in order to achieve the modelling aims.

As described above, the model was initially built by Ofcom and provided an end-to-end calculation at a high level. Ofcom decided that it would be useful to further develop some aspects of the model prior to consultation and asked Analysys Mason to undertake this development work. In order to explain the way in which the model has been developed, as well as its current design, the remainder of this section is structured as follows:

- we begin, in Section 3.1, by describing the key characteristics of the model originally developed by Ofcom
- in Section 3.2 we provide an overview of the modifications made to the model by Analysys Mason.

3.1 Key characteristics of the model initially developed by Ofcom

The model initially developed by Ofcom, which is built in Python 3, produced charts describing the percentage of repair and provision jobs completed within specified target times for a range of

8 To meet the SLA, CL2 repair jobs must be completed by the end of the next working day, while CL1 jobs must be completed by the end of the next working day plus one. Both MPF and WLR products can theoretically be designated CL1 or CL2, depending on the arrangement between the retailer and Openreach.

9 The model considers only appointed provision jobs and not those which do not require a customer appointment.

10 A small percentage of jobs fails because the customer or the RSP does not show up for the appointment. Such jobs are modelled, but are excluded from the measures above.

resource-utilisation levels. Below we discuss some of the key characteristics of this, subsequently updated, version of the model.

Demand

The original Ofcom model generates two queues of randomly arriving jobs, with the daily arrival rate assumed to follow a normal distribution.[11] The jobs in each queue could in theory be either provision or repair jobs, or a mixture of the two, but Ofcom had only tested scenarios in which all of the jobs in both queues are repair jobs. In practice, this meant that one queue represented CL1 jobs, while the other represented CL2 jobs.

Each queue could be further sub-divided into two configurable sub-queues in order to assess the impact of Openreach's notion of a 'glass ceiling'. This notion puts forward limits on the number of (non-MBORC)[12] jobs that can be completed within the SLA targets owing to factors not related to human resources, such as the availability of specialised tools or the time taken to get permits to dig up a road.

Resources and scheduling

The modelled queues in Ofcom's original model can be served with either a constant number of resources or one which varies day to day. Furthermore, the resources serving the queues can either be pooled, or split into two separate groups, with each serving one of the queues. In terms of how the resources are allocated to the jobs (which it calls 'scheduling'), the model makes use of the following approaches to allocate resource to queues on a daily basis (rather than, for example, choosing which resources to use as each job is completed):

- **Unprioritised** – the available resources are determined by the input resource utilisation for each queue (with or without a random uniformly distributed element).
- **Priority resource scheduling** – a form of sharing resources across the queues in which a CL2 queue gets all the resource its needs to complete jobs immediately, up to a maximum configurable level of the total resource (e.g. 90%). The CL1 jobs get the remainder of the resources. In this case, the input resource utilisation sets the overall resource volume shared between the two queues.

Service times and quantum of time modelled

As with Openreach's previous QoS model[13], job service times are assumed to be constant in Ofcom's model, with the working day in effect divided into a set of time slots (which are each 2.5 hours in length) in which one job can be completed. This means that, for example, travel time and the effects of variations in job length are not explicitly modelled.

---

11 In reality, the daily number of arriving jobs may not be normally distributed. We understand from Ofcom that analysis suggested that the normal distribution was a reasonably good fit to the actual data.

12 Some jobs are designated by Openreach as 'matters beyond our reasonable control': loosely, an equivalent to Force Majeure.

13 This is the model that Ofcom used the outputs of as part of its decision in 2014 on charge controls for the current charge control period.

Geography

In terms of geographical distribution of jobs, the model generates results for only a single region of undefined size, which could represent:

- the whole of the UK
- a General Manager (GM) or a Senior Operations Manager (SOM) region
- any single, contiguous region.

Modelling at the right scale is important because statistical aggregation effects ('laws of large numbers') lead to large collections of jobs being more efficiently served; in reality, smaller areas will be less efficient but far more realistic – for example, resources in Aberdeen are not available to fulfil job demand in London, and *vice versa*.

How the job completion time is estimated

Ofcom's initial model estimated job completion times by using that day's resource level and the queue length when a job enters the queue.

While this approach should work well for considering a single queue when resource levels are constant, it could introduce errors where the resource level fluctuates (as tomorrow's resource allocation would affect the outcomes). In particular, this may be an issue where resource levels are different when a job enters the queue, from when it is served (i.e. just after it reaches the head of the queue). Furthermore, this approach may not correctly estimate job completion times where the profile of jobs of different types entering the queue over time is variable. For example, if after a CL1 job has entered the queue an unusually large number of CL2 jobs arrive, which have earlier job completion deadlines, there may be a delay to completing the CL1 job if there is prioritisation in operation. These potential inaccuracies might have become significant once the model was adapted to address all job types (i.e. provision jobs and both CL1 and CL2 repair jobs) using a shared pool of resources; as we will describe later, we have adopted a significantly different approach.

Resource

Each resource is assumed to be able to complete four 2.5 hour jobs, every working day. No allowance is made for sickness, holiday, training etc. This means that these resources are not directly comparable to an FTE technician, although it is assumed that if 10% more resources are needed, then the same scaling factor would apply to FTE technicians.

Resource levels

Ofcom's initial model set the number of resources available to carry out jobs on each day (the resource level) indirectly via a resource utilisation (RU) factor. In particular, the resource level was set equal to the mean volume of job arrivals (over some period) divided by the RU.[14]

'Tools' jobs

In Ofcom's initial model, the effects of certain jobs taking a long time because they need specific skills or tools were modelled by having a percentage of each care level modelled with a separate queue, and explicitly starving these queues of resources, severely limiting the performance of these jobs (and increasing the time taken to serve these jobs).

The net effect of this is that there is a 'glass ceiling' – a performance level which is less than 100% which cannot be reached even if the resource level is very high, owing to factors not related to resource levels (such as the availability of specialised tools or the time taken to get permits to dig up a road).

3.2 Overview of the modifications made by Analysys Mason

Ofcom recognised a number of limitations with its initial model and that further development of this model, in specific areas, was therefore needed. Additional capabilities and functionality were required to make the model more reflective of the real-world situation it is seeking to capture. In particular, Ofcom set out within our scope of work three main areas where it was desirable to extend the model:

- adapt the model such that both repair and provisioning jobs can be considered together, using a shared pool of resources
- allow for a collection of regions/areas to be considered simultaneously, and for resources to be shared between these different regions (where appropriate)
- adapt the model such that it is capable of using actual historical fault and provision data for job arrivals rather than random-number-generated arrivals from a specified distribution.

In order to implement these changes, some fundamental modifications were required to the way in which the model works. We categorise these model developments into seven main categories, listed below and described individually in more detail in the sub-sections that follow:

- modelling at the SOM region level, with outputs for each GM region and nationally
- resource sharing between regions (to help neighbouring regions or regions in crisis)
- job tracking from logging to completion
- measurement of volume of resources
- handling of failed jobs and customer 'no-shows'

14 This approach is not intrinsically problematic, but we note it since Analysys Mason's update to the model has fundamentally changed this approach.

- allocation of the daily level of resource between provision and repair jobs (and between CL1 and CL2)
- data used and required adaptations to that data to support the scenarios of interest.

There are also a number of modelling assumptions that we have retained from Ofcom's initial model. These are:

- the use of time periods, whereby each working day is divided into a number of separate periods
  - however, the model now considers five periods per day[15] (of which four are working periods, while the fifth is a non-working period which simply serves to log jobs that arrive when staff are not working). Note that under normal circumstances, no jobs take place on a Sunday and only a smaller volume of repair CL2 jobs occur on a Saturday[16]
- daily setting of actual resource levels available to serve each of the modelled queues
- a constant number of served jobs per day per resource
- a constant job duration.

Modelling at the SOM region level, with outputs for each GM region and nationally

We model each of 56 SOM regions separately, and exclude the two Northern Ireland SOM regions as the input data for these areas was incomplete.[17]

Where the input data has not been provided at SOM or more granular level, we have multiplied the national-level data by a suitable proxy ratio to generate a (hypothetical) SOM-level input. This chiefly affects provisioning demand data. The real data may not be distributed in this way (it might be more unevenly distributed over time, for example) and the realism of the modelling could be improved if the provisioning demand data were available in a more geographically granular form (e.g. tagged by main distribution frame (MDF)).

Each of the 56 SOM regions modelled lies within one of 9 GM regions modelled (again, Northern Ireland is excluded).

We also store information about the adjacency of the SOM regions such that neighbouring regions can in some circumstances share resources (see below).

15 As opposed to four periods per day that were previously modelled.

16 When varying the care level mix to estimate the impact of future changes on care levels, we allow both repair care levels to share resources on a Saturday. If we did not, then the total resource levels needed are not a fair comparison for 100% CL1, because this is not getting the benefit of the amount of resource which is normally deployed on a Saturday (normally serving CL2 jobs, with an indirect impact on CL1 performance at later dates).

17 We understand from Ofcom that Northern Ireland tends to achieve higher performance levels than most other GM areas. In particular, provision and repair CL1 performance is above average for Northern Ireland's 2 SOM areas, with repair CL2 performance above average for WLR and FTTC jobs but below average for MPF.

The output data is gathered, and histograms are generated, both nationally, and at the GM region level. We also explicitly output data on the worst GM region, because Ofcom seeks that Openreach deliver the required performance in each GM area.

Resource sharing between regions (to help neighbouring regions or regions in crisis)

In order to capture some efficiencies that Openreach is able to realise due to variations in job arrival rates across different regions at different times, regions can lend resources to a neighbouring region in some circumstances. However, such 'lending' does lead to inefficiencies, which arise due to travel time (e.g. at the beginning and end of the day), and also the lack of familiarity of staff with the network locations if resources are 'borrowed' from another region.

This resource sharing is done on a day-by-day basis and (in the model) is used only for repair jobs. This will therefore have only an indirect effect on provisioning.

We have implemented an approach that allows (optionally) to:

- *Help a neighbour* – Resources from one region can work in a neighbouring region where both the 'donor' region has sufficient resources to meet its own levels of demand and the 'recipient' region does not.
- *Help in the event of a crisis* – resources from any region (irrespective of whether they are neighbours) can be sent to assist the worst-performing region in the country (i.e. where the queue is longest relative to the permanently available resources), which might occur in reality, for example after cases of severe weather.

Neighbours are defined using a matrix of inter-region "centroid"[18] distances and a measure of how urban the region is. Any pair of regions for which the inter-region centroid distance is less than a given threshold are classed as neighbours. Currently, the thresholds are set at 50, 80, and 120km for decreasing degrees of 'urbanness'.

After having worked out a given region's resource balance (see below), the model looks to help the single worst-off neighbour for each potential donor region in turn[19]. Help is only given if the donor has short-enough queues and sufficient resources that in aggregate its provision, CL1 and CL2 jobs are expected to meet their targets. There is no point in transferring resources if the donor is not going to meet its own targets, given the inefficiency of transferred resources (discussed below).

The resource donated is based on the minimum of what is needed to restore performance in the recipient region, and the amount that the donor can afford to give before endangering its own

---

18 These points and the distances between them were generated using MapInfo, and are not strictly speaking centroids.

19 We do not recalculate the resource balance after a region accepts help, although this aspect of the modelling could be changed at a later date. The SOM regions are considered in alphabetical order; though we suspect this ordering does not change the result (as we do not recalculate the resource balance after receipt of help). Each region selects its worst -off neighbour, and that region can in theory gain resources from several neighbours.

performance. The help given is less effective than a resource working in its home region: the effective resource transferred is multiplied by an adjustable factor, set to 75%,[20] to reflect the cost of travel time or the relative unfamiliarity of the donated resource with the region.

Sharing of resources with regions in crisis works similarly as for neighbouring regions, but does not consider neighbours; instead, the model looks to help the single-worst region on that day, defined as the region with the longest queues relative to its resources and targets. Each potential donor provides a certain amount of resources based on that region's ability to contribute, and based on the total needed to restore that region to having enough resource. The algorithm is controlled by four parameters regarding the resources relative to targets and current queue length: the level below which the worst-off region has to be to get any help; the maximum level that help will be given to restore; the level below which a donor will not be reduced, and the level above which a potential donor has to be to contribute.

Job tracking from logging to completion

We have implemented a different approach to that used by Ofcom in its initial model: rather than estimating job-completion times, we measure them, which involves tracking each job from entering the queue to completion. In our model, each service class (repair CL1, repair CL2 and provision) has its own queue for each geographical region considered.[21]

Job tracking is done by storing, for each job within the queue, the time at which that batch entered the queue and the target time for that job to be completed (according to the relevant care level and the date and time at which it entered the queue). The period in which the job is served (i.e. leaves the queue and is completed) is also logged when the job is completed. There is a complication related to 'no-shows' and 'fails', which is discussed below.

A histogram is subsequently produced, based on the completion time relative to the corresponding target for each job type (provisions, repair CL2 and repair CL1). The model maintains separate data stores for each resource level tested, and reports performance statistics by year, service class and resource level at the end of the model run.

Measurement of volume of resources

A single national/overall resource level is specified for each run of the simulation. During model initialisation, this resource level is divided across the SOM regions using a proxy for expected demand (related to the number of faults in a region in 2011/12 and 2011/12 working system size).

20 That is, donated resources are only 75% as effective as when working in their home region. For example, if 20 resources were donated, this would be modelled as 15 additional resources in the recipient region.

21 In fact, largely for performance reasons, we track batches of jobs of the same job type which enter the queue in the modelled SOM regions in the same 'period', which is an equivalent approach.

Should there be regions which are short of resources, the 'neighbour sharing' and 'crisis sharing' mechanisms discussed above can come in to play, but these would be less efficient than giving a SOM region sufficient resource to start with.

A possible improvement in this area may be that in reality Openreach has a different way of allocating resources between SOM regions, which may also be able to vary slowly over time (e.g. by varying the recruitment levels in different areas), should the need for resources change over time. If real resource allocation data were available, this could be used to set the ratios of the resources allocated to each SOM region.

Handling of failed jobs and customer 'no-shows'

To better capture the effect of no-shows or failed jobs on the levels of resource required to hit the set performance targets, the percentage of no-shows and failed jobs for each service class is fed into the model. Failed jobs are classified as either 'major fails', or 'minor fails':

- *Repair minor fails* are modelled as being able to re-enter the queue at the front of the queue the next day when, for example, the correct tools can be brought to the site.
- *Provision minor fails* are modelled as being able to re-enter the provisioning booking system the next day, in effect at the back of the queue. In practice, this means that they usually have a significant wait to be re-appointed.[22] This is one area in which the modelling could in future be improved, but this would require significant changes to the model (e.g. an additional queue for urgent provision minor fails).
- *Major fails* are modelled as not being able to re-enter the queue for a period of five days; for example, due to the need for civil works. Unlike repair minor fails, repair major fails do not "jump straight to the front of the queue" when they re-enter the queue.[23]

Where a no-show occurs, the job re-joins the end of the queue the next day and resource will need to be allocated to re-attempt it. However, for both provision and repair, provided the job was appointed within the target period, it is not deemed to have missed the target (since the failure was not Openreach's fault). Such jobs are therefore excluded from the histograms measuring the proportion of jobs meeting the required targets.

For provision jobs, any fail (major or minor) is considered to cause the job to miss the CDD target (as the CDD is by definition the day on which the job was being attempted). This means that the CDD metric is given by the model inputs (i.e. 1 – (provision % minor fail + provision % major fail)).

For repair jobs, any major fail will result in the target for job completion within the SLA being missed. For repair minor fails, the impact is more nuanced: the job re-joins at the head of the relevant queue and can be re-attempted the next day. In theory, those jobs attempted on a day before the

22 This is because the model never cancels other provision appointments that have already been made

23 The model could be modified such that repair major fails jump straight to the head of the queue once they re-enter it, but this would not directly affect the results as such jobs will already have failed to meet the SLA.

deadline can still be completed within the targeted timeframe. However, for those first attempted on the day of the deadline, a minor fail will always cause the job to miss the target.

The option to trigger no-shows and major/minor fails can be turned on or off in the model. If turned on, the percentage of jobs that are no-shows, minor fails and major fails for each service class can be specified. Because we model "batches" rather than individual jobs, we use a binomial distribution to estimate the number of jobs in each "batch" that fail or are a no-show; this is equivalent to each job having a specific probability of being a fail or a no-show.

We use a similar probabilistic method in a number of places within the model in order to avoid systematic errors due to rounding to the nearest integer when dealing with relatively rare events (e.g. if the batch is of 40 jobs in a given SOM region and period, then a 1% chance of a no-show would lead (once rounded to the nearest integer) to 0 no-show jobs, unless a probabilistic mechanism is used). As a result, the model outputs do vary slightly if different random seeds are used.

Allocation of the daily level of resource between provision and repair jobs

The updated model has been designed to include both repair and provision jobs and to model the interactions between them, notably in relation to the shared pool of resources.

Four types of job are modelled: 3-day appointed provision (AP3), 10-day appointed provision (AP10), repair CL1 and repair CL2. The significance of the separation of the two types of provision is due to a requirement to wait at least 10 days before scheduling an appointment for certain jobs (as part of anti-slamming regulation) – these are the AP10 jobs, whereas an AP3 job can be scheduled after only 3 days.[24] As a result, AP10 jobs cannot meet the 10-day FAD target if the entire 10 days cannot be used. However, we believe that in practice Openreach will appoint on the tenth day and include the AP10 jobs on that basis (i.e. assuming a 9 day minimum delay) in the results presented later in this report.

There are three service classes (Provision, CL2 repair and CL1 repair). AP3 and AP10 provision jobs, which only differ by their minimum wait, are of the same service class (provision) and use the same queue. Each service class uses its own first-in, first-out (FIFO) queue for each SOM region. Provision jobs and repair jobs use different types of queues: provision jobs have a separate queue for each future period, which can be considered as the 'appointments diary' for that period, whereas the two repair queues for each region are not specific to a single period and jobs will be served when they reach the front of the queue.

The data provided to us by Ofcom suggests that around 99% of provision jobs are in the AP3 category. However, this is likely to underestimate the number of AP10 jobs, and we therefore assume, based on Ofcom analysis of Openreach data, that 10% of those categorised as AP3 jobs are in fact AP10.

24 In the model AP3 jobs are offered the first available slot for provision jobs more than 3 days after arrival in the queue, whilst for AP10 jobs the first available slot more than 10 days after arrival in the queue.

In the model there are 16 different job types, covering AP3, AP10, CL2 and CL1 with normal jobs, no-shows, major fails and minor fails modelled for each. All jobs start as “normal”.

The three service classes all share a common pool of resources in each SOM region and therefore an algorithm is needed to determine the daily number of resources allocated to each. Openreach’s resource allocation approach is complex, and it is clearly not possible to replicate the actual resource allocation and scheduling algorithm within this model. Instead, we follow a simpler approach. In essence, under this approach resources are allocated to the three service class queues such that the expectation is that jobs will meet their target.

At the start of the day, the algorithm decides how many resources to allocate to each service class for that day. Either a fixed ratio or the resource-balancing-allocation mode may be specified by the user:

- In the fixed ratio mode, the level of resource is set to a ratio specified by the user, based on the total resource level modelled.
- In the balancing mode, which we focus on in the remainder of this report, the model seeks to balance the performance across all job types relative to each job type’s performance target in days. The ideal minimum level of resource for each job type is the length of the queue at midnight divided by the product of working periods per day and the target days for that job type. Resources are pro-rated by the ratio of these target resource levels. For example, if repair CL2 has a queue of 100 with a target of 1 day and repair CL1 has a queue of 160 with a target of 2 days, then CL2 needs 25 resources[25] and CL2 needs 20, so repair resources will be allocated in the ration 25/45:20/45.

Provisioning resources are set well in advance of the day on which the appointments take place, because the jobs are ‘booked’ well in advance of this day. On each day, this method is applied once to work out the target provisioning resource level for future periods. In each region we keep count of the total number of appointed provision jobs as yet uncompleted (i.e. the total provision queue length); this (and the FAD target and the number of jobs per day) is used (together with the repair queue lengths, targets, and jobs per day) to set the desired target level of provision resource for (in principle) all future dates[26].This level can be reset each day, though if the level declines, previously made appointments are still respected.

The resource balance is then applied again to determine the repair balance given the resource previously committed to provisioning for that day (i.e. the maximum of the number of appointments booked in the past for each period on that day).

For provisioning, we maintain a separate queue for each future period in each region. Arriving jobs are appointed by placing them into the first free provisioning appointment after the required minimum (3 working days or 9 working days, respectively): we do not assume that the customer

25 This is because there are 4 working periods in a day.

26 The model only records this target level for periods that so far have appointments, and applies the current value to any period which gains its first appointment. We do not use averages of demand volumes, but the queue lengths do (in effect) mean that short duration fluctuations in demand are unimportant in setting the level of resources.

will decline the first offered date. This is an area in which the model could be modified in the future, although we expect that such a change will not greatly alter the results, for the following reasons.

- First, once the appointment has been offered within a suitable time frame, the FAD target has been successfully met even if the appointment is declined, and this is correctly captured by the model. Thus, more customers can in principle meet the target with the same level of resource if there are declined appointments.
- Second, although in reality appointments may be turned down, the total number of jobs that must be carried out is captured correctly and where an appointment is turned down the relevant slot would be left open for the next job reaching the front of the queue. It is possible, though unlikely when the system is under load, that there might be un-booked appointment slots, which could lead to idle resources. On balance we believe that the model is still a good representation of the real-world resourcing requirements.

Once appointed, the model will meet that appointment (i.e. it will keep that resource for provisioning in that period). If a future period is 'full' (i.e. the number of jobs appropriate to the level of resource allocated in that period have already been appointed), the model will begin booking into subsequent periods. The algorithm alters the number of resources allocated to future provisioning jobs, but does not ever cancel appointments already made (so the level of resource cannot be reduced below the number of jobs already 'booked' in that period). For example, if the algorithm sets the target provisioning resource to 9 and there are already 10 appointments in that period, it will not reduce the provisioning resource demand in that period (i.e. 10 jobs will still be carried out in that period); but if it sets it to 11, then an additional appointment could be made to that period.

Repair jobs arriving during the day can be served in periods after their arrival if there are resources to do so, and as the queue is FIFO they will always be served after jobs of the same care level due on the day.

This method does not prioritise repair CL2 jobs at the expense of CL1 jobs, or *vice versa*. It uses the actual queue lengths at the start of the day, so if there is enough resource to not starve CL1 jobs on average, then CL2 jobs will also meet their targets. However, model outcomes (at least for resource levels around 80% targets) do seem to result in higher performance levels against targets for provisioning and CL2 jobs compared to CL1 jobs; this is partly due to the way resources are allocated, which can fail to put sufficient urgency on CL1 jobs due on that day in some circumstances.[27]

As we are seeking to hit specified targets in all three service classes at the same time, it would be inefficient to over-perform in one service class if the other service classes were not hitting the target. Accordingly, we have also built in the ability to bias the resource allocation algorithm. The bias factors are multipliers which are applied to the target days (these bias factors are recorded in the model outputs). A larger target decreases the resource allocation, so a multiplier of 1.1 would reduce

27 It is also partly due to the different types of metric. The FAD metric is met if the job is attempted, even if it fails. The repair metrics are only met if the job successfully completes, meaning that all major fails and some minor fails will not meet the job completion within SLA target. Because resources are allocated based on queue length with the aim of allocating adequate appointment slots for each job type, and not based on the chances of the job failing, the resulting performance of repair jobs will tend to lag behind the provisioning FAD metric.

the short-term resource allocation by 10%, all other things being equal. In the longer term, the algorithm will balance out, as if insufficient resource is being allocated, the queues will grow and this will cause additional resources to be allocated to that service class.

Data used and required adaptations to that data to support the scenarios of interest

Real fault data is read-in by SOM region, excluding the two Northern Ireland SOM regions, for 5 years, from 2011/12 to 2015/16 inclusive. We understand that the data includes all provisions and faults, regardless of whether they are covered by an MBORC designation.

A number of assumptions have been made, along with some modifications to the data to address issues that arise:

- Real provisioning demand data is also read-in, though only for appointed provision jobs. However, because the data is not subdivided in terms of SOM regions, the demand is pro-rated across SOMs by working system size in 2011/12, dropping any jobs assigned to the Northern Ireland SOM regions. The availability of granular provisioning data at the SOM level would improve the model.
- The AP3 and AP10 job mix has been varied, as described above, since the data provided does not appear to accurately identify all AP10 jobs. These changes are parametrised and therefore the AP3/AP10 job split can be adjusted within the model.
- Resource allocation by region is pro-rated by a factor that reflects both repair (based on 2011 calendar-year fault data) and provisioning demand (based on working system size on 1 April 2011) weighted together. We have not tested the impact of alternative allocations of resources between regions, for example based on fault data and working system size in other years.
- Provision jobs that fail on the first appointment miss the CDD target and as such the proportion of major and minor fails limits the performance level against this target. Ofcom has requested a scenario with 95% performance against this target, which is only achievable if fails do not exceed 5%. Data provided to us by Ofcom suggests that 1.4% of provision jobs are major fails, but 6% are minor fails. We have reduced the proportion of minor fails to 3.6% to ensure that Ofcom's preferred CDD scenario is possible to achieve.
- We have also made a similar assumption in relation to the number of "minor fails". At one stage of model construction, all minor fails would fail to meet the SLA. This meant that it was necessary to reduce the proportion of repair minor fails from 5% to 3%such that the target of 93% repair performance was possible to achieve. Note that by reducing the "minor fail" percentage to 3% for repair jobs we have reduced the overall number of repair fails to 6%. As the target approaches the "glass ceiling", there is less and less ability for any other job to miss the target in order for the required performance level to be achieved, which leads to very large increases in the required number of resources. The handling of the time information related to 'fails' has been improved since this assumption was made, but we have retained the assumption of 3% minor fails.

In addition, some properties of the data have been parameterised to allow for scenario and sensitivity testing, results of which are provided in Section 5 and Annex B of this report:

- The job volume mixes can be varied for each repair care level to assess the impact of different care level mixes (i.e. how many additional resources are required for greater proportions of repair jobs being CL2 *versus* CL1). In these cases, we allow Saturday working for both CL2 and CL1 repair; the reason for this is that if we did not, the levels of "resource" are not directly comparable for the 100% CL1 case. (See Section 5.3)

- The proportion of FTTC provisioning jobs can be adjusted. We sensitivity test reductions of 10% and 20%, but leave the data unadjusted in this regard for our main scenario results. (See Annex B)

- The volume of faults can be adjusted. We have tested the impact of a 10% reduction in fault volumes. (See Annex B)

- The adjustment to the proportion of minor fails described above (from 5% to 3%) can be undone to reflect a case in which Openreach is not able to reduce the proportion of repair jobs failing at the first attempt. We have tested the impact on required resources and resource "deltas" with the proportion of minor fails restored to 5% for repair jobs. (See Annex B)

4 Summary of key inputs to, outputs from, and limitations of the model

Much of Section 3 describes modifications made to the Ofcom model affecting the mechanics of how jobs are scheduled, resources are allocated and performance against targets is measured. In this section we begin by providing a brief summary of the wider functionality of the model, including its key inputs and outputs (Section 4.1). We then comment on the extent to which the model constitutes a reasonable approach to estimating the outputs of interest to Ofcom, in particular highlighting a number of limitations that remain within the current implementation of the model (Section 4.2).

4.1 Key inputs and outputs

In simple terms, the model reads-in data relating to job volumes and the time at which each job was logged, along with resource numbers and target definitions (in relation to how quickly jobs need to be appointed and completed). It then performs a range of calculations to schedule jobs and record the time between each job being logged and it being completed for repair jobs, and whether an appointment was available within a target number of days for provision jobs. The outputs are performance levels against targets (i.e. a percentage of jobs for which the relevant target is met).

As such, the model does not calculate a number of resources required to achieve a given performance level, but rather a performance level that can be achieved with a set number of resources. The number of resources required for a particular level of performance therefore needs to be calculated by interpolating between pairs of results generated using different numbers of resources.

Although the model runs the field force simulation at a SOM level, it aggregates the results up to a GM level to measure the performance level against targets in each GM region. Regulation requires that a specified performance level is achieved within each GM region as well as nationally. As such, the performance level in the model is actually calculated as being the worst performance level achieved in any of the 9 GM regions modelled against each target.[28]

We have used the following data as inputs to the model:

- Openreach fault data from 2011/12 to 2015/16, given in terms of job arrival times for each care level by SOM region
- Openreach provisioning data from 2011/12 to 2015/16, given in terms of job arrival times for communications provider switches and FTTC provisions
- a list of SOM and GM region IDs and names, and a matrix of distances between the approximate centres of the SOM regions
- target times for completion of provision and repair jobs, and target FAD for provision jobs
- the proportion of jobs which are no-shows or major/minor fails

28 The performance level against each target is equal to the worst-performing GM region in relation to that target. It is possible that the worst-performing GM region may be different for the different targets.

- a list of resource levels to be considered in the simulation.

The model then formats the data for use in the field-force simulation, which runs through all years considered in chronological order. At the end of the simulation, the model reports the performance level against the target for provisions and repair CL1 and CL2 jobs, for each year and for each tested resource level. A summary of the model process flow is shown in Figure 4.1 below.

Figure 4.1: Model flowchart [Source: Analysys Mason, 2017]

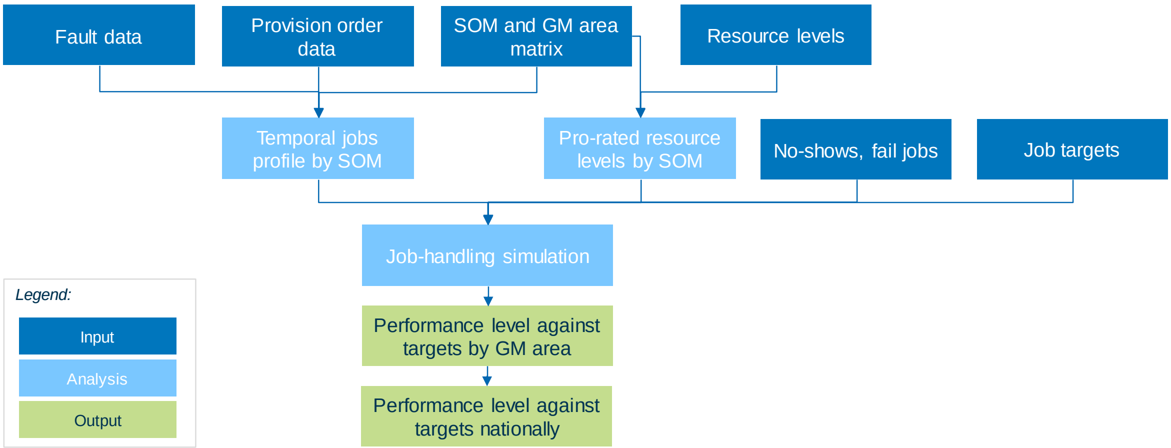

4.2 Limitations of the model

The functionality of the model has been kept relatively simple in order to produce results within the available time that are robust enough for Ofcom to use for the purposes of its consultation. We believe that the model provides a reasonable reflection of the resourcing challenges facing Openreach. However, because of the relatively simple nature of the model, it is subject to certain limitations. We identify and briefly discuss some of the more significant limitations below.

Job durations

The duration of every job is assumed to be a constant 2.5 hours. While this is likely to be close to the mean job duration in each modelled year, in reality, there will be some variance. Not taking the variations in job duration into account could, on some days, have a material impact on the number of jobs that can be completed for a given resource level within a region.

Incorporating a distribution of job durations would be time-consuming to implement and greatly add to the complexity of the model (because each job needs to be handled individually, and because if the job duration varies the service of jobs does not happen period by period, but in a more asynchronous fashion). We note that the previous model produced by Openreach made a similar assumption of constant job duration.

Volume of resources

A constant volume of resources is assumed for each modelled scenario, with no day by day resource fluctuations (e.g. due to sickness, holidays, part-time working) other than reduced staffing on Saturdays. Whilst in reality the resourcing levels may fluctuate, it is unlikely that this would have a significant impact on the resource deltas between different scenarios.

Resource allocation

► *Pro-rating resources across regions*

We assume a fixed split of resources between SOM regions based on a weighted sum of that region's fault rate and working system size as of 1 April 2011. The relative working system size between SOM regions may have changed since 2011, or the ratio of provisions to working system size may have changed, or the fault rate per line may be higher in certain regions, or may be changing over time. The modelled approach could therefore allocate resources between SOM regions in a sub-optimal way: it is possible that Openreach could attain a higher proportion of jobs meeting the SLA targets by allocating resources more appropriately, but equally we do not wish to bias the results by using the outcomes from the model after the fact to set the resource levels (i.e. in effect using foresight of jobs that will actually arise in a given year). As such, we have not tested the impact of alternative allocations of resources between regions, for example based on fault data and working system size in other years. If real resource per region data were available, then this could be used to set the appropriate ratios of resources between regions.

► *Allocating resources between repair and provision*

Allocation of resources between repair and provision may not always lead to optimal performance against the respective targets for the three job types. For example, resources are allocated between CL1 and CL2 repair queues based on the total length of the queue, using the approach described in Section 3.2. As such, no account is taken of the number of CL1 jobs in the queue whose deadline is the same day (i.e. because they were not addressed on the previous day despite having been in the queue). This could lead to some over-allocation of resources to CL2 and under-allocation of resources to CL1 in some circumstances. We could modify the model to track this variable and modify the resource allocation to allow for it (e.g. in setting the minimum resources).

While the bias mechanism incorporated within the resource allocation algorithm does serve to modify the allocation of resources between queues, it is an additional external variable and it would be computationally expensive to 'goal seek' it to the optimal set of values in each scenario.

► *Skill level of Openreach technicians*

The impact of technicians with different skill sets is not considered by the model. All technicians are assumed capable of carrying out all jobs. In reality some technicians will not be capable of carrying out all types of job and this will place some additional constraint on Openreach's resourcing. As a result the model may underestimate the number of resources required to achieve a given performance level against targets. However, the impact on the resource deltas required to achieve a higher performance level than in the base case is unclear.

Job-handling and glass-ceiling effects

The percentage of provision jobs completed by CDD is not explicitly modelled, as it is entirely determined by the model inputs about the number of 'fails'.

The model tracks when minor fails originally entered the queue. This means that at high levels of resource, some CL2 jobs first attempted on their day of arrival can be re-attempted and completed the next morning, thereby meeting the SLA target. In principle this means that the 93% performance level could be hit without modifying the percentage of minor fails (with sufficient resources allocated), although we have not generated results for this scenario (i.e. all the results presented in this report use the lower "minor fails" percentage).

MBORC

We note that the model makes no allowance for MBORC. If the input data could be identified as such, then these job outcomes could be excluded from the calculated results. The model considers all jobs that were logged in each year and measures performance against targets across all of those jobs. As such, the model is likely to be conservative, because some of the modelled jobs would actually be exempt from the measure of performance due to an MBORC designation. In effect, fewer resources than those calculated may actually be needed since performance level against a given target is likely to be higher were MBORC jobs excluded. The impact on resource deltas between different scenarios is however unclear.

In any event, we understand that Ofcom has historically based its decisions on a similar modelling approach, whereby all jobs are taken into account when determining performance levels and targets, with designations of MBORC occurring thereafter.

5 Model outputs

As explained above, the model aims to calculate percentage deltas in the number of FTE resources needed to complete different percentages of jobs within specified targets, or with varying percentages of repair jobs at each care level, relative to a *base case*. This base case needs to be defined such that the performance levels achieved (i.e. proportion of jobs completed within the target time) matches as closely as possible to Openreach's actual performance target in that year.

The input that is varied in order to determine an appropriate base case is the number of resources. Note that since actual performance levels and job volumes vary year by year, a base case needs to be defined separately for each year in which the model is to be run. Although we have been provided with data for a five-year period running from 2011/12 to 2015/16, we understand that the years 2015/16 and 2014/15 are of primary interest to Ofcom. We understand that 2011/12 is of secondary interest in order to compare results to those obtained using Openreach / E&Y's modelling at the time of setting the charge controls for the 2013/14 to 2015/16 period.

It turns out that many more resources are needed for 2014/15 and 2015/16 than for 2011/12. This appears to be because the volume of jobs is higher in 2015/6. We have run a separate set of simulations just for 2011/2 to test the lower levels of resource appropriate to 2011/2.

In this section we first set out the base cases we have defined for the years of interest (Section 5.1). We then introduce the scenarios to be considered (Section 5.2) and present the results obtained from the model for each of those scenarios (Section 5.3). A further breakdown of some results by GM area is provided in Annex A, while Annex B presents results for some specific sensitivity tests, as requested by Ofcom.

5.1 Base case

In order to determine the most appropriate base case for each year of interest, we have run the model with different resource levels to determine performance against the three measured minimum standards in each case. A 12-day FAD is used since this was the target in place during these years.

The following table shows these performance levels for a range of different resource levels, for the years 2014/15 and 2015/16. As the table shows, performance levels of 80% provision appointments within 12 days and 80% repair within SLA was achieved for between 7250 and 7500 resources in 2014/15 and between 8000 and 8250 resources in 2015/16, as we describe further below.

Figure 5.1: Performance levels in each year from different numbers of resources (bias factors: provision 1.2, CL2 0.8, CL1 0.8) [Source: Analysys Mason, 2017]

Year	Job type	Number of resources													
		6250	*6500*	*6750*	*7000*	*7250*	*7500*	*7750*	*8000*	*8250*	*8500*	*8750*	*9000*	*9250*	*9500*
2014/15	Provision (12-day FAD)	3.7%	7.8%	22.8%	54.7%	78.1%	96.1%	98.8%	99.8%	99.9%	100.0 %	100.0%	100.0%	100.0%	100.0%
	Repair CL2	1.6%	10.8%	44.2%	72.8%	87.7%	92.5%	94.4%	95.0%	95.1%	95.2%	95.3%	95.3%	95.4%	95.4%
	Repair CL1	0.0%	3.8%	34.0%	63.8%	84.1%	91.1%	95.0%	96.1%	96.4%	96.5%	96.5%	96.6%	96.6%	96.6%
2015/16	Provision (12-day FAD)	0.7%	1.2%	2.2%	10.4%	24.1%	39.9%	52.4%	68.2%	86.7%	96.8%	99.6%	99.7%	100.0%	100.0%
	Repair CL2	0.0%	0.1%	3.3%	17.2%	35.4%	48.3%	60.4%	78.5%	88.2%	92.2%	94.6%	94.7%	95.2%	95.3%
	Repair CL1	0.0%	0.0%	1.2%	11.5%	27.9%	41.1%	49.8%	71.1%	83.3%	91.3%	95.3%	96.0%	96.4%	96.6%

We note that the current performance targets applied to Openreach are as follows:[29]

- provision appointment availability (12-day FAD) – **80%**
- provision by CDD – **90%**
- repair within SLA – **80%.**

Given our model outputs, we consider that the following interpolated resource levels provide the best fit as a base case for each year, when compared to the current targets:

- 2014/15 – 7276
- 2015/16 – 8182

We have also calculated a base case resource level for 2011/12 in order to replicate actual performance in that year, namely:

- provision appointment availability (12-day FAD) – **65%**
- provision by CDD – **90%**
- repair within SLA – **77.8%.**

This leads to an interpolated base case resource level for 2011/12 of 4361.

5.2 Scenarios to be considered

Ofcom has asked us to initially consider 20 different scenarios. In each case, the output is the percentage resource delta – compared to resource volumes from the base case for the relevant year – that is required to achieve the performance levels specified in that scenario. Each scenario consists of a combination of input parameters taken from the following three dimensions:

- a range of minimum standard values (2 options)
- a range of different product and care-level volume mixes (5 options)
- provision, repair and resource-volume profiles that arose in different years (2 options).

The different options for the minimum standard values, care-level mix of repair jobs and job-volume temporal profiles are shown in Figure 5.2, Figure 5.3 and Figure 5.4, respectively.

29 We note that Openreach's actual performance over this period is likely to be somewhat higher, at least for appointment availability, and may also vary by year. However, for the purposes of defining a base case, we consider target performance to be a reasonable proxy – it is after all the delta from the base case performance that we are interested in, rather than the resourcing volume in the base case itself.

Figure 5.2: Minimum standard options [Source: Ofcom, 2017]

Option[30]	Appointment availability (10-day FAD)	Repair within SLA
1 ('Base Year')	80%	80%
2 ('Year 3')	90%	93%

Option	Volume mix
1	100% CL1, 0% CL2
2	60% CL1, 40% CL2
3	50% CL1, 50% CL2
4	40% CL1, 60% CL2
5	0% CL1, 100% CL2

Figure 5.3: Care-level volume mix [Source: Ofcom, 2017]

Option	Resource–volume profile
1	2015/16 (largely a benign year but with bad weather in the second half of the year)
2	2014/15 (a benign year)

Figure 5.4: Job-volume temporal profiles [Source: Ofcom, 2017]

In addition we have briefly considered headline results for 2011/12, for comparison with previous modelling relied upon by Ofcom in its 2014 decisions related to the current charge control period. In particular we calculate the resources required to achieve the performance levels of minimum standard option 1 (80/80%) and compare this to the base case resources for 2011/12 described above.

5.3 Scenario results

2015/16

We have adjusted the resource levels to understand the resource delta from the base case that is required in order to meet each combination of minimum standard and care-level volume mix. The resulting resource levels for 2015/16 are shown in the following tables. These resource levels are calculated such that the worst performing GM area achieves the required performance level. Performance levels in other GM areas may therefore exceed the minimum standard.

30 The names for each minimum standard options relate to a glide path from 2015/16 to 2020/21 that we understand Ofcom is considering in relation to its upcoming consultation

Figure 5.5: Resource levels by minimum standard option and care-level volume mix for 2015/16 [Source: Analysys Mason, 2017][31]

M Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
Current performance 12 day FAD	8182					
1 (80/80) 10 day FAD	8232	8182	8213	8212	8216	8378
2 (90/93) 10 day FAD	8567	8365	8575	8656	8661	8839

Figure 5.6: Percentage resource deltas by minimum standard option and care-level volume mix for 2015/16, relative to 12-day FAD, 80% attainment [Source: Analysys Mason, 2017]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
1 (80/80)	0.6%	0.0%	0.4%	0.4%	0.4%	2.4%
2 (90/93)	4.7%	2.2%	4.8%	5.8%	5.9%	8.0%

31 The actual care level mix in 2015/16 was 44% CL1 and 56% CL2. Bias values of provision = 1.2, CL2 = 0.8, CL1 = 0.8 were used.

It can be observed that under minimum standard option 1, for care-level mixes 1 to 4, fewer resources are needed than for the actual care-level mix for minimum standard option 1. Conversely, more resources are required for care-level mix 5. Note however, that the resources required for the actual mix are not directly comparable to those required for the other care level volume mixes. This is because the calculation for all of these other care level mixes assumes that CL1 jobs can be addressed on a Saturday, as described in Section 3.2.

To observe the impact of the care level volume mix we therefore focus on the differences in resources required only within the care level volume mix scenarios other than the actual mix. However, the effect of changing the care-level mix is non-linear. There appear to be at least two reasons for this:

- For minimum standard option 2, where the repair within the SLA target is set to 93%, having no CL2 jobs requires fewer resources. We suspect this is due to it being harder to meet the CL2 target close to its "glass ceiling", because sufficient resource needs to be allocated such that some fraction of 'minor fail' cases can be reattempted on the next day within the SLA. One possible improvement would be to re-bias the resources (change the bias factors) such that CL2 gets more resources when aiming for the 93% target.
- Between care-level volume mixes 2, 3 and 4 the resource deltas do not increase uniformly for either minimum standard option. This may be related to the use of random numbers at certain points within the model, which can be illustrated by the variation between the outputs of the model for the same input parameters if different random seeds are used. In essence, we would expect the resources required for these three difference care-level volume mixes to be close, and therefore small variations arising from such differences could be sufficient to affect the ordering of resource levels required under these scenarios. Further investigation may be useful

As the minimum standard is increased (option 2), the number of resources required also rises.

2014/15

The same analysis has been performed for 2014/15, with the results shown in Figure 5.7 and Figure 5.8. The observed trends are broadly similar to those for 2015/16. There is a slight difference in the impact of varying the care-level volume mix. Overall, the number of resources required for any scenario remains lower than for the corresponding scenario in 2015/16: it is easier to achieve a given target in 2014/15 due to the number or profile of jobs.

Figure 5.7: Resource levels by minimum standard option and care-level volume mix for 2014/15 [Source: Analysys Mason, 2017] [32]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
Current performance	7276					
1 (80/80)	7336	7256	7291	7312	7339	7404
2 (90/93)	7614	7401	7714	7739	7752	7952

Figure 5.8: Percentage resource deltas by minimum standard option and care-level volume mix for 2014/15, relative to 12-day FAD, 80% attainment [Source: Analysys Mason, 2017]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
1 (80/80)	0.8%	-0.3%	0.2%	0.5%	0.9%	1.8%
2 (90/93)	4.6%	1.7%	6.0%	6.4%	6.5%	9.3%

32 The actual care level mix in 2014/15 was 49% CL1 and 51% CL2. Bias values of provision=1.2,CL2=0.8,CL1=0.8 were used.

We understand that Ofcom is interested in deriving a factor to estimate the percentage change in resource needed for a given percentage change in the care-level volume mix. Due to the non-linearity of the resource deltas required to achieve each minimum standard under different care-level volume mixes, there is no uniquely correct factor of this type. In particular, measuring the gradient of a straight line between different pairs of points will give different answers, depending on the pair of points chosen.

We recommend that Ofcom either produce a line of best fit using all points (i.e. the five care-level volume mixes[33]) under each minimum standard option in each year, or do the same but excluding the 100% CL1 case, which would avoid the difficulties associated with the results for care-level volume mix 1.

2011/12

The base case for 2011/12 (which has a lower target level of performance and a 12 day FAD) required 4361 resources. Increasing the performance levels to 80% repair within SLA and 80% of provision jobs able to offer a FAD within 12 days increases the resource requirement to 4558. This corresponds to a resource delta of 4.5 %.

By way of comparison we note that in 2014 the resource delta arising from the previous modelling (the Openreach E&Y model) was 3.9%.

33 Using the actual care level mix as an extra data point when trying to estimate the impact of changing the care level mix is probably not appropriate because the model makes different assumptions about Saturday working for this data point (no CL1 repair on Saturday).

6 Conclusions

The model has been developed in order to run the scenarios requested by Ofcom, most notably in relation to varying both the minimum standards (i.e. performance level against targets) and the care-level volume mix for repair jobs.

We have tested the impact of these variations on resource requirements for Openreach in both 2014/15 and 2015/16 against two base cases: a minimum standard of 80% provision FAD (for a 12-day FAD) and 80% repair within the SLA target. For 2014/15, the base case requires 7276 resources, while in 2015/16, a less benign year, the corresponding number is 8182.

The results of our scenario runs against these base cases are shown for 2015/16 and 2014/15 in Figure 6.1 and Figure 6.2, respectively.

Figure 6.1: Percentage resource deltas by minimum standard option and care-level volume mix for 2015/16, relative to 12-day FAD, 80% attainment [Source: Analysys Mason, 2017]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
1 (80/80)	0.6%	0.0%	0.4%	0.4%	0.4%	2.4%
2 (90/93)	4.7%	2.2%	4.8%	5.8%	5.9%	8.0%

Figure 6.2: Percentage resource deltas by minimum standard option and care level volume mix for 2014/15, relative to 12-day FAD, 80% attainment [Source: Analysys Mason, 2017]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
1 (80/80)	0.8%	-0.3%	0.2%	0.5%	0.9%	1.8%
2 (90/93)	4.6%	1.7%	6.0%	6.4%	6.5%	9.3%

Although the relationship is non-linear, there is a general trend that the required number of resources increases both as the proportion of repair jobs at CL2 increases and as the minimum standard for provision and repair is raised. The increase in minimum standard has slightly more impact on the required number of resources (and hence resource deltas) than the increases required to address higher proportions of CL2 jobs.

The observed trends are broadly similar between 2015/16 and 2014/15, though varying the care-level volume mix requires a somewhat higher resource delta for minimum standard option 2 in 2014/15.

A line of best fit could be used to help determine a ‘gradient’ to establish changes in the level of resources needed for a (small) unit change in care-level volume mix; the 100% CL1 case may be an outlier and might be excluded from this calculation. We note that some further investigation may be required to understand the behaviour in relation to the care-level volume mix more fully.

We have also tested (see Annex A) the variation in performance across GM areas and have observed that for high minimum standards (e.g. 90/93%) there is little variation, but for lower minimum standards (e.g. 80/80%) the worst-performing GM area may exhibit significantly lower-than-average performance. We believe that this is due to the “S-curve” nature of performance against resources: the gradient of performance against resource allocated is high when the performance is 80% but low when the performance level is above 90%. The volume of additional resource that Openreach will require to meet a set minimum standard in every GM area rather than just on the basis of a national average, will therefore vary according to the level of the minimum standard.

If the minimum standard is high (e.g. 90/93%) then the resource delta to achieve that standard in every region rather than “just” nationally may be relatively low. Conversely if the minimum standard is low (e.g. 80/80%) then the resource delta to achieve that standard in every region rather than just nationally may be relatively high.

Annex A Analysis of performance-level variations by GM area

In this annex we illustrate the variation in performance levels across GM areas. Regulation requires the performance levels specified by Ofcom to be achieved in each GM area. The results presented in Section 5.3 show the resource delta required such that all GM areas meet or exceed the target performance. In essence, this means that the worst-performing GM area for each of the targets will achieve a performance of just above target, but other GM areas may significantly exceed the performance targets.

In Figure A.1 we set out the performance level against each target by GM area for a fixed number of resources. We have chosen the 8250-resource scenario as this scenario requires the minimum number of resources such that performance in the worst-performing GM area exceeds the 'base year' targets of 80% provision FAD and 80% repair within SLA in 2015/16 for the actual care-level volume mix.

Figure A.1: 2015/16 performance level by GM area with 8250 resources (hits target performance of 80% FAD, and 80% repair SLA in the worst-performing GM area, current CL mix), bias parameters: provisioning= 1.2, CL2=0.8, CL1=0.8 [Source: Analysys Mason, 2017]

Care-level volume mix option and resource numbers	Targets	2015/16 performance level by GM area								
		East Anglia	London	North East	N. Wales and N. Midlands	North West	Scotland	South East	S. Wales and S. Midlands	Wessex
Actual levels	Provision (FAD)	93.1%	91.9%	90.0%	85.9%	82.4%	89.4%	91.0%	86.8%	87.8%
	CL2 Repair SLA	92.0%	91.9%	89.9%	89.3%	86.5%	89.7%	91.1%	88.2%	90.4%
Number of resources = 8250	CL1 Repair SLA	90.9%	91.7%	86.2%	86.0%	81.0%	87.8%	89.7%	87.2%	87.3%
1 (100/0)	Provision (FAD)	97.1%	95.1%	95.2%	93.5%	85.2%	93.6%	97.4%	91.6%	93.5%
Number of	CL2 Repair SLA	nan	nan	nan	nan	nan	nan	nan	nan	nan
resources = 8250	CL1 Repair SLA	95.1%	95.7%	93.1%	92.5%	91.0%	93.4%	95.4%	91.3%	93.2%
2 (60/40)	Provision (FAD)	95.0%	93.2%	92.3%	90.4%	83.6%	90.2%	94.0%	88.3%	89.4%
	CL2 Repair SLA	90.9%	91.9%	90.0%	89.0%	85.7%	88.5%	90.9%	87.1%	87.0%
Number of resources = 8250	CL1 Repair SLA	93.3%	94.1%	91.5%	91.1%	87.9%	90.5%	93.5%	89.7%	88.6%
3 (50/50)	Provision (FAD)	93.8%	92.5%	93.1%	89.3%	83.4%	90.8%	93.4%	88.6%	89.9%
Number of	CL2 Repair SLA	91.0%	91.0%	89.7%	87.9%	85.5%	88.3%	90.3%	87.2%	86.6%
resources = 8250	CL1 Repair SLA	93.0%	92.8%	91.3%	90.2%	87.1%	89.8%	92.5%	89.5%	88.0%
4 (40/60)	Provision (FAD)	93.3%	91.8%	90.6%	87.6%	83.2%	89.1%	93.4%	86.9%	90.1%
	CL2 Repair SLA	89.8%	90.8%	88.6%	85.6%	84.3%	86.7%	90.2%	84.5%	87.2%
Number of resources = 8250	CL1 Repair SLA	91.8%	92.3%	89.7%	87.6%	86.1%	88.2%	92.5%	86.4%	88.8%
5 (0/100)	Provision (FAD)	90.3%	90.4%	87.2%	84.2%	72.4%	85.9%	88.7%	84.7%	86.8%
Number of	CL2 Repair SLA	83.9%	86.6%	82.6%	80.1%	73.1%	79.6%	83.3%	79.7%	80.3%
resources = 8250	CL1 Repair SLA	nan	nan	nan	nan	nan	nan	nan	nan	nan

The results show that there is reasonably significant variation between GM areas. The North West is the worst-performing area in this case, with London, East Anglia or the South East typically the best-performing ones. The North West appears to perform around 5 percentage points below the average GM area, and around 8 percentage points below the best-performing area. Note that for some specific scenarios and service classes, the differences can be greater. As noted above, this might indicate that when dividing the resources across SOM regions more resources ought to be allocated to the North West (and Openreach may already do this).

Figure A.2 provides a similar set of performance levels, but this time the targets are set to ‘Year 3’ levels (i.e. 90% provision FAD, and 93% repair SLA). The fixed number of resources used, such that the higher minimum standard is met in the worst-performing GM area for the actual care-level volume mix, is therefore higher than in Figure A.1 at 9250. With this number of resources, the performance against each target in each GM area is correspondingly greater.

Figure A.2: 2015/16 performance level by GM area with 8750 resources (hits target performance of 90% FAD, and 93% repair SLA in the worst-performing GM area, current CL mix), bias parameters: provisioning= 1.2, CL2=0.8, CL1=0.8 [Source: Analysys Mason, 2017]

Care-level volume mix option and resource numbers	Targets	2015/16 performance level by GM area								
		East Anglia	London	North East	N. Wales and N. Midlands	North West	Scotland	South East	S. Wales and S. Midlands	Wessex
Actual levels	Provision (FAD)	99.9%	99.8%	99.8%	99.5%	99.1%	99.4%	99.8%	99.3%	99.9%
	CL2 Repair SLA	94.8%	94.7%	94.9%	94.6%	94.4%	94.7%	94.8%	94.8%	94.8%
Number of resources = 8750	CL1 Repair SLA	95.9%	95.9%	95.8%	96.0%	95.3%	95.5%	96.0%	95.8%	95.5%
1 (100/0)	Provision (FAD)	99.9%	100.0%	100.0%	99.7%	99.9%	99.9%	100.0%	99.9%	100.0%
	CL2 Repair SLA	nan	nan	nan	nan	nan	nan	nan	nan	nan
Number of resources = 8750	CL1 Repair SLA	96.1%	96.4%	96.3%	96.0%	96.3%	96.4%	96.4%	96.3%	96.2%
2 (60/40)	Provision (FAD)	99.9%	100.0%	100.0%	99.4%	99.5%	99.6%	99.9%	99.4%	99.6%
	CL2 Repair SLA	94.3%	94.4%	94.6%	93.8%	94.2%	94.0%	94.1%	94.0%	93.7%
Number of resources = 8750	CL1 Repair SLA	96.0%	96.2%	96.0%	95.7%	96.0%	95.9%	96.1%	95.7%	95.4%
3 (50/50)	Provision (FAD)	99.9%	99.8%	99.9%	98.9%	99.0%	99.5%	99.9%	99.3%	99.6%
	CL2 Repair SLA	94.4%	94.4%	94.5%	93.4%	93.7%	94.1%	94.2%	94.0%	93.8%
Number of resources = 8750	CL1 Repair SLA	96.0%	96.1%	96.0%	95.1%	95.6%	95.9%	96.0%	95.7%	95.6%
4 (40/60)	Provision (FAD)	99.9%	99.8%	99.9%	99.6%	99.1%	99.5%	99.8%	98.9%	99.6%
	CL2 Repair SLA	94.4%	94.4%	94.5%	94.0%	93.9%	94.1%	94.1%	94.1%	93.9%
Number of resources = 8750	CL1 Repair SLA	96.0%	96.1%	95.9%	95.9%	95.7%	95.7%	96.0%	95.8%	95.5%
5 (0/100)	Provision (FAD)	99.8%	99.6%	99.7%	98.1%	98.1%	98.9%	99.7%	98.3%	99.5%
	CL2 Repair SLA	94.2%	94.0%	93.8%	92.3%	92.8%	93.4%	93.9%	92.4%	93.2%
Number of resources = 8750	CL1 Repair SLA	nan	nan	nan	nan	nan	nan	nan	nan	nan

This time we see that the variation between GM areas is significantly smaller, with almost no difference in performance. This is due to the target in 'Year 3' being significantly harder to reach and indeed close to the glass ceiling for repair jobs. As such, exceeding the minimum standard in any GM area is much more difficult (the performance/resource gradient is very low).

What these results imply is that were Openreach to be granted leeway in relation to one or more 'difficult' GM areas, this would likely make it easier (i.e. require fewer resources) to meet more modest minimum standards, such as 80% repair within the SLA target. However, for more challenging minimum standards, such as 93% repair within the SLA target, it would make little difference to Openreach. As such, and perhaps slightly counter-intuitively, granting such additional leeway could in fact increase the resource deltas required to increase minimum standards.

Annex B Further sensitivity testing of results

In this annex we provide the results of a number of sensitivity tests requested by Ofcom.

B.1 Impact of stripping out different percentages of FTTC provision jobs (0%, 10% and 20%) on interpolated resources to hit targets in 2015/16

We understand that Ofcom initially had some concern over whether simultaneous provide FTTC jobs have been double counted in the data. Although this concern was resolved, to test whether any such inaccuracies would have a material impact on our results, we have calculated the resources that would be required in the base case and the minimum standard options 1 and 2 with 10% and 20% fewer FTTC provision jobs. We show the results alongside those for the actual mix of jobs shown in the data in the table below. The calculated resource deltas to the relevant base case show that there is a relatively small difference between the three cases.

Figure B.1: Sensitivity testing of stripping out FTTC provision jobs [Source: Analysys Mason,2017]

Minimum standard option	Number of resources			Resource delta vs. base case		
	Actual mix	10% less FTTC provision	20% less FTTC provision	Actual mix	10% less FTTC provision	20% less FTTC provision
Base (80/80) 12 day FAD	8182	7933	7682			
1 (80/80) 10 day FAD	8232	7992	7718	0.6%	0.7%	0.5%
2 (90/93) 10 day FAD	8567	8333	8052	4.7%	5.0%	4.8%

B.2 Impact of a 10% reduction in fault volumes on interpolated resources to hit targets in 2015/16

We have also sensitivity tested the impact of a 10% reduction in fault volumes. The results are shown in

Figure B.2 below and indicate that with 10% fewer repair jobs to be carried out, the resource delta to achieve minimum performance standard 2 remains similar.

Figure B.2: Sensitivity testing of stripping out repair jobs [Source: Analysys Mason, 2017]

Minimum standard option	Number of resources		Resource delta vs. base case	
	Actual mix	10% less repair	Actual mix	10% less repair
Base (80/80) 12 day FAD	8182	7829		
1 (80/80) 10 day FAD	8232	7889	0.6%	0.8%

Minimum standard option	Number of resources		Resource delta vs. base case	
	Actual mix	10% less repair	Actual mix	10% less repair
2 (90/93) 10 day FAD	8567	8206	4.7%	4.8%

B.3 Impact of increasing the percentage of repair minor fails to 5%

We have also sensitivity tested the impact of setting the level of repair minor fails to 5%.

The results are shown below and indicate that with 5% repair minor fails, the resource delta to reach 93% is (as would be expected) rather higher. This is because a fraction of the minor fails need to pass the SLA by being initially attempted on the day before the due date.

Figure B.3: Percentage resource deltas for repair having 5% minor fail by minimum standard option and care-level volume mix for 2015/16, relative to 12-day FAD, 80% attainment [Source: Analysys Mason, 2017]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
1 (80/80)	1.1%	-0.1%	0.6%	0.6%	0.9%	2.6%
2 (90/93)	8.8%	2.4%	9.5%	10.2%	8.9%	10.5%

Figure B.4: Percentage resource deltas for repair having 5% minor fail by minimum standard option and care level volume mix for 2014/15 [Source: Analysys Mason, 2017]

Minimum standard option (% performance against 10-day FAD / repair within SLA)	Actual mix	Care-level volume mix 1 (%CL1 / %CL2): (100/0)	Care-level volume mix 2 (60/40)	Care-level volume mix 3 (50/50)	Care-level volume mix 4 (40/60)	Care-level volume mix 5 (0/100)
1 (80/80)	0.7%	0.1%	0.5%	0.8%	0.8%	2.6%
2 (90/93)	6.5%	1.8%	11.0%	11.3%	11.4%	12.3%

Annex C Glossary

CDD	Contractual delivery date
CL	Care level
FAD	First available date
FIFO	First-in first-out
FTE	Full time equivalent
FTTC	Fibre to the cabinet
GM	General Manager
LLU	Local Loop Unbundling
MBORC	Matters beyond our reasonable control
MDF	Main distribution frame
MPF	Metallic Path Facility
RU	Resource utilisation
SLA	Service level agreement
SOM	Senior Operations Manager
WLR	Wholesale Line Rental