www.analysysmason.com

Final Report for Ofcom

Quality of Service model assessment

22 November 2013

Ref: 37846-476

Contents

Analysys Mason Limited
St Giles Court
24 Castle Street
Cambridge CB3 0AJ
UK
Tel: +44 (0)845 600 5244
Fax: +44 (0)1223 460866
cambridge@analysysmason.com
www.analysysmason.com
Registered in England No. 5177472

1 Introduction

Ofcom is currently reviewing Wholesale Line Rental (WLR) and Metallic Path Facility (MPF) charge controls for the period April 2014 to March 2017. The charges are cost-based, and as a part of this work Ofcom is interested in understanding the resources involved in ensuring a given quality of service (QoS) from Openreach in terms of meeting its targets for provision of new connections and fixing reported faults.

Openreach has commissioned Ernst & Young to develop a Process Model (the 'Openreach model' or the 'model') of the resource requirements within Openreach's Service Delivery division and how these vary with the QoS provided. Analysys Mason has been commissioned by Ofcom to carry out an assessment of the model to support a separate public consultation on the model's outputs.

This document is our final report, covering our assessment of the model. It is not a standalone description of the model or a user guide and so should be read in conjunction with Ofcom's consultation document and with a description of the model methodology (which we understand will be provided to Ofcom by Openreach).

1.1 Scope of assignment

Our assignment is split into two broad parts:

- **Part 1 – an audit of the model**: a description and critique of the model, covering the model design, its inputs, assumptions, operation, limitations and how accurately it reflects the operational organisation and processes being modelled.
- **Part 2 – testing of the model**: a review of the modelling scenarios specified by Ofcom, describing how they have been implemented by Openreach, the inputs, assumptions and limitations of the modelling and a verification of the results provided by Openreach.

In both cases we also make recommendations for changes that we consider necessary, and provide an outline analysis of the possible impact on the results.

1.2 Structure of this document

In summary, the remainder of this document is laid out as follows:

Part 1: model audit	• Section 2 describes the way in which the model is constructed and operates. It outlines the three main sections of the model and the ways in which they interact when the model calculates results for a given scenario.

- Section 3 discusses the structural integrity of the model together with the assumptions highlighted in Section 2, and examines in more detail the role of the gamma distribution and other key methodological choices and parameters.

Part 2: model testing

- Section 4 contains a discussion of the results of the model both in the baseline scenario and as the required QoS metrics are adjusted from the baseline. The results are generated based on the actual set of repair and provision jobs carried out by Openreach in both 2011/12 and 2012/13.
- Section 5 contains a discussion of the changes made to the model by Openreach to enable it to run scenarios with hypothetical alterations to the profile[1] of repair jobs carried out by the Service Delivery division in 2011/12 and 2012/13, as well as the results from those scenarios.

Conclusions

- Section 6 provides our overall conclusions for both parts of our work. In this section we set out our overall views on the model, including what we consider to be the assumptions and limitations of greatest concern.

1 Specifically to the mix of repair jobs designated to each care level, a concept we expand upon in Section 2.

2 Model construction and operation

This section describes the construction and operation of the model. It provides details of the granularity of data at relevant points throughout the model and also highlights any assumptions that are discussed further in later sections.

It should be noted that this section of the document is to be read in conjunction with the model description provided by Openreach. The two together provide an holistic description of the model, its construction and operation.

2.1 Overview of model

The model itself is built using a combination of Microsoft Excel and Simul8 and can be understood as following a three-stage process:

- calculation of the input parameters for a given scenario, carried out in Excel
- Simul8 discrete event simulation of the scenario, and export of the Simul8 results to Excel
- post-processing of the simulation results, carried out in Excel.

Figure 2.1 on the next page shows the various components of these three parts and how they interact with one another within the model itself. Each of the three parts is discussed in more detail in the following sub-sections.

The model considers repair and provision jobs undertaken during financial years 2011/12 and 2012/13. Analysis is conducted at a General Manager area (GM) level and the final output of the model is the resource level (communicated as the number of full-time equivalents, or FTEs, in the Service Delivery division of Openreach) required to achieve a specified percentage level of performance[2] against service level agreement (SLA) during an historical time period, for both repair and provision jobs.

The performances (specified separately for repair and provision jobs) can be either the actual performance achieved in that historical time period or an increased or decreased targeted set of performance levels (for repair and provision). By calculating the resource required to achieve the actual performance for a specified period of time (the ‘Baseline performance’) and then the resource required to achieve another specified performance level, the resource requirement delta between the two scenarios can be calculated. In other words, the proportional increase in resource that would lead to a given increase in performance (or the proportional decrease in resource that could achieve a lower performance) can be obtained.

2 That is, the percentage of total repair or provision jobs which are completed within the time limit specified by the SLA.

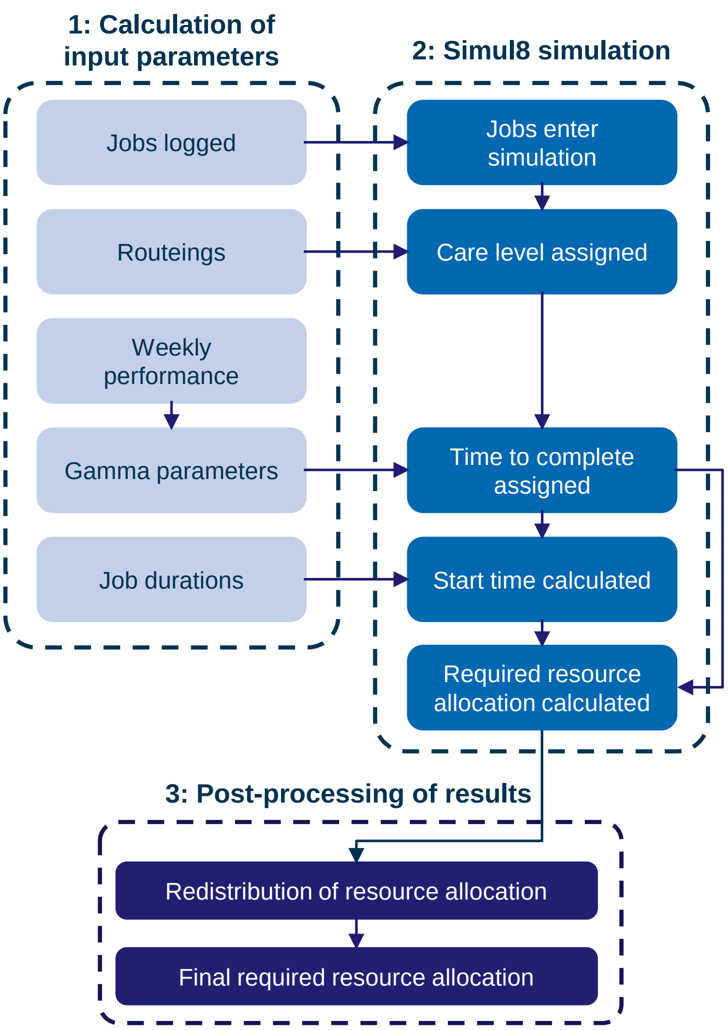

Figure 2.1: High-level flowchart of model operation [Source: Analysys Mason, based on Openreach model, 2013]

It is this proportional increase in resource requirement (the resource requirement delta) that is of most interest in this report.

We note at this stage that the model itself is not a ‘queuing model’ in the traditional sense. In a traditional queuing model the incoming jobs arrive in a queue at a specified rate and are held in the queue waiting to be completed by a fixed number of servers (resources). The inputs to such a model are the arrival rate into the queue and the amount of resource (servers); and the output is the time at which jobs are completed, and by extension, the performance achieved. However, the Openreach model works in reverse: the input to the model is the performance to be achieved, together with a pre-defined number of jobs and the time at which they are logged, from which the time at which jobs must be completed, and hence started, is calculated. The distribution of the job start times in effect drives a required level of resource. This required level of resource is therefore the output of the model.

The progression of a job through the model, together with the terminology we use throughout this report to describe the timings and durations associated with jobs is set out in Figure 2.2 below.

Figure 2.2: Progression of a job through the model [Source: Analysys Mason, 2013]

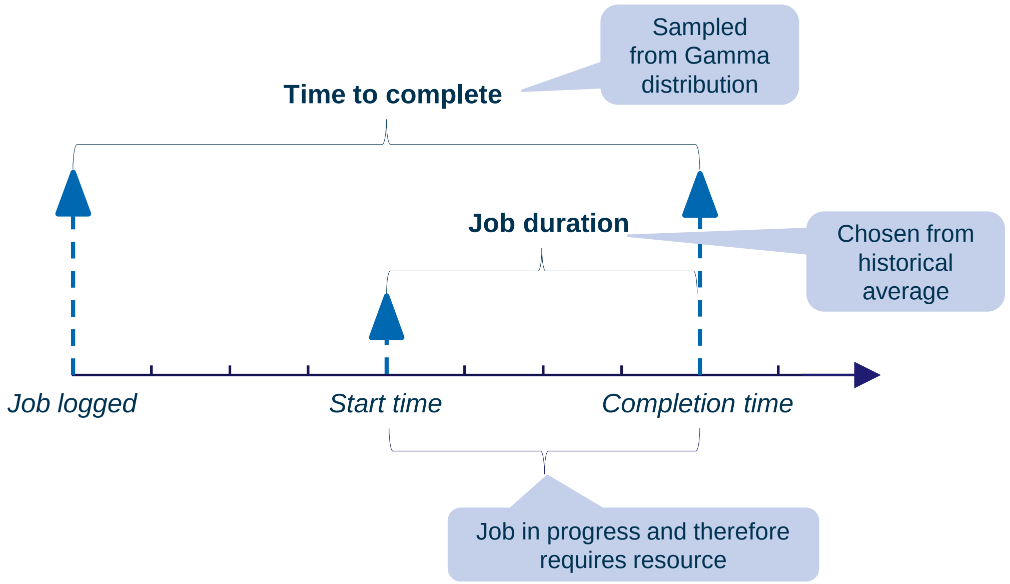

As with any model, the Openreach model is only an approximation of the real world. Openreach operates a very complicated near-real-time process for determining resource requirement and matching Service Delivery staff to jobs. This model does not, and does not seek to, fully replicate that process, but instead models the problem in a different way as described above. We return to comment on the applicability of the model to the problem it is trying to solve in our conclusions (Section 6), after considering various aspects of its functionality in more detail in the following sections.

2.2 Calculation of input parameters

The following sections describe how the various input parameters required by the Simul8 discrete event simulation model (shown in Figure 2.1 above) are generated in the input Excel workbook.

2.2.1 Jobs logged

The 'jobs logged' input data contains the number of calls logged per half day (morning and afternoon) by GM (General Manager area) and site type[3]. The job data is historical data of actual jobs carried out by Openreach in the financial years 2011/12 and 2012/13.

This historical data is then used to calculate the rate at which jobs will enter the simulation during a half-day period for a given GM and site type.

This calculation assumes that jobs are to be logged at regular intervals throughout each half-day period, and as such, the output parameter from this calculation is the number of minutes that pass between each pair of jobs entering the simulation at a given site type within a given GM.

3 Site type refers to the locations where the faults are cleared within Openreach's network (D Side, E Side, etc.)

2.2.2 Routeings

Openreach classifies repair jobs with different SLAs under a series of headings commonly referred to as the ‘care level’. There are four different care levels in the classification system:

- Care level 1 (CL1) – the fault must be repaired by 23:59 on the next working day plus 1, Monday to Friday, excluding Public and Bank Holidays. For example, a job logged on a Tuesday must be repaired by the end of Thursday
- Care level 2 (CL2) – the fault must be repaired by 23:59 on the next working day, Monday to Saturday, excluding Public and Bank Holidays. For example, a job logged on a Tuesday must be repaired by the end of Wednesday
- Care level 3 (CL3) – a fault logged by 13:00 must be repaired by 23:59 the same day. A fault logged after 13.00 must be repaired by 12:59 the next day, seven days a week, including Public and Bank Holidays
- Care level 4 (CL4) – a fault must be repaired within six hours of being logged, any time of day, any day of the year.

The routeings input data contains the split of jobs across care levels in each month by GM and site type.

This historical data is then used to calculate the proportion of jobs that will enter the simulation at each care level for each GM, site type and month within the scenario period. Within the simulation this proportion of jobs of each care level is then used in conjunction with the “jobs logged” data to create a flow of jobs of each care level into the system.

2.2.3 Generation of weekly performance targets

The data and corresponding calculation relating to weekly performance creates a vital parameter that enables the overall performance level against a specified target (actual or modified SLA) for the scenario in question. This calculation forms the basis on which the target performance for a given scenario is achieved by defining the required manipulation of weekly performances.

This process begins with the actual historical performance achieved week by week for each care level, in each GM. From this data, the actual performance achieved across the scenario period can be calculated by simply looking at the total number of jobs logged and the total number of jobs completed within the given SLA. This actual performance can then be compared to the target performance for that scenario to obtain the number of additional jobs that must now be either completed within the SLA (if the target performance is higher than actual performance) or not completed within the SLA (if the target performance is lower than actual performance).

In the case of a target performance higher than the actual performance, these additional completed jobs are distributed equally across each week within the scenario period. However, should any given week exceed a specified upper limit on performance (the ‘glass ceiling’ – set at 90% performance and discussed in more detail in Section 3.4) then any jobs over this upper limit are redistributed to the ten weeks[4] with the lowest performance. This process of redistribution is iterated until either the overall performance is within a specified tolerance, or a specified maximum number of iterations have been completed.[5] This process is discussed in Section 3.4, where we examine the impact of the glass ceiling, as well as the choice of ten weeks when redistributing jobs.

Figure 2.3 below illustrates the variation in performance across simulation weeks. It shows the actual (baseline) performance, the adjusted performance (adjustment to a target of 85%), where performance above the glass ceiling of 90% can be seen, illustrated by the dotted line on the chart, and the redistributed performance, where imposition of the 90% limit results in an increase in the performance of lower-performing weeks.

Figure 2.3: Illustration of performance adjustment and redistribution (illustrative only) [Source: Analysys Mason, based on Openreach model, 2013]

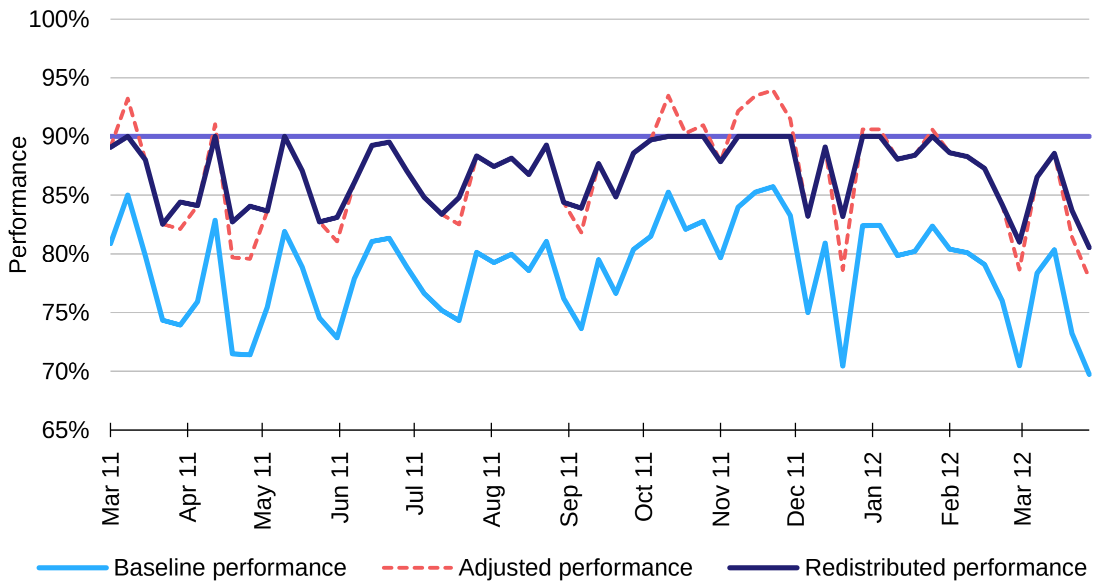

The final output of this process is the manipulated (i.e. target) week-by-week performance levels for each care level in each GM. These represent a set of target performance levels to be achieved in the scenario. They are then used to select gamma distribution parameters for the simulation, as described in the following section.

2.2.4 Gamma parameters

The model uses a gamma distribution to sample the time taken to complete jobs (the exact process by which this occurs is described in Section 2.3). As such, appropriate gamma distribution

4 There is a 'number of weeks to redistribute performance over' parameter in the model which is, by default, set to ten weeks, but it does allow different numbers of weeks to be tested.

5 The model has a default parameter value of 25 for the maximum number of iterations. To the best of our knowledge, the only purpose of this parameter is to allow the computation time to be speeded up.

parameters (alpha and beta) must be calculated to create a distribution representing the actual distribution of times to complete jobs for different SLAs and levels of performance.

The process by which these gamma parameters are calculated, their fitness for purpose, and their overall effect on the model are discussed in detail in Section 3.2.

The input Excel workbook contains a list of alpha/beta pairs, with each representing the model's distribution of time to complete jobs for a given SLA and performance target (defined in performance target increments of 1%). These sets of alpha/beta pairs are then mapped onto the target weekly performance levels (calculated in the previous section) to provide the simulation with gamma distribution parameters for all care levels and GMs in any given week.

2.2.5 Job durations

The job duration parameters that are used within the simulation are the average number of hours that a job takes to complete, once started, by GM, site type and month. There is also functionality within the model to alter these job durations by a specified percentage (to either increase or decrease the job duration).

2.3 Simul8 discrete event simulation

Once the required input parameters have been calculated, the model initiates the Simul8 discrete event simulation. Each simulation looks at a period of dates for one GM; to produce a national result requires calculation for nine GMs. Before the simulation can begin it goes through a simple process of setting the scenario parameters: the start date, run time and GM are read in and all Simul8 internal spreadsheets are cleared and populated with the input parameters needed for the current scenario. After these processes are complete, the simulation is started and jobs begin to enter the system.

Jobs flow through the simulation as 'entities'[6] and move between 'activities'[7] that perform a specific task (e.g. care-level assignment). All entities enter the simulation at 'entry points'; in this case, the entry points are the site types at which jobs are logged.

2.3.1 Jobs entering simulation

Jobs enter the simulation by site type[8]. Each site type logs jobs at regular intervals. These intervals are chosen from the set calculated in the way described in the previous section, according to the site type, GM and half day of the current scenario.

6 Entities are the items that flow through the system; in this instance they are single jobs that have been logged by a customer.

7 Activities process entities as they move through the system, by performing predefined actions on them (e.g. routeing the entities on to the next process or holding them for a defined period of time).

8 Site type refers to the locations where the faults are cleared within Openreach's network (D Side, E Side, etc.).

2.3.2 Care-level assignment

It is at this stage that the care level of each job is assigned. The proportion of jobs assigned to each care level from a site type during a given month and GM is dictated by the routeing parameters (from the input calculations). It is this process which ensures that the correct number of jobs are completed at each care level.

2.3.3 Time-to-complete assignment

Each job's care level, along with the week of the simulation, is used to look up the relevant gamma distribution parameters from the input Excel workbook. This alpha and beta pair of parameters is then used to create a unique gamma distribution from which the time to complete the job is randomly sampled.

2.3.4 Start-time calculation

The completion time for each job is calculated from the time at which the job is logged and the time-to-complete data. With the completion time of each job in the system known, the simulation then uses the average job durations previously calculated to obtain the start time of each job. Each job in the system now has a start and completion time; in other words, a period of time during which the job requires resources.

2.3.5 Calculation of required resource allocation

As all the jobs entering the system now have a start and completion time, it is possible for resources to be allocated. As the simulation runs and jobs become active, resources are created to satisfy demand. A resource is created when there is no available resource at a skill level equal to (or higher than) that required for the new job. This process assumes no prior knowledge of the jobs occurring later in the day and begins anew for each day of the simulation period.

An implication of this approach is that once a resource is allocated to a job, the resource cannot be swapped to a different job. This might be useful, for example, if a new job which can only be handled by a high skill level resource enters the system and only lower skill level resources are available, whilst some higher skill level resources are engaged on jobs which could be carried out by lower skill level resources. In such a situation the model would still create a new high skill level resource for the new job rather than re-allocating resources between jobs. On balance though we believe that this is likely to be a reasonable approach.

Resources are allocated in this way throughout each day of the simulation period. Once the simulation is over, the final output that is exported is the day-by-day resource requirement by skill level.

2.4 Post-processing of results to redistribute resources

The output of the Simul8 discrete event simulation is the day-by-day resource requirements; however, the desired output of the model is the required resource level to obtain a given performance across a specified time period. In order to obtain this, the simulation results are post-processed.

This is done within the output Excel workbook, through a process of resource redistribution. The key concept is that the number of resources (technicians) of a given skill level are not necessarily all required to work on jobs requiring that level of skill on all days, or at all times of a given day. Openreach is therefore able to re-allocate more highly skilled technicians, when they are available, to jobs which only require a technician of a lower skill level.[9] This re-allocation of more highly skilled resources to lower skilled jobs is referred to in the model (and this report) as resource redistribution.

The approach followed in the model considers the day of the year on which the maximum number of technicians of a given skill level are required (the 'Maximum day').[10] Summing the Maximum day requirement across all skill levels would certainly provide enough resources to meet the target performance level. However, if a resource redistribution approach is followed then less resource may be required.

The resource redistribution algorithm uses one parameter, N, to determine its function. N is defined such that the resource redistribution approach is based on the N days with the highest resourcing requirements for each skill level (the approach referred to by Openreach as the 'Top *N* days'). For N=1 the approach is referred to as the 'Maximum day' redistribution approach. The output workbook calculates results for four different values of the parameter N: 1 (i.e. the Maximum day approach), 10, 20 and 25. All four of these results are then presented in the final output worksheet.

The algorithm begins by taking the maximum number of technicians of the highest skill level (CSE[11]) required across the scenario period; this number of CSE technicians is then set as the final required number by the model. This assumption leaves an amount of idle CSE resource on all days other than the day which generated the maximum requirement; that is, an idle level of resource that can carry out jobs of any lower skill requirements.

The exact number of technicians that are then assumed to be left over and available for re-allocation on any given day is calculated as the averaged CSE resource requirement on the top *N* days minus the CSE resource used on that day. The process is repeated for each skill level in turn, running from highest skill level (CSE) to lowest (OMI[12]).

9 Note that the skill levels within Openreach are strictly hierarchical, and so a technician of the highest skill level is always capable of carrying out any job designated as requiring a lower skill level.

10 Note that this maximum day may be (and almost certainly is) a different day for each skill level.

11 Customer Services Engineer.

12 One-Man Installer.

Figure 2.4 and Figure 2.5 illustrate this process for the Maximum day redistributed and the Top 25 days redistributed approaches respectively, for a single day in a particular scenario (London GM, 80% performance targets for both repair and provision). As can be seen in Figure 2.4, starting from the highest skill level (CSE) on the left the resource level increases as the maximum required at any given skill level is added, and decreases as the daily requirement of that skill type (and those above it) is used.

Figure 2.4: Chart illustrating redistribution of resources based on Maximum day approach (London GM, 80% performance target for repair and provision)[13] [Source: Analysys Mason, based on Openreach model, 2013]

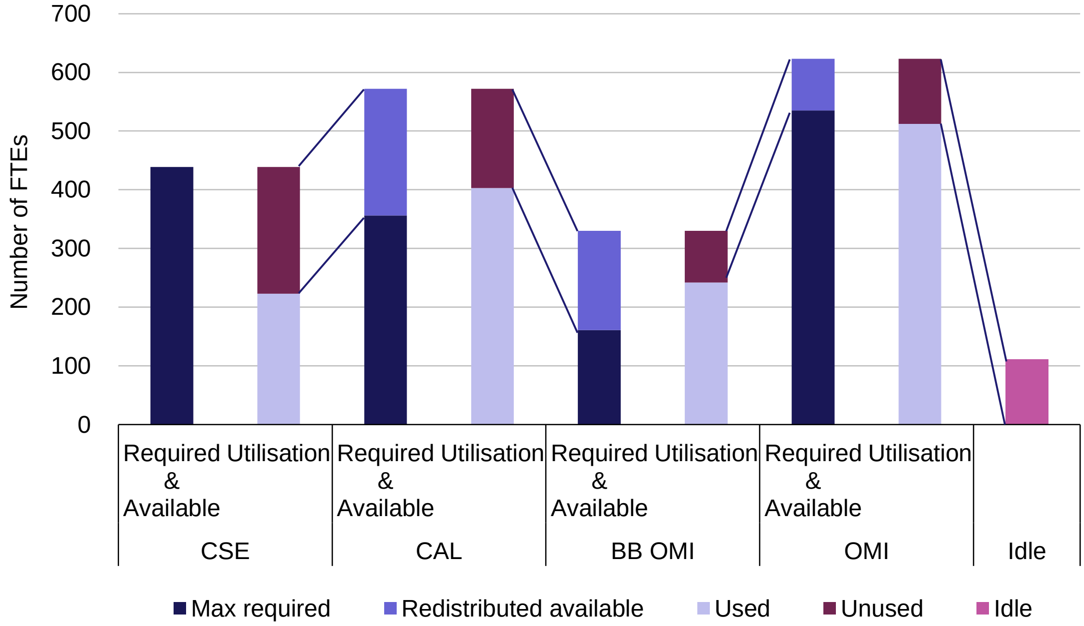

The model calculates the maximum number of CSE resource required on any day over the year, and assume that this number of resources is always available to complete CSE or lower skilled jobs. In the example above, the maximum number of CSE engineers ever required is around 430, and on the particular day shown in the figure, there are only around 220 CSE jobs. The model therefore calculates that there are around 210 CSE resources unused. Following the Maximum day redistributed approach all 210 of these unused resources are then made available to do the next lowest skilled jobs (i.e. those requiring at least a CAL engineer). Resources at the CAL level will only be added where there are no redistributed CSE resources available. This process applies all the way down the chain to the lowest skilled resource level. On any given day it is likely that a number of resources will be left unused, or "idle".

13 CAL = Customer Apparatus and Line, BB OMI = Broadband One Man Installer.

Figure 2.5: Chart illustrating redistribution of resources based on Top 25 days approach (London GM, 80% performance target for repair and provision) [Source: Analysys Mason, based on Openreach model, 2013]

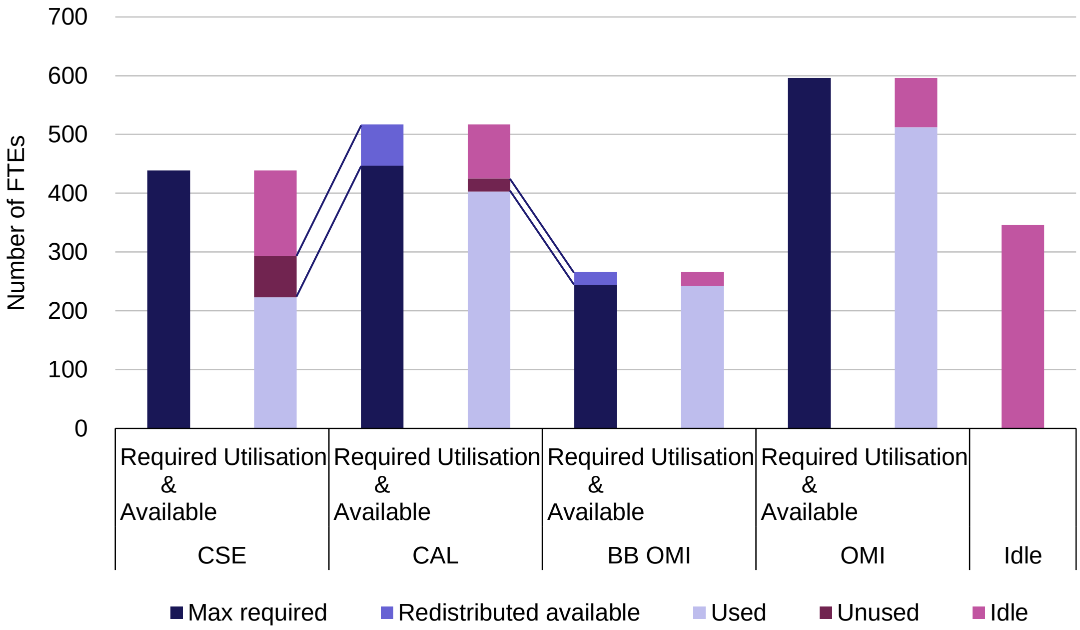

As Figure 2.5 illustrates, when considering the Top 25 days redistributed approach, the number of resources assumed to be available for redistribution is lower; i.e. the level of resource assumed to be available to carry out jobs of a lower skill level is reduced due to an assumed higher level of idle workers at each skill level. This in turn leads to an increase in requirement for workers of a lower skill level, and hence a higher total resource requirement. This, to us, is counter-intuitive: a point we discuss further in Section 3.3.

With reference to the example in Figure 2.5, the model again calculates the maximum number of CSE resources required (around 430). Again around 210 are unused on this day, however, this time the average of the top 25 days requirement is used to identify how many CSE resources are assumed to be available on any day to undertake lower skill level jobs. In this example, that is around 70 resources, which are redistributed to the CAL skill level jobs as before. The remaining 140 unused CSE resources are not redistributed and are left idle.

Whilst the Maximum day redistributed approach assumes that resources are only idle when the number of jobs at all skill levels below are lower than the available level of higher skilled resource, the Top 25 days redistributed assumes a greater proportion of idle time in the model. Indeed the number of idle resources can be calculated, as set out in the example above, for each skill level. The total number of idle resources is then the sum of those at each skill level. The total idle resource and those at each skill level is shown in pink in Figure 2.5.

The effects of this parameter on resource levels, both absolute and in relation to the baseline, are discussed in Section 3.3.

2.5 Final results of the model

The final results of the model are for each GM and consist of the total number of workers required for the modelled period as well as the performance actually achieved across all four repair care levels, as well as provision.

3 Structural integrity and assumptions of the model

This section discusses the structural integrity of the model and its fitness for purpose. We also detail the methodology by which we understand the gamma distributions have been generated, and examine the effects of the main assumptions which are made (as identified in Section 2).

3.1 Structural integrity of the model

There are three main points to consider when discussing the integrity of the model:

- Is the choice of modelling tool appropriate?
- Does the model function in the way it is designed and intended?
- Does the model create consistent, replicable results?

We discuss each of these points in turn below.

Choice of modelling tool

The choice of modelling tool itself seems appropriate to the task it is intended to perform.

The use of a Simul8 discrete event simulation appears justified, if a little slow (or involving overly intensive computation). Excel is adequate for the tasks it is used to perform.

Functionality

Based on our investigations of the Simul8 Visual Logic programming statements, Excel formulae and the general architecture of the model, these appear to function as intended. This conclusion was reached following a comprehensive review of all Visual Basic code, as well as a sample of the Simul8 Visual Logic statements that control the major activities of the simulation.

Consistency/replicability

We are interested in the consistency of results, because random numbers are used within the model (samples from statistical distributions).

Simul8 allows for the random number generator to be seeded and the seed fixed so that subsequent runs of the same scenario, on the same machine, will yield identical results. This has been confirmed by the replication of results from identical scenarios.

However, we understand that due to design decisions within Simul8, these random numbers are unique to individual releases of the software; as such, the same scenario run on different machines with the same version of Simul8 (but a different detailed release) can yield a slightly different result.

This can be seen when comparing the results of runs on our installation of the software with those of Openreach / Ernst & Young. However, as discussed in Sections 4 and 0 of this report, our results have been identical to those shared by Openreach / Ernst & Young to two decimal places.

We would expect that the end result should not greatly depend on the choice of random seed, but note that this has not been tested extensively.

3.2 The gamma distribution, the way it is changed at different target performance levels, and its impact

As described in Section 2.2.4, the model uses a randomly sampled gamma distribution to estimate the time to complete jobs. The gamma distribution is a continuous probability distribution which is commonly used in applications relating to queuing models and the calculation of weighting times. The exact probability density function depends on the choice of the parameters alpha and beta, however, the distribution can commonly be thought of as resembling a skewed normal distribution. Examples of different gamma distribution probability density functions are shown in various diagrams throughout this section.

The gamma distribution is used in the Openreach model to define the probability of any given time-to-complete occurring; with a statistically large enough number of jobs (as there are in this model) this probability distribution becomes the number of jobs completed at any given time-to-complete.

Use of a gamma distribution is not *a priori* unreasonable (as such distributions do arise in queuing theory).

How the baseline gamma distribution is generated

Openreach has not provided us with details of the exact method by which an initial (baseline) gamma distribution was chosen to fit to the empirical data. However, a histogram of the empirical data is available and the model has the functionality to run scenarios using this empirical distribution in place of the gamma distribution, allowing comparisons to be made, as described in the following sub-section.

Comparison of baseline gamma, and empirical distribution

As the model contains the functionality to run scenarios using the empirical distribution as well as the gamma distribution it allows direct comparison to be made between the two. Figure 3.1 shows the empirical distribution of times to complete as a histogram (the form in which the data is found in the model)[14] and the corresponding baseline gamma distribution. In the example shown in Figure 3.1, the data relates to all care level 1 repair jobs (which have a next working day plus one time to complete SLA) in 2011/12 across all GMs.

[14] The model provides this data by GM, while Figure 3.1 shows the total distribution across all GMs, which required processing of the raw data by Analysys Mason.

Figure 3.1: Comparison of empirical distribution of times to complete and the corresponding gamma distribution for care level 1 repair jobs in 2011/12 [Source: Analysys Mason, based on Openreach model, 2013]

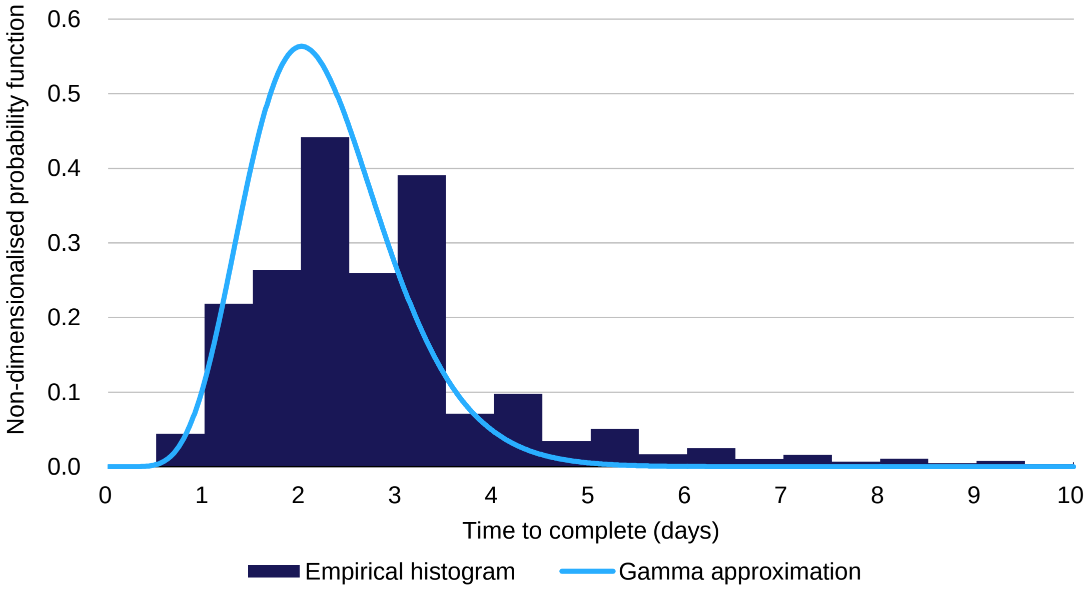

As can be seen in the diagram, the gamma distribution approximation appears a reasonable but by no means perfect fit to the empirical data.

When considering provision jobs the approximation is not as good. Figure 3.2 shows the time-to-complete distribution for all provision jobs (which are assumed to have an implicit 13-day time to complete SLA) in 2011/12 as both the empirical distribution (shows as a histogram) and the gamma approximation. As can be seen, the approximation is very rough and does not accurately represent a large number of jobs which have short times to complete. This is because the methodology of fixing the mode of the gamma distribution at one day prior to the SLA target (discussed further in the following section) leads to the approximation shown, which concentrates a large proportion of job completions just before the SLA target (13 days in this instance).

A can be seen in the diagram, the gamma distribution approximation is far from a good fit to the empirical data.

Figure 3.2: Comparison of empirical distribution of times to complete and the corresponding gamma distribution for provision jobs in 2011/12 [Source: Analysys Mason, based on Openreach model, 2013]

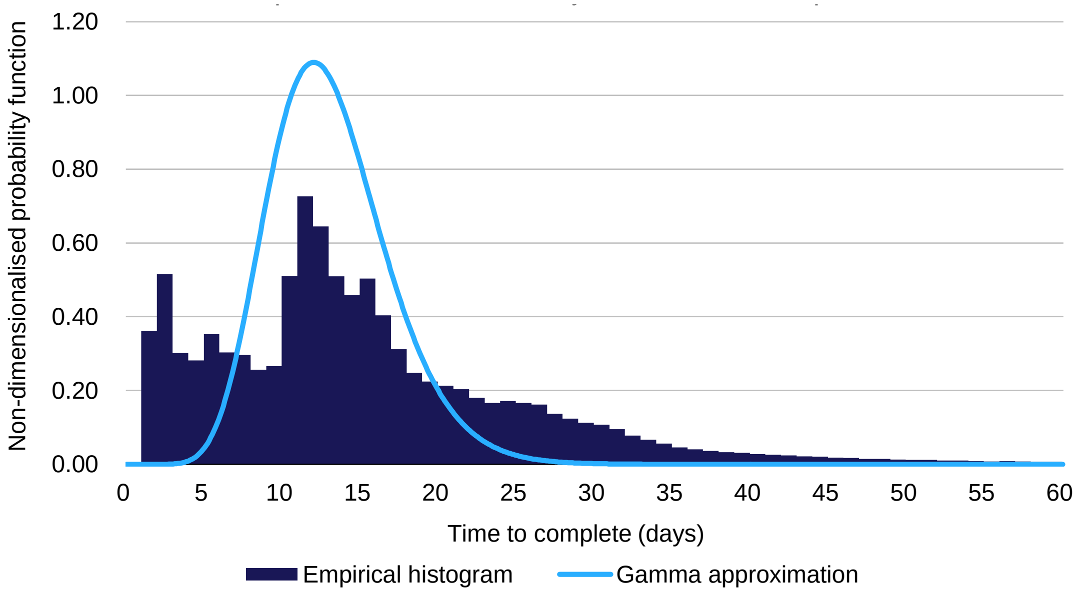

Despite this, it is not clear that any other statistical distribution would have been a better choice. In particular, any statistical distribution will have a similar peak: the mode is the highest value, by definition. We therefore focus on the impact of using the gamma distribution.

To assess the effect that this approximation has on the result, the model has been run using the empirical data for a baseline performance scenario; results of this are shown in Figure 3.3. It can be seen that the baseline gamma distribution produces results very close (just over 1% delta) to those obtained when using an empirical distribution for 2011/12 data, thereby validating the shape of the baseline gamma distribution in this specific case.

Two further scenarios were run for the London area only, to compare results arising from the two different distribution approaches (see Figure 3.3):

- First, repair performance was set to 80% and provision kept at the baseline level. It can be seen that the result is very close to those obtained when using gamma distributions
- Second, repair and provision performance targets were both set to 80%. Again, the result is close to that obtained when using the gamma distribution, although not quite as close as under the previous scenario.

	Gamma	Empirical
National baseline resource requirement (FTEs)	✂	✂
London baseline resource requirement (FTEs)	✂	✂
London; repair performance 80%, baseline provision	✂	✂
	-1.4%	-1.45%
London; repair and provision 80%	✂	✂
	4.47%	4.15%

Figure 3.3: Comparison of baseline gamma distribution and 2011/12 empirical data resource requirements [Source: Analysys Mason, based on Openreach model, 2013]

How the gamma distribution is modified for different scenarios

The original gamma distribution that has been fitted to the empirical data is adjusted for non-baseline scenarios to give the desired properties, namely for an adjusted level of performance to be achieved for a particular care level SLA.

However, given that a gamma distribution requires two parameters, there are multiple possible pairs of these parameters that would yield distributions with the same required performance level for a given care level SLA.

In order to obtain a unique pair of parameters to be used in the model, an extra constraint was needed to ensure a unique pair could be found; the extra constraint used by Openreach / Ernst & Young concerns the mode of the distribution.

The two constraints imposed on the distribution by Openreach / Ernst & Young are:

- The cumulative probability at the care level defined SLA time-to-complete cut-off (i.e. the probability that the job is completed in less time than the SLA allows for) is fixed; i.e. the performance achieved within the SLA for that distribution is specified.
- The position of the mode of the distribution is fixed; the mode was chosen to be fixed at one day before the SLA; i.e. an SLA of two days would have a distribution with a mode of one day, with the exception of jobs with a one-day SLA (which have a mode of half a day). This was reportedly based on observations made from the empirical data for different care levels.

With these two constraints, unique alpha–beta pairs can be calculated and hence a unique gamma distribution can be specified. Figure 3.4 shows the variation of the distribution of times to complete for different SLAs for a given performance (in this instance, 80% performance achieved). The shorter the SLA, the higher the peak on the distribution; in the model, this would drive a higher peak required resource.

Figure 3.4: Variation of time to complete distribution with SLA for an 80% performance scenario [Source: Analysys Mason, based on Openreach model, 2013]

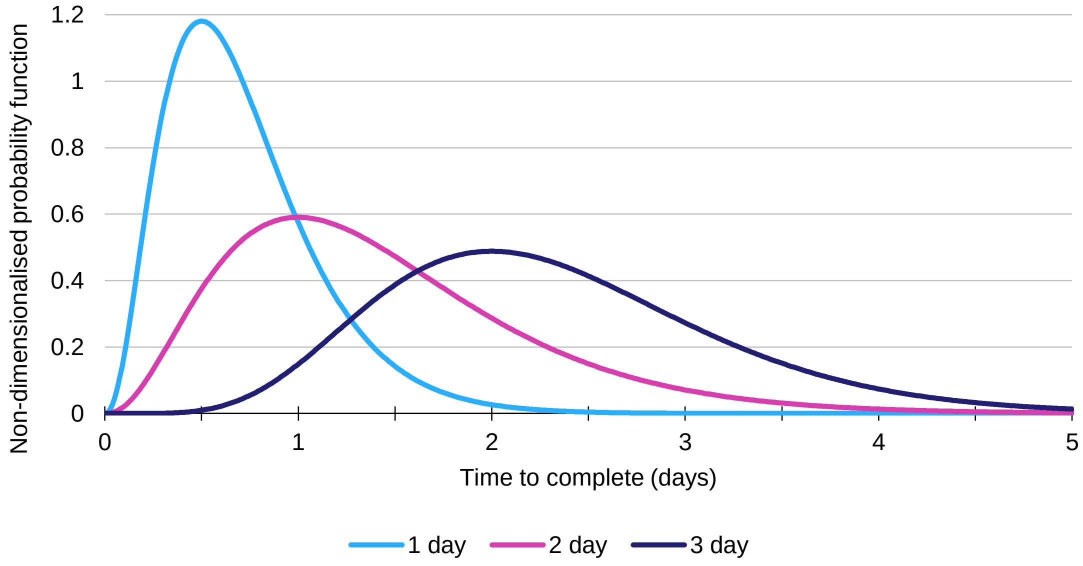

Figure 3.5 shows the variation of the distribution of times to complete for different performances achieved for a given SLA (in this instance, two days). As can be seen in the examples chosen, the mode of the distributions remains one day ahead of the SLA, and as the performance increases so does the peak of the distribution. In the model, this increase in the height of the peak will drive an increase in the required resource.

Figure 3.5: Variation of time to complete distribution with performance for a 2-day SLA scenario [Source: Analysys Mason, based on Openreach model, 2013]

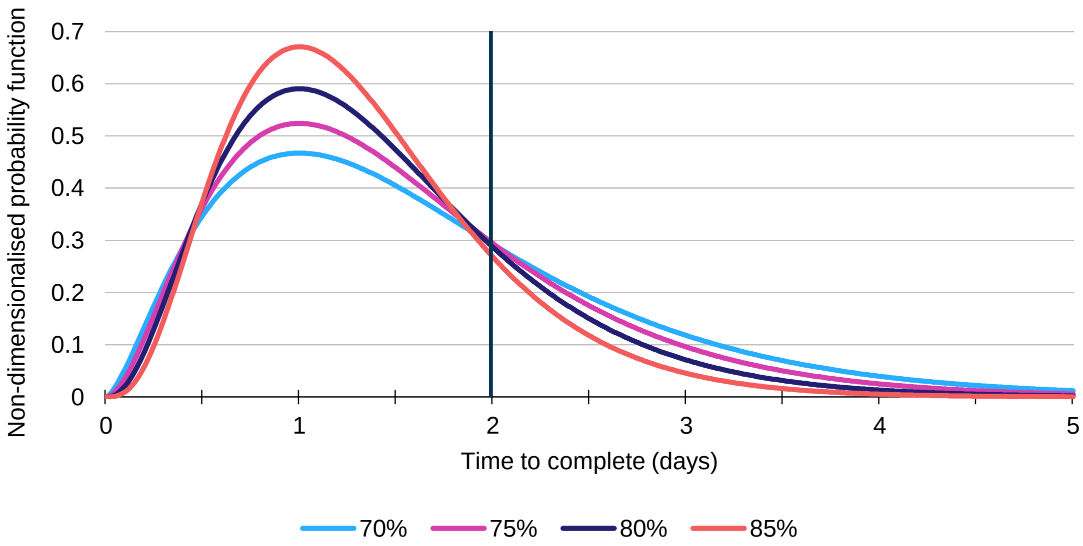

However, this methodology relies heavily on the assumption that the mode of the distribution remains at a fixed point. While that may have been an accurate observation from the empirical data (although we have not seen direct evidence of this), there is nothing *a priori* to suggest that this should continue to be the case when the performance is adjusted. There does, of course, have to be a second constraint placed on the distribution, in order for a unique alpha–beta pair to be calculated, but it is unclear whether a fixed mode is the most appropriate choice.

In particular for provision jobs it seems that fixing the mode 1 day prior to the SLA when high performance targets (e.g. 85%) are used would have the effect of overestimating the increase in the modal value. This is because a very large number of provisioning jobs will be forced to complete in the one day period prior to the SLA cut-off. The high modal value equates to a high peak resourcing requirement and will therefore drive a high number of resources.

As a demonstration of the potential non-unique solutions for alpha and beta, Figure 3.6 shows four distributions each representing the same performance achieved within the SLA (in this instance, 80%). It depicts the distribution calculated by fixing the mode at one day ahead of the SLA (the approach followed by Openreach), as well as two other distributions where alpha has either been increased or decreased by a small proportion and a corresponding beta calculated. As shown, they yield distributions with differently placed peaks, not only in relation to the care-level SLA, but more importantly, in magnitude. As it is the magnitude of the distribution's peak that in the end drives the required resource[15], this is of great importance.

It is clear from this that alpha–beta pairs could have been chosen to either decrease, or increase, the peak of the distribution. While choosing a constraint such as the mode does not explicitly bias the model, we do not have enough evidence to determine whether an assumption that the mode is unchanged is actually borne out in reality.

Figure 3.6: Distributions with the same performance properties but different alpha–beta pairs for a 2-day SLA and 80% performance target [Source: Analysys Mason, based on Openreach model, 2013]

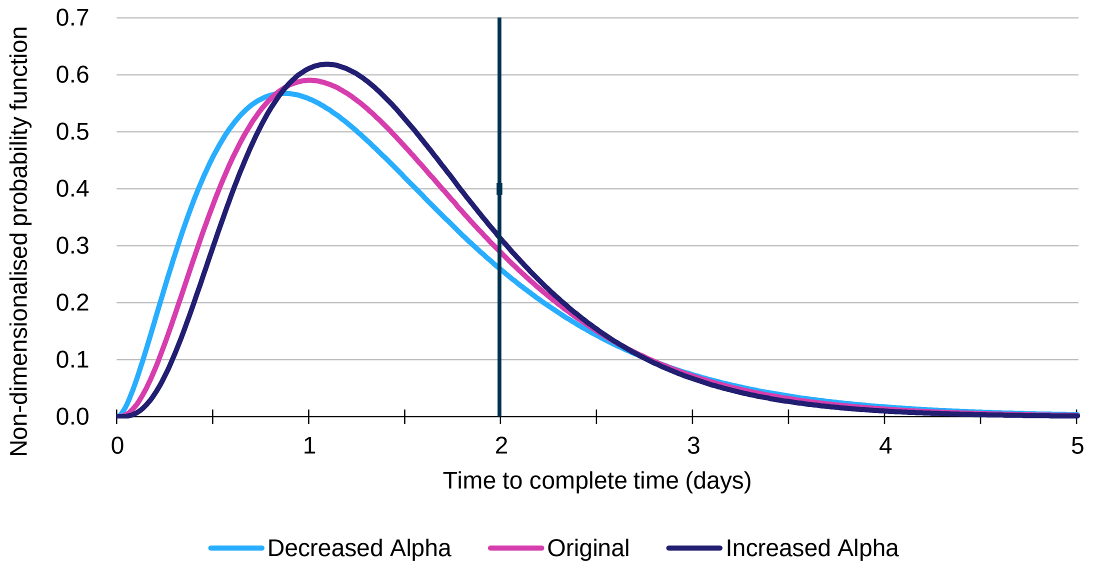

15 Temporal fluctuations in the job arrival rate have been smoothed out by the way in which jobs are assumed to arrive at a constant rate per site type, leaving half-day to half-day fluctuations in job arrival rate and the shape of the distribution as the key parameters.

Importantly the magnitude of the effect of fixing the mode of the gamma distribution on the calculated resource requirements is unclear. This is an important point because, even though we do not consider that the decision to fix the mode is necessarily unreasonable, an alternative choice could have been considered equally reasonable and could have yielded very different outputs across the range of scenarios we consider in later sections.

How the modified gamma distribution is then used

In the model, the gamma distribution parameters are chosen from a look-up table (a list) giving an alpha–beta pair for a specified SLA and performance target. The index used provides a different pair of parameters for target performances varying in increments of 1%. It would have been possible to include a larger number of pairs of gamma distribution parameters to provide distributions for different performance targets at a granularity of greater than 1%. However, given the number of values being chosen for each scenario (the target performance level for each site type and care level in each week) and the granularity of the data, it is unlikely that this 1% granularity (rather than having greater granularity) creates a significant bias in the results.

3.3 Redistribution of resource allocation

As described in Section 2.4 earlier, the process of resource redistribution resulted in a certain number of resources remaining idle on the majority of days. We discuss this issue in more detail in this section. We also mentioned that the process of redistributing using the Top *N* days approach leads to an increase in the required resource due to the need for an increased number of workers at lower skill levels. The effects of this parameter on the total resource required and, more importantly, the resource delta between scenarios is also described in this section.

The impact of idle resources

As described in Section 2.4, the way in which resources are allocated and subsequently redistributed leads to an amount of idle resource being available on the majority of days. In reality this resource could be used to carry out jobs that are in the system and, as such, increase the performance achieved on those days.

As a whole this would have the effect of increasing the overall performance of a given scenario above the level which is specified in the scenario, though it is not possible to calculate the exact impact of this effect. By way of example, Figure 2.4 demonstrated that for a scenario in the London GM with ✂ total FTEs serving that GM, ✂ were idle on the day which was being considered. This constitutes around 7.4% of the total resource but the impact of this level of idle resource, even if known for every individual day in the modelled period, cannot easily be translated into an increased performance level based on the existing functionality of the model. This translates to a level of uncertainty in the model results discussed in Section 4; it is clear that the absolute level of resource is overestimated but the effect of this on the resource delta is not known.

We note that the outputs of most interest are the resource requirement deltas between scenarios, but an x% overstatement of resource requirements in any pair of scenarios may still affect the delta.

The approach to redistribution using the top N days average

As was described in Section 2.4, the redistribution algorithm assumes that the level of resource available to redistribute at any given skill level is the average resource requirement in that skill level for the top *N* days, but still takes the maximum day requirement for each skill level to obtain a total resource level for that scenario. In essence it understates the amount of resource available for redistribution to jobs of a lower skill level, leading to a greater total resource being required.

This methodology is counter intuitive, as, the greater the value of *N*, the higher the level of idle resource that is created for any given day in the scenario.

We consider that a more logical method would be to use the average resource requirement for the top *N* days as both the resource to redistribute as well as the required resource per skill level. Of course this would lead to there being days in which the final resource level is insufficient to match the requirement (although there are also many days with idle resources), but it would go some way towards eliminating the situation where a few key days may possibly dictate the magnitude of the result.

Without this adjustment we consider the top *N* days scenarios for *N*>1 not to be a useful measure of resource requirement.

Figure 3.7 illustrates the weekday resource requirement in terms of total full-time equivalents (FTEs), summed over all skill levels, under different methods of redistribution for a given scenario (80% repair performance and 80% provision performance for the London GM).

As can be seen, the ‘max day redistributed’ approach leads to a resource level which sits just above the peak daily resource requirement; the current methodology for redistributing the top 25 days averaged leads to a resource level which is (counter-intuitively) higher than this, in line with what has been discussed above. The revised (AM) methodology for redistributing the top 25 days averaged leads to a resource level which is lower than both the max day redistributed approach and the peaks for the few days where the daily resource requirement is highest.

Figure 3.7: Weekday resource usage relative to total required resource with different redistribution methods (illustrative only) [Source: Analysys Mason, based on Openreach model, 2013]

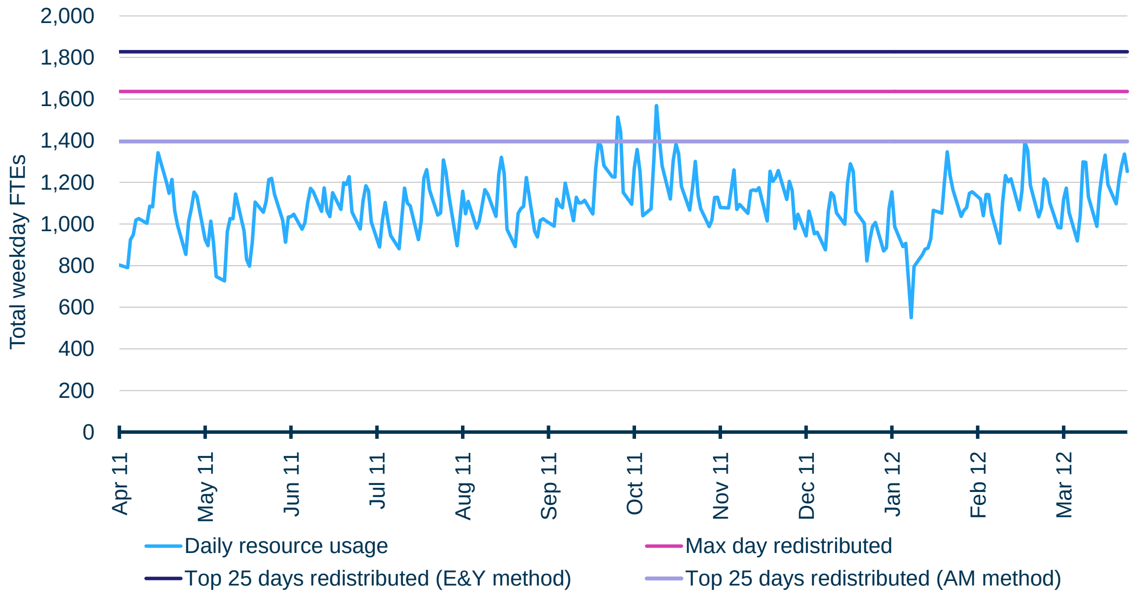

However, it is not the absolute resource levels that are of greatest interest, but rather the *increase* in required resource from the baseline.[16] And as such these methodologies must be applied to create deltas in resource requirements from their respective baselines. Figure 3.8 shows the absolute resource levels, as well as the delta from baseline resource[17], for a scenario with the following parameters:

- national geography (i.e. all GM areas)
- 2011/12 dataset
- 85% repair performance
- 12-day provision SLA
- 85% provision performance
- no adjustment to job durations from raw data provided for both provision and repair jobs.[18]

16 As described in Section 2, the baseline scenario is the scenario in which actual historic performance levels are achieved. In other words, the actual performance levels for repair and provision for the year in question are entered as targets and the model calculates the resulting resource requirement in terms of FTEs.

17 Note that the baseline resource requirements are different for each of the different redistribution approaches shown in Figure 3.8. Other than the redistribution approaches, each baseline scenario shown is identically defined.

18 We note that when Openreach runs scenarios with 85% performance targets average job durations are inflated, as discussed in Section 4. The scenarios shown here were chosen only to illustrate the effects of the redistribution parameter as they give large resource deltas which allow trends to be seen more easily.

		Redistributed top N days			
		Max	10	20	25
Ernst & Young method	Baseline	✂	✂	✂	✂
	85%	✂	✂	✂	✂
	Delta	8.11%	5.10%	3.87%	3.78%
AM method	Baseline	✂	✂	✂	✂
	85%	✂	✂	✂	✂
	Delta	8.11%	4.49%	2.57%	0.55%

Figure 3. : Results table for changes in redistribution parameter for different methodologies [Source: Analysys Mason, based on Openreach model, 2013]

As can be seen from these results, when considering the original Openreach / Ernst & Young methodology, as the number of redistribution days (N) increases, the absolute resource required also increases (as has been previously noted). However, the delta from the baseline resource level actually decreases; in this instance from 8.11% in the case of the maximum day redistributed, to 3.78% in the case of redistribution of the top 25 days averaged. When considering the revised (AM) methodology of redistribution the absolute resource required decreases (as has been previously noted), as does the delta from the baseline; in this instance from 8.11% to 0.55%, in the case of redistribution of the top 25 days averaged.

If looking at the result for greater values of N is considered important, then the above illustration shows that a very different view of the resource delta is generated depending on whether the Openreach/E&Y or AM redistribution method is chosen.

We do not mean to suggest that the revised AM approach is necessarily the correct one, or indeed to suggest an appropriate value for N: the tested values appear to be arbitrarily selected by Openreach / Ernst & Young. However, these results do show that this parameter can have quite strong effects on the final result and that although a suitable value for N is not easy to justify, taking the max redistributed approach (i.e. N=1) may result in an overstatement of resource deltas for a change in SLA or target performance.

3.4 Other assumptions we have investigated

Other key assumptions made within the model concern the weekly performance calculation and the job durations (detailed in Section 2.2), which are considered in more detail in the following sub-sections.

For the weekly performance calculations, two sets of parameters have a direct impact on the calculated model inputs:

- the ‘glass ceiling’, a cap set on the maximum performance achievable in any given week and the way in which jobs above the glass ceiling are redistributed (across a specified number of weeks)
- the granularity of the job duration data.

These assumptions and their effects are detailed below.

3.4.1 Weekly performance calculation

The glass ceiling

The glass ceiling has a default setting of 90%; there are, however, several weeks within the historical data where actual performance exceeds this threshold. This would suggest that a higher glass ceiling could potentially be justified. To test the effects of this parameter the model was run with an increased glass ceiling of 92%.

For speed of computation these sensitivities were only initially run for four GM areas (in this instance, London, South East, South West & South Midlands, and Wessex) for the 2011/12 year and for only one performance scenario: 80% repair performance and 80% provision performance. Figure 3.9 shows the results of this sensitivity test in the form of a percentage increase in resource required from the baseline. As can be seen, when the glass ceiling is increased from 90% to 92% the resource requirement delta from the baseline also increases, from 6.18% to 6.48%; similarly, when the glass ceiling is reduced from 90% to 88% the resource requirement delta from the baseline also decreases, from 6.18% to 6.05%.

	88% glass ceiling	90% glass ceiling	92% glass ceiling
Resource requirement	✂	✂	✂
Delta from baseline	6.05%	6.18%	6.48%

Figure 3. : Results for alternative glass ceiling parameter values [Source: Analysys Mason, based on Openreach model, 2013]

These results represent the increase in resources required above the baseline in order to achieve the target performances.

The first thing to note is that different levels of 'glass ceiling' do not have a very significant effect on the resource delta at 80% levels of performance.

The second is to note that this is counter-intuitive: apparently removing a constraint makes it harder to meet the required performance. However, this result is consistent with the way in which the model uses the glass ceiling parameter. By increasing the glass ceiling, the performance in certain weeks is not just *allowed* to reach that level but actively *required* to reach it. With these high-performing weeks now required to reach an even higher performance, their resource requirement increases. Conversely, when the glass ceiling is lowered the performance redistribution algorithm described in Section 2.4 spreads the required performance among the bottom-performing weeks; by observation, this leads to a decrease in resource requirement, possibly due to the fact that increasing the performance of a poorly performing week requires less incremental resource than increasing the performance of an already high-performing week, at least in the case as modelled.

When this sensitivity is conducted on a national basis however, the results are a little different. The effect seen when increasing the glass ceiling to 92% is similar, with the delta from the baseline increasing by 0.7% (up from the 0.3% seen in the four GM test). However, when decreasing the glass ceiling parameter to 88% the result is also an increase in the delta from the baseline of 0.56% (up from the -0.13% seen in the four GM test). The reasoning set out above does not appear to apply in this example and is most likely due to this parameter having more subtle effects within the model that are very hard to trace. In particular, for small changes in the parameter the output may be heavily dependent on how resource requirements are changed in a small number of weeks which may drive the overall peak requirement.

Number of weeks over which performance is redistributed

The number of weeks over which performance is redistributed has a default value of ten weeks; that is, those jobs that cannot be performed in some weeks due to the performance glass ceiling are offset by increasing the performance in the ten lowest-performing weeks. To ascertain the effects of altering this parameter, the model was run with both higher and lower values.

As with the glass ceiling sensitivities these were carried out over a geographical area of four GMs (in this instance, London, South East, South West & South Midlands, and Wessex) for 2011/12 and for the following performance scenario: 80% repair performance, 80% provision performance. The results shown in Figure 3.10 are the absolute resource requirements as well as the percentage increases in resource required from the baseline. As can be seen, movement in either direction away from the default of ten weeks yields a decrease in required resource. This effect is far more prominent when changing from 10 weeks to 5 weeks, where the resource delta decreases from 6.18% to 4.96%. This result is in line with what is expected, as fewer resources are required to increase performance in lower-performing weeks. However, the result that the resource delta also decreases when the parameter is increased from 10 weeks to 15 weeks is more difficult to understand.

	5 weeks	10 weeks	15 weeks
Resource requirement	✂	✂	✂
Delta from baseline	4.96%	6.18%	5.97%

Figure 3. : Results for change in number of redistribution weeks [Source: Analysys Mason, based on Openreach model, 2013]

Figure 3.11 illustrates the variation in weekly performance across the simulation period for the three different redistribution parameter values, as well as highlighting the weeks in which the key resource requirement days sit. As can be seen, by varying the resource redistribution parameter, localised changes in the performance distribution can be seen in the lower-performing weeks. However, the key resource requirement days remain in the same two weeks. This implies that this parameter has more subtle effects on the resource requirements that are extremely difficult to trace through the model.

Figure 3.11: Variation in performance for different redistribution parameters, with respective key resource weeks highlighted [Source: Analysys Mason, based on Openreach model, 2013]

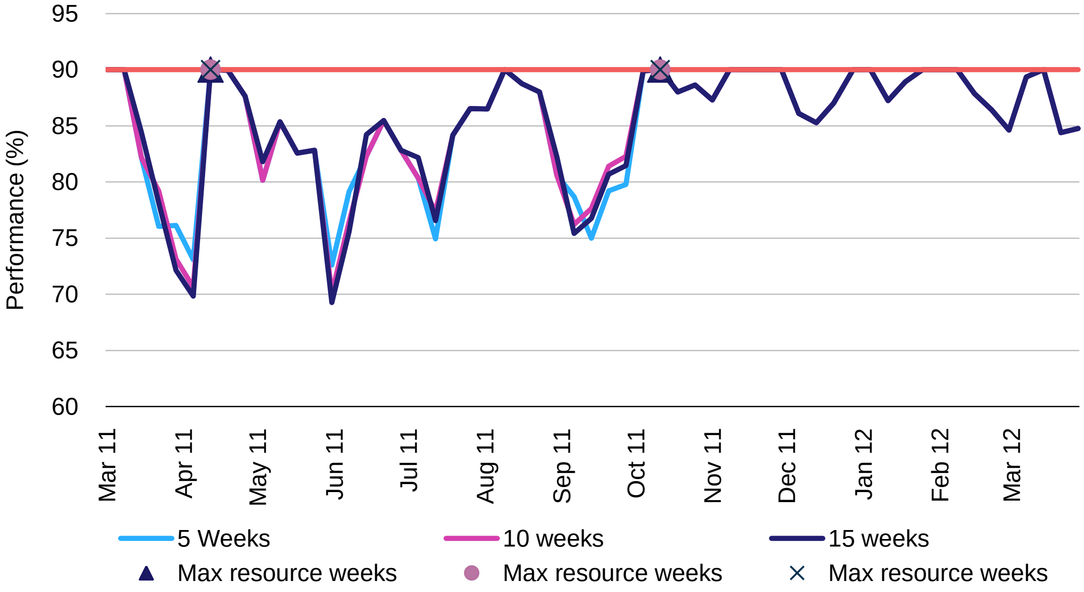

Although the choice of 10 weeks seems potentially arbitrary, it does not have a hugely significant impact on the results, unless parameter values of substantially less or substantially more than 10 weeks are used. And as noted above, the exact effects of this parameter throughout the model are very difficult to fully trace and understand.

Overall impact of the glass ceiling and associated redistribution

We were keen to understand the overall impact of including a glass ceiling and the associated redistribution of jobs. As a result, we ran two further scenarios which were not reported by Openreach. These two scenarios test the two parameters discussed above (i.e. the glass ceiling and the number of redistribution weeks) in an effort to understand the extreme ends of the possible ranges these parameters can take, and to some extent view how the model might behave if they were not present at all.

The results of these scenarios are shown in Figure 3.12 below. These scenarios, as with the previous sensitivity results shown in Figure 3.9 and Figure 3.10, were carried out over a geographical area of four GMs (in this instance, London, South East, South West & South Midlands, and Wessex) for 2011/12 and for the following performance scenario: 80% repair performance, 80% provision performance.

To test the extreme limits of the glass ceiling the relevant parameter was raised from 90% to 99%.[19] This resulted in the resource delta from the baseline increasing from 6.18% to 10.69%, a significant increase.

To test the extreme limits of the number of redistribution weeks the parameter was raised from 10 to 40 weeks. This was designed to approximate the behaviour which would be observed if the redistributed jobs were allocated smoothly across all (or at least the majority of[20]) weeks, rather than just across a set of the weeks with the lowest performance. This parameter choice resulted in the resource delta from the baseline increasing from 6.18% to 12.37%, again a significant increase.

	90% glass ceiling, 10-week redistribution	99% glass ceiling, 10-week redistribution	90% glass ceiling, 40-week redistribution
Resource requirement	✂	✂	✂
Delta from baseline	6.18%	10.69%	12.37%

Figure 3. : Results for extreme glass ceiling and redistribution parameter values [Source: Analysys Mason, based on Openreach model, 2013]

These results are of course the very extremes of the possible ranges that these parameters can take, and as such potentially stress the model to unrealistic levels of operation. However, they do go some way to showing how important this element of the methodology is.

In summary, the inclusion of a glass ceiling appears to significantly reduce the resource requirement delta for increasing performance above the baseline, and redistribution of jobs where the glass ceiling is exceeded across a small number of weeks seems to have the same effect.

However, that being said, we do agree with Openreach that there is likely to be a limit to the number of jobs that can be completed (as a result of, for example, the need for major network intervention for a small percentage of jobs). Small variations in the glass ceiling parameter do not have a huge effect on the results.

3.4.2 Granularity of job duration data

Within the model job durations are assigned to jobs depending on the GM, site type and month in which the job is logged. The job durations are the average of all the jobs historically completed within that time period and at that location.

However, averaging this set of data removes a large amount of granularity from what should be a distribution of job durations. It stands to reason that the mean of this distribution (as used in the model) may be significantly different from the largest extreme of the data set.

[19] We had some concerns about whether the model would behave correctly with the parameter set to 100% and therefore chose a value of 99%, as we believe that it adequately illustrates our point.

[20] Again, we had some concerns that the model might behave in a way which we could not explain if this parameter were pushed too far. For example, if the parameter were set to, say, 51 weeks but there were, say, 5 weeks where performance had been pushed above the glass ceiling, some strange effects could have been observed.

It could be argued that the variability in job durations is already captured in the distribution of times to complete jobs. Also it would probably add significant complexity to the model to have a distribution of job durations rather than fixed values (at a given location in a given period) – especially if the job duration and overall duration (queuing time plus job time) were correlated. However, the exact impact of this simplification on the model results is not known, as the data needed to check this was not available to us.

4 Resource modelling for increased performance

This section discusses the results of the model when both provision and repair performance targets are changed from the baseline. The results are for both the 2011/12 and 2012/13 datasets and the provision of SLAs of both 13 and 12 days. All results are shown as percentage deltas from their respective baselines.

In all cases the results we obtained from running the model agree with those supplied to us by Openreach, to at least the level of accuracy provided (typically one decimal place).[21] However, as noted below, we have run some additional scenarios for which Openreach did not provide equivalent results for comparison.

This section is structured as follows:

- Section 4.1 discusses the adjustments made to input parameter values for each scenario
- Section 4.2 sets out the results from scenarios run using the 2011/12 historical dataset
- Section 4.3 sets out the results from scenarios run using the 2012/13 historical dataset
- Section 4.4 contains a brief summary of our conclusions from these results.

4.1 Adjustments to input parameter values

For most scenarios Openreach has provided results based on average job durations which reflect the underlying historical data. However, for scenarios with a targeted performance of 85% for either provision or repair, Openreach has instead chosen to modify the average job durations using the parameter discussed in Section 2.2.5. In particular, Openreach has provided a range of results for 85% performance scenarios with average job durations inflated by 5–10% for repair jobs and by 0–5% for provision jobs.

Openreach's justification for giving this parameter values in this range stems from analysis of real durations of jobs undertaken outside a technician's normal SOM (Senior Operation Manager) area; it was stated that jobs undertaken outside the SOM area by these technicians have a job duration of at least 16% higher than the average. Under situations of increased performance, such as those scenarios with a target of 85% repair or provision performance, Openreach indicated that around an additional 1% of total job volumes are executed by technicians from outside their normal SOM area, and that, according to Openreach, due to the "higher volatility of these jobs this proportion would rise to c. 6% of total jobs". This implies to us that, averaged across all jobs, the increase in job duration associated with greater target performance should in fact be around 1%. This can be seen simply by calculating the average increase in job duration needed across all jobs to simulate a

[21] In general the results are extremely close to those calculated by Openreach. For some of the more extreme variations from the base case our results differed from Openreach's at the second decimal place. However, this may only be as a result of running the discrete event simulation model on a different machine with a different installation of Simul8.

16% increase in job duration in 6% of jobs. As such a range of 5–10% for increases in job duration seems to us to be an unrealistically aggressive inflation of an important parameter.

For the 85% performance scenarios, we therefore show three results, reflecting different choices for the average job duration parameters. Figure 4.1 displays a key for those results with performances of 85% and their associated job duration inflations. The lowest value is that obtained with no increase in job durations, the middle value uses a 5% job duration increase on repair jobs only, and the highest value uses a 10% job duration increase on repair jobs and a 5% job duration increase on provision jobs.

Inflation of duration of provision job	Inflation of duration of repair job	Resource delta from baseline
0%	0%	Lower bound % delta
0%	5%	Middle value % delta
5%	10%	Upper bound % delta

Figure 4.1: Key for results with performances of 85% [Source: Analysys Mason, 2013]

4.2 Results for the 2011/12 dataset

Figure 4.2 displays the baseline resource requirement for the 2011/12 period, expressed in both peak required FTEs and kilo-man-hours (KMHs).

	Peak FTEs	Peak KMH
Model baseline, 2011/12	✂	✂

Figure 4. : Baseline resource requirements for 2011/12 [Source: Analysys Mason, based on Openreach model, 2013]

Figure 4.3 displays the monthly baseline resource requirement for the 2011/12 period by skill type as well as the maximum by day, maximum sum by skill type[22] and the maximum KMH. The full set of data for all 12 months has been provided by Openreach, but due to computational effort only a sub-set (the first two months, April and May) has been verified through independent model runs. For this two-month sub-set we found no discrepancies with the results provided by Openreach.

The monthly maximum figures are calculated by summing the daily requirements for each GM, then calculating the maximum of this sum over the month. The peak requirements, as shown in Figure 4.2, are calculated by summing the peak of each GM.

This process leads to the peak FTE requirement over the whole year being high when compared to the monthly maximum values, as shown in Figure 4.3, as the peak FTE requirement is the sum of

22 The maximum sum by skill type is found by taking the maximum day for each skill type and summing the values; this maximum therefore does not occur on a single day and is therefore higher than the maximum by day value.

the maximum values regardless of which days they may fall on. This analysis ties in with the explanations on the impact of idle resources set out in Sections 2.4 and 3.3.

This inflation of the peak FTE requirement goes some way towards explaining the seemingly low ratio between FTEs and KMHs shown in Figure 4.2. In particular, a quite large proportion of the ✂ FTEs shown for the baseline in Figure 4.2 are likely to be idle on many days, given the significantly lower numbers shown in the "Max by day" column of Figure 4.3. There may be other potential reasons for a low ratio of FTEs to KMHs, but the results currently show that ✂ FTEs perform ✂ KMHs of work, which would imply approximately ✂ hours of work per annum per FTE. For a 47 working-week year this would imply just ✂ hours per week, which seems low. As a result we expect that a large part of the reason for this low ratio will be the way in which resource requirements have been driven by the peak requirements for each skill type and result in large numbers of apparently idle resources.

Figure 4.3: Baseline resource requirement by skill type for 2011/12 [Source: Analysys Mason, based on Openreach model, 2013]

2011/12	Con-tractor	OMI	BB OMI	CAL	BB CAL	CSE	Max by day	Max sum by skill	Max KMH
	(FTEs)								(KMHs)
April	✂	✂	✂	✂	✂	✂	✂	✂	✂
May	✂	✂	✂	✂	✂	✂	✂	✂	✂
June	✂	✂	✂	✂	✂	✂	✂	✂	✂
July	✂	✂	✂	✂	✂	✂	✂	✂	✂
August	✂	✂	✂	✂	✂	✂	✂	✂	✂
September	✂	✂	✂	✂	✂	✂	✂	✂	✂
October	✂	✂	✂	✂	✂	✂	✂	✂	✂
November	✂	✂	✂	✂	✂	✂	✂	✂	✂
December	✂	✂	✂	✂	✂	✂	✂	✂	✂
January	✂	✂	✂	✂	✂	✂	✂	✂	✂
February	✂	✂	✂	✂	✂	✂	✂	✂	✂
March	✂	✂	✂	✂	✂	✂	✂	✂	✂

Openreach has used the model to calculate the resource deltas required to achieve increased performance levels against the 2011/12 baseline (79% repair, c. 65% provision for a 13-day provision SLA). It should be noted that these results rely on a notional 13-day provision SLA, as there was no explicit provision SLA during this period. As such we understand that it is Openreach's view that the 2011/12 data is not as representative as the 2012/13 data.

Figure 4.4 contains the resource deltas for performance increases with a 13-day provision SLA. As described above, we provide a range of results for the 85% performance scenario, including a lower end of the range with no alteration to the average job durations (a scenario for which results were not provided by Openreach).

SLA of 13 days		Provision		
		75%	80%	85%
	75%	1.3%	3.7%*	6.0% 6.3% 10.1%*
Repairs	80%	1.6%	3.3%	6.0% 6.0% 9.3%
	85%	3.3% 6.7% 10.0%	5.1% 9.0% 12.3%	7.1% 11.0% 17.5%

Figure 4. : Resource deltas based on 2011/12 data and 13-day provision SLA [Source: Analysys Mason, based on Openreach model, 2013]

Note: The results marked with asterisks are considered by Openreach to be affected by outlying data points; Openreach therefore believes that these are likely to be over-estimates of the relevant results.

The scenarios with performances of 85% display three numbers; a key to these results is shown in Figure 4.1.

Figure 4.5 contains the resource deltas for performance increases with a 12-day provision SLA.

SLA of 12 days		Provision		
		75%	80%	85%
	75%	2.2%	3.7%	5.5% 5.7% 9.6%
Repairs	80%	2.6%	3.9%	6.3% 6.3% 9.7%
	85%	5.0% 8.4% 11.8%	6.5% 10.3% 13.9%	8.1% 12.4% 19.0%

Figure 4. : Resource deltas based on 2011/12 data and 12-day provision SLA [Source: Analysys Mason, based on Openreach model, 2013]

Note: The scenarios with performances of 85% display three numbers; a key to these results is shown in Figure 4.1.

The percentage figures given in this section are the results of simulations which exclude all jobs classified as MBORC (matters beyond our reasonable control) from the performance calculation. If MBORC-affected jobs were to be included in the performance calculation this would increase the displayed resource requirement deltas by a very small amount.

In general, as might be expected, the resource delta increases as the performance targets are increased. Typically, as the absolute performance targets get higher, the resource requirement delta for a one-percentage-point performance increase will be greater, as one would expect.

Also, as the provisioning SLA is reduced to 12 days, there is an increased amount of resource required to hit a given performance target.

It is worth noting that for the 85% performance scenarios the range of results obtained, depending on the choice of job duration time inflation parameter, is considerable. As discussed above, although we recognise that some level of job duration increase might be expected as performance levels increase, the 5–10% range suggested by Openreach for repair and the upper end of the 0–5% range for provision seems aggressive. Therefore we would expect that results towards the bottom of the ranges we show in the tables above for the 85% performance scenarios are likely to be most realistic.

4.3 Results for the 2012/13 dataset

To obtain results using the 2012/13 data, Openreach decided that changes needed to be made to the empirical data in order for the results to be fully representative. These changes are described in this section, together with the results.

4.3.1 Changes made to 2012/13 provision data

Due to the nature of the observed performance distribution for provision jobs in the 2012/13 data, Openreach decided to adjust the dataset. The 2012/13 provision performance was, for a few weeks in the year, very poor. This led to a backlog of provision jobs within the system, which in turn caused an unrepresentative spike in resource requirement later in the year. To avoid this being modelled and resulting in an unrealistically high resource requirement, the provision performance distribution was (completely) flattened; this reduces job backlog in the latter part of the year by ensuring that performance stays constant and does not drop as low as was historically seen.

We note that the adjustment to performance data for 2012/13 will have the effect of pushing up resource requirement deltas when trying to improve performance, relative to a normal 'peaky' performance distribution. This is because performance in the lowest-performing weeks is now higher and as such it is relatively more expensive to increase performance, because improving performance in higher-performing weeks requires additional resources.

As Figure 4.6 illustrates, provision performance is now assumed to be constant throughout the scenario period, rather than exhibiting the large variation that the empirical data showed.

Figure 4.6: Distribution of provision performance across the scenario period [Source: Analysys Mason, based on Openreach model, 2013]

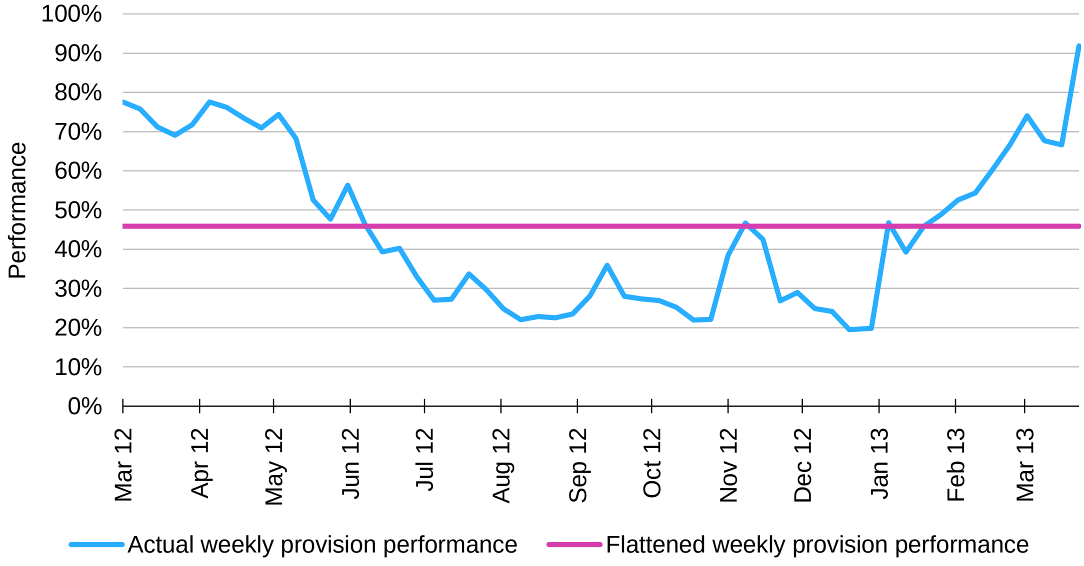

4.3.2 Results

Figure 4.7 displays the baseline resource requirement for the 2012/13 period as both peak required FTEs and KMHs.

	Peak FTE	Peak KMH
Model baseline 2011/12	✂	✂

Figure 4. : Baseline resource requirements for 2012/13 [Source: Analysys Mason, based on Openreach model, 2013]

Figure 4.8 displays the monthly baseline resource requirement for the 2012/13 period by skill type as well as the maximum by day, maximum sum by skill type and the maximum KMH. Again, the full set of data for all 12 months has been provided by Openreach, but due to computational effort only a sub-set (the first two months, April and May) has been verified. Again, for this two-month sub-set we found no discrepancies with the results provided by Openreach.

Similar to 2011/12, we observe that the ratio of FTEs to KMHs appears low. If anything, the differences between the ✂ FTEs and the values in the "Max by day" column of Figure 4.8 are even greater than for 2011/12, implying an even higher level of idle resource.

Figure 4.8: Baseline resource requirement by skill type for 2012/13 [Source: Analysys Mason, based on Openreach model, 2013]

2012/13	Con-tractor	OMI	BB OMI	CAL	BB CAL	CSE	Max by day	Max sum by skill	Max KMH
	(FTEs)								(KMHs)
April	✂	✂	✂	✂	✂	✂	✂	✂	✂
May	✂	✂	✂	✂	✂	✂	✂	✂	✂
June	✂	✂	✂	✂	✂	✂	✂	✂	✂
July	✂	✂	✂	✂	✂	✂	✂	✂	✂
August	✂	✂	✂	✂	✂	✂	✂	✂	✂
September	✂	✂	✂	✂	✂	✂	✂	✂	✂
October	✂	✂	✂	✂	✂	✂	✂	✂	✂
November	✂	✂	✂	✂	✂	✂	✂	✂	✂
December	✂	✂	✂	✂	✂	✂	✂	✂	✂
April	✂	✂	✂	✂	✂	✂	✂	✂	✂
May	✂	✂	✂	✂	✂	✂	✂	✂	✂
June	✂	✂	✂	✂	✂	✂	✂	✂	✂

Resource deltas to achieve increased performance levels against the 2012/13 baseline performances (63% repair, 42% provision for a 13-day provision SLA) are displayed below.

Figure 4.9 contains the resource deltas for performance increases with a 13-day provision SLA.

POS 13 days		Provision		
		75%	80%	85%
Repairs	75%	14.1%	17.2%	- 20.5% 24.1%
	80%	14.9%	18.1%	20.7% 20.8% 26.3%
	85%	- 20.4% 24.5%	21.1% 22.9% 27.1%	23.3% 24.9% 29.5%

Figure 4. : Resource deltas based on 2012/13 data and 13-day provision [Source: Analysys Mason, based on Openreach model, 2013]

Note: The scenarios with performances of 85% display three numbers; a key to these results is shown in Figure 4.1.

Figure 4.10 contains the resource deltas for performance increases with a 12-day provision SLA.

POS 12 days		Provision		
		75%	80%	85%
Repairs	75%	16.3%	19.3%	
	80%	16.8%	20.0%	20.1% 22.0% 24.6%
	85%		22.0% 24.3% 26.6%	24.9% 26.7% 31.5%

Figure 4. : Resource deltas based on 2012/13 data and 12-day provision[23] [Source: Analysys Mason, based on Openreach model, 2013]

Note: The scenarios with performances of 85% display three numbers; a key to these results is shown in Figure 4.1.

The percentage figures given in this section are the results of simulations excluding all MBORC-affected jobs from the performance calculation. If MBORC-affected jobs were to be included in the performance calculation it would have the effect of increasing the displayed deltas by a very small amount.

As can be seen from the results, the same patterns are generally exhibited as for the 2011/12 data when performance targets are increased or the SLA is decreased. However, in general the percentage deltas to the 2012/13 baseline are much higher than for the 2011/12 equivalent scenarios. This reflects the fact that the baseline performance in 2012/13 was much lower than in 2011/12, which means that much bigger differentials from the baseline are being sought.

4.4 Summary of conclusions for scenarios of increased performance

From the results discussed in this section there are three particular points of note:

- The difference between results from the 2011/12 dataset and those from the 2012/13 dataset are significant, though which is most representative we cannot say. It should be noted in this context that the adjustments made to the 2012/13 provision data will have the impact of increasing resource requirements.
- The job duration inflations lack clear justification; our understanding of the explanation provided by Openreach suggests that a much lower inflation parameter, if any, would be most appropriate. This would have the effect of lowering the resource deltas relative to those provided for the 85% scenarios by Openreach.
- The choice of redistribution parameter has a significant impact on the final results, and we have seen no clear justification as to which parameter is most appropriate.

23 There are some scenarios (85% provision, 75% repair performance and 75% provision, 85% repair performance) that were not run; this decision was made by Ofcom in an effort to minimise the number of scenario results required from Openreach and so reduce computation time.

5 Resource modelling for change in care level mix

This section considers the changes made to the model by Openreach to enable it to run scenarios with alterations in the care-level mix, as well as the results from those scenarios.

All results displayed in Sections 5.2 and 5.3 are resource requirement deltas relative to the respective baseline values for 2011/12 and 2012/13. In contrast to the previous section there is now a unique baseline for each performance target. The results in this section therefore show the impact of a change in care-level mix relative to this baseline of the same performance target; that is, for a given performance there are results for varying care-level mixes that are shown relative to the baseline *mix* (as opposed to baseline *performance*, which was measured against in the previous section).

Various scenarios are tested with different proportions of the total number of care level 1 (CL1) and care level 2 (CL2) jobs. These scenarios, defined by the proportion of CL1 jobs, are as follows:

- 100%, one extreme of the scale
- 60%, the 2011/12 basecase mix of jobs (only calculated for 2011/12 scenarios)
- 57%, the 2012/13 basecase mix of jobs (only calculated for 2012/13 scenarios)
- 55%, the 2013/14 forecast for mix of jobs
- 48%, the 2016/17 forecast for mix of jobs derived from Ofcom WLR and MPF product volume forecast
- 35%, an intermediate point
- 0%, the other extreme of the scale.

All Openreach results in this section are calculated using the 'Top 10 days redistributed' method of resource redistribution, as described in Section 2.4 and critiqued in Section 3.3, as opposed to the previous approach of consistently using the 'Maximum day' method. No satisfactory explanation for this change has been provided.

It should also be noted that the results displayed for scenarios of 85% performance do not include any adjustments to average job durations, unlike the previous scenarios which Openreach have considered.

The remainder of this section is organised as follows:

- Section 5.1 discusses the changes made to the model and its functionality to allow for scenarios involving variations in care-level mix to be modelled, together with our views on the appropriateness of these changes.
- Section 5.2 presents the results of scenarios using the 2011/12 historical dataset
- Section 5.3 presents the results of scenarios using the 2012/13 historical dataset
- Section 5.4 presents and discusses the results of 2012/13 scenarios using alternative resource redistribution parameters and methodologies

- Section 5.5 contains a brief summary of our conclusions from the results.

5.1 Model adjustments to enable changes in care-level mix

In order to calculate resource requirements for scenarios with different mixes of jobs across the care levels, some adjustments were required to the model. These changes, made by Openreach, are as follows:

- Changes to the 'routeings' parameters, as discussed in Section 2.2.2
- Changes to the 'Gamma parameters', as discussed in Section 2.2.4
- Exclusion of provision jobs from the model
- Exclusion of Saturday working for CL2 jobs.

These changes, the justifications provided, the validity of these justifications, and their possible effects on the results of the model are discussed in the following sections.

Changes to the 'routeings' parameters

The routeings input parameters contain the split of jobs across care levels. By manipulating this data the split of jobs across care levels can be targeted to a specified value between the two extremes: all CL1 and CL2 jobs completed at the CL1 SLA; or all CL1 and CL2 jobs completed at the CL2 SLA.

Openreach has indicated that this change was made to allow for this split between the two care levels to be parameterised within the model.

Changes to the 'Gamma parameters'

Changes were made to the gamma parameters used to create 'time to complete' distributions. The gamma parameters were changed to have the effect of aligning the peak of the CL1 and CL2 time to complete distributions. Figure 5.1 illustrates the changes made to the gamma distributions in order to align the peaks of the distributions. The basecase lines show the original gamma distributions used for CL1 and CL2 jobs (modelled as separate distributions) for an 80% performance target; the 'mix' lines show the distributions of times to complete used in the care-level mix scenarios.

The justification given for this change was that it was done in order to achieve a meaningful comparison of the extreme cases and to take into account what might likely happen in the real world if these extremes were to be realised.

Whilst it is perhaps fair to say that some change in these distributions would be observed in the extreme cases in the real world it is much harder to justify such extreme changes for all scenarios with changes in care level mix; especially given how the position of the mode of the distribution was used as a constraint in the previous version of the model (as discussed in Section 3.2).

This change is therefore not fully justified as it significantly changes the modelling process, even for scenarios where only a small change in care level mix is being considered.

Figure 5.1: Comparison of previously used gamma distributions and those used in scenarios with care-level mix changes [Source: Analysys Mason, based on Openreach model, 2013]

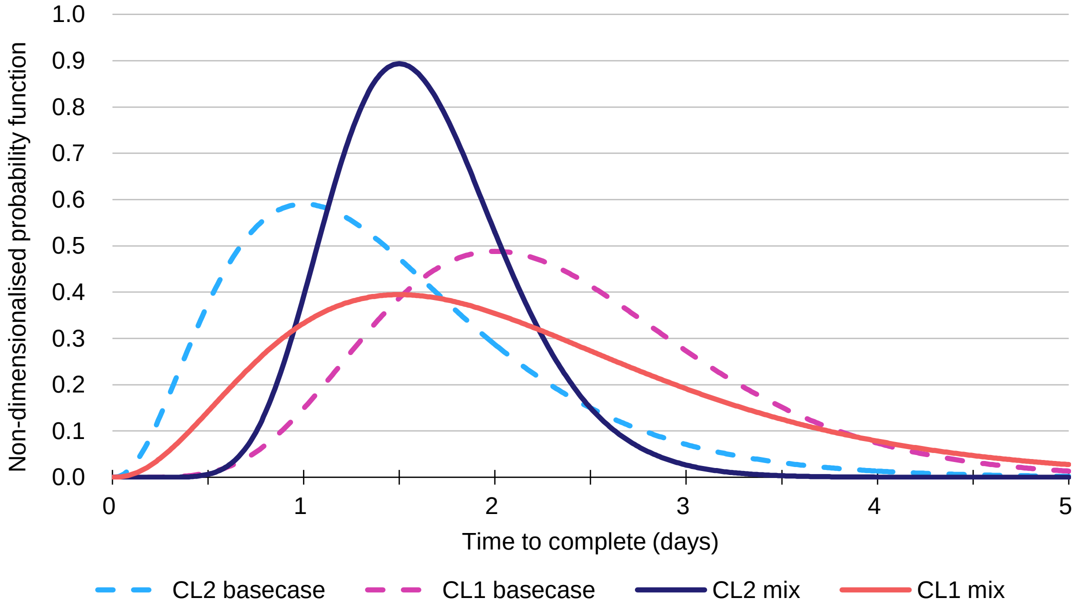

As it is a departure from the original methodology used to ascertain the resource requirement deltas for changing performance targets (as described in Section 4), the care-level mix scenario results must be considered relative to their own baseline requirements, as described above.

Exclusion of provision jobs from the model

The raw jobs data sets themselves were adjusted to exclude all provision jobs from the simulation.

The justification given for this was that these care level mix scenarios are considering the effects of changing the mix of repair jobs and so removing provision jobs would aid in comparing scenarios without the possibility of them being affected by exogenous factors brought about by the inclusion of provision jobs.

This justification is valid to a point. As the care level mix changes, significant numbers of jobs are completed on different days (days that may or may not already have high numbers of provision jobs being completed). This means that the way in which resource requirements change with changing care level mix may be heavily influenced by a few statistically significant days which are affected to a large extent by the historical distribution of provision job completions throughout the year (a factor based on the raw input data).

However, this is a departure from reality which ignores the potential economies of scope that arise from having both repair and provision jobs in the same system, and so must be considered when examining the results of these scenarios.

Exclusion of Saturday working for CL2 jobs

The final change made by Openreach is the exclusion of Saturday working for CL2 jobs. Previously, jobs assigned a CL2 SLA have had the ability to be completed on Saturdays, whereas jobs assigned a CL1 SLA have not. This ability has now been removed for CL2 jobs. As such, all jobs assigned to either CL1 or CL2 must be completed on weekdays. Therefore, jobs that would previously have been completed on a Saturday are distributed equi-proportionately across weekdays.

The justification provided by Openreach for this change is that in reality the number of workers employed to work the Saturday shift is much smaller than for a weekday shift; as such, in scenarios of increased CL2 jobs potentially more jobs would be assigned to a Saturday than would be achievable in the real world.

As with the exclusion of provision jobs this justification is valid to a point but makes a departure from reality which must be considered when viewing the results of these scenarios. In particular, the modification has the effect of increasing the peak of the busiest weekday, and by extension the peak resource requirement; though the effect this has on the resource deltas is unknown, as the resource increase will occur in both the relevant baseline and any performance scenario being tested.

We also note that if the resource requirement delta were to be used as a proxy for cost differences between different care-level mix scenarios then care would need to be taken. In particular, our view is that Saturday jobs would have to be included in order to conduct a proper analysis of costs, to account for the cost premium associated with Saturday working. A simple methodology for this would be to include Saturday jobs and to post-process the model results to enable some cost premium to be applied to these resources.

5.2 Results for the 2011/12 dataset

Again, in all cases the results we obtained from running the model agree with those provided to us by Openreach, to at least the level of accuracy provided (typically one decimal place)[24].

Figure 5.2 shows the results of varying the split of CL1 and CL2 jobs in the simulation for the 2011/12 dataset.

24 In general the results are extremely close to those calculated by Openreach. For some of the more extreme variations from the base case our results differed from Openreach's at the second decimal place. However, it is expected that this is only as a result of running a discrete event simulation model on a different machine with a different installation of Simul8.

2011/12	% CL1 jobs	Performance		
		75%	80%	85%
	100%	-10%	-11%	11.2%
Basecase	60%	0%	0%	0%
13/14 f/cast	55%	3.3%	1%	0.4%
16/17 f/cast	48%	3.8%	1%	1.5%
	35%	5.7%	7%	3.5%
	0%	13.1%	12%	9.1%

Figure 5. : Resource deltas based on 2011/12 data [Source: Analysys Mason, based on Openreach model, 2013]

5.3 Results for the 2012/13 dataset

Figure 5.3 shows the results of varying the split of CL1 and CL2 jobs in the simulation for the 2012/13 dataset.

2012/13	% CL1 jobs	Performance		
		75%	80%	85%
	100%	-4.7%	-5.9%	-9%
Basecase	57%	0%	0%	0%
13/14 f/cast	55%	0.4%	1.1%	1%
16/17 f/cast	48%	2.3%	3%	2%
	35%	5.6%	5.4%	5%
	0%	10.4%	13%	15%

Figure 5. : Resource deltas based on 2012/13 data [Source: Analysys Mason, based on Openreach model, 2013]

Note that the percentage values in each column are additive. For example, for the 85% performance scenario, there is a 9% resource requirement delta (9% of base case resource levels) to move from the case where 100% of jobs are CL1 to the base case with only 57% of jobs at CL1. To move to 55% CL1 jobs there would be a total resource delta of 10% (relative to base case levels) and to move to 35% CL1 the total resource delta would be 14%.

Figure 5.4 displays these results as a scatter graph with trend lines added to each series. Given that the relationships shown are linear, these trend lines give an indication of the increased volume of resources that CL2 jobs require, according to the Openreach model.

As can be seen, the gradient of these lines varies from 15% to 24% with the performance target. In this case, with the x-axis running from 0% to 100% the gradient is simply the difference between y-axis values at x=0% and x=100%. This represents the resource delta between job mixes of 100% CL2 jobs and 100% CL1 jobs.

As such, we understand that Openreach believes that a gradient of approximately 20% is a fair conclusion for the change in resource volume arising from an increase in the proportion of CL2 jobs.

Figure 5.4: Summary of care-level mix results [Source: Analysys Mason, based on Openreach model, 2013]

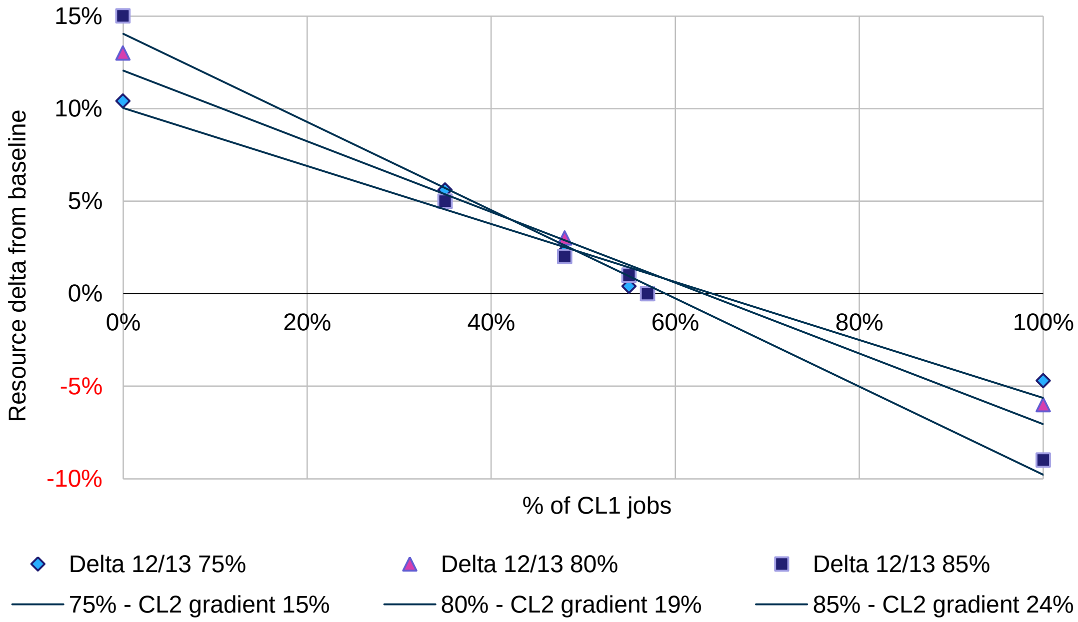

If these results are to be accepted, then the conclusion drawn is reasonable. However, these results rely on:

- the various model adjustments set out in Section 5.1 which we do not believe are fully justifiable; and
- the use of the Top 10 days redistributed approach for calculating the total number of resources required and hence the delta between the number of resources in the baseline and the alternative care-level mix scenarios.

As explained in Section 3.3, we do not consider the Top 10 days redistributed results (as derived in the Openreach model) to be particularly meaningful. In Section 5.4 below, we therefore set out the results of the equivalent scenarios based on the Maximum day redistributed approach and based on Analysys Mason's adaptation of the 'top N days' redistributed approach for various values of N.

We also reiterate our comment about there being better ways to include the effects of Saturday working for CL2.

5.4 Results for 2012/13 using alternative redistribution approaches

As previously mentioned, the results in Sections 5.2 and 5.3 were obtained using the Top 10 days redistributed approach. However, it is far from clear that this approach is most appropriate. First, it is a departure from the Maximum day redistributed approach used by Openreach for the generation of other results. And second, we do not consider the implementation of the approach in the Openreach model to be meaningful, as explained in Section 3.3.

In this section results from a particular set of scenarios are set out using different redistribution methods; in doing so the effects of this parameter in this context can be analysed.

Figure 5.5 displays the results of varying the split of CL1 and CL2 in the simulation for the 2012/13 dataset when the (repair) performance target is set to 80%. The figure displays the results obtained from using different redistribution parameters and methodologies:

- the Top 10 E&Y method (whose results were displayed in the previous section)
- the Top 10 AM method (which was discussed in Section 3.3 and is an adjusted version of the E&Y method)
- the Maximum day method (which in effect uses a redistribution parameter of 1).

2012/13	% CL1 jobs	Redistribution method		
		Max day	Top 10 E&Y	Top 10 AM
	100%	-6.38%	-5.9%	-5.98%
Basecase	57%	0%	0%	0%
13/14 f/cast	55%	0.22%	1.10%	0.41%
16/17 f/cast	48%	0.88%	3.00%	0.98%
	35%	2.92%	5.40%	4.56%
	0%	7.59%	13.00%	11.87%

Figure 5. : Comparison of resource requirement deltas under various redistribution approaches [Source: Analysys Mason, based on Openreach model, 2013]

As can be seen, the resource requirement delta for scenarios with a greater proportion of CL2 jobs is lower in both the Maximum day and Top 10 AM redistribution approaches than in the Top 10 E&Y approach reported by Openreach.

These data points are then shown in Figure 5.6 with associated trend lines. Given that these trends are linear they give an indication of the increased volume of resources that CL2 jobs require for the different redistribution approaches.

Figure 5.6: Summary of care-level mix results for various redistribution approaches in an 80% performance scenario [Source: Analysys Mason, based on Openreach model, 2013]

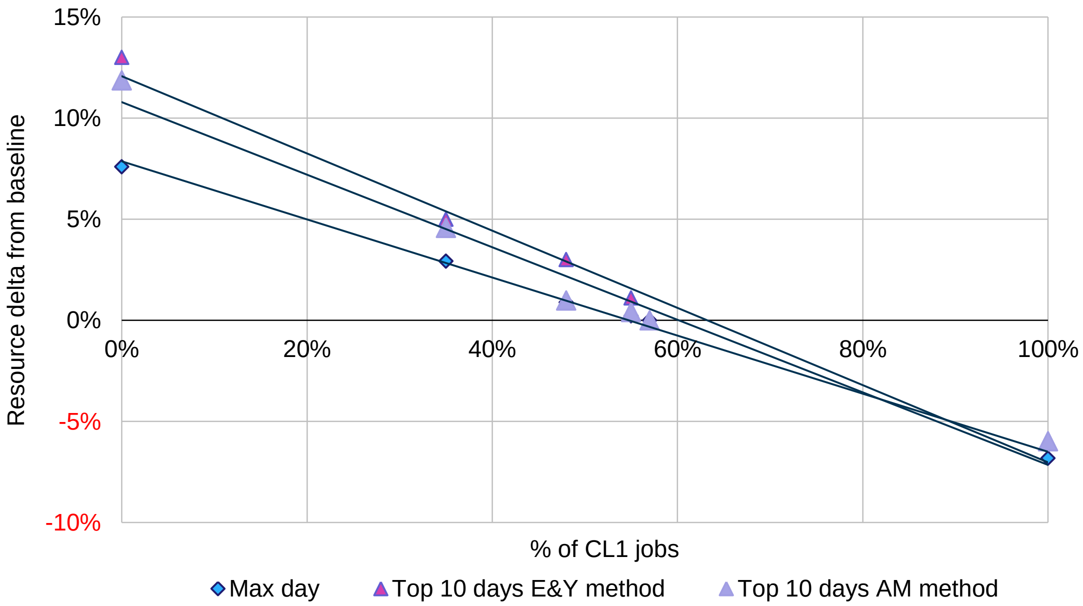

As stated in the previous section, the Top 10 E&Y data yields a gradient of 19% for all jobs at CL2 instead of CL1. The Top 10 AM method yields a similar gradient of 18.6%; as such there does not appear to be a great distinction in the two methodologies in this particular instance, even though the approach by which the result is arrived at for the Top 10 E&Y approach does not appear particularly meaningful. However, the Maximum day method (which was used consistently for the results shown in Section 4) yields a gradient of only 14%, significantly lower than the other two methods.

We therefore conclude that the choice of redistribution approach can make a significant difference to the results and that there is no obvious reason for the approach to be different in the care-level mix scenarios from that in all other scenarios reported on by Openreach.

5.5 Summary of conclusions for scenarios with changes in care-level mix

Four main points arise from the results discussed in this section:

- The adjustment to the way in which the gamma distributions are manipulated is a departure from the original methodology used to ascertain the resource requirement deltas for changing performance targets. This departure has an unknown and potentially significant effect on the results.
- The exclusion of provision jobs from the model, whilst isolating the effects caused by changes in care level mix, is a departure from reality and reduces any potential economies of scope that may be attainable across repair and provision jobs. This departure has an unknown effect on

the results and any effect would most likely be caused by a statistically significant few days which would be driven by the raw input data itself.

- The exclusion of Saturday jobs, which for CL2 is a departure from reality. Although we understand that the adjustment is made to aid in the comparison of FTE requirements between different scenarios it will have the effect of increasing the weekday peak and hence the resource requirements. Also, if resource requirement deltas were used as a proxy for cost difference between scenarios then care would be needed as this approach does not account for the increased cost of Saturday working.
- The results are heavily dependent on the relatively unexplained choice of the Top 10 redistribution approach.

6 Conclusions

Our overall impression of the model is that Openreach / Ernst & Young have built a reasonable initial model to try and address a very challenging issue. The model generally seems to be well built and is without significant errors in the coding and implementation. The modelling tool (a combination of Excel and Simul8) also seems appropriate to the task it is intended to perform, albeit arguably rather slow. Results are consistent and replicable, yielding identical results on the same machine and very similar results between different machines.

6.1 Conceptual issues

The approach taken to developing the model is not standard for a queuing model. Rather than taking resource levels as an input and calculating performance, the Openreach model uses performance to be achieved as an input and calculates the required resources. This not a problem *per se*, but it does seem to create situations where quite significant assumptions are required. The greatest of these is the manipulation of the gamma distribution used to represent the waiting time in the queue.

Again, the use of a gamma distribution is not a problem *per se*, and in the baseline cases it seems to reflect the empirical data well. However, the way in which the gamma distribution is manipulated when performance targets are modified, by fixing the mode of the distribution, lacks justification and could potentially lead to significant errors.

Fixing the mode of the distribution is not necessarily *wrong;* another constraint is needed in order to select a unique alpha–beta pair of parameters for the distribution. However, others could have been chosen with similar levels of justification and such choices could reasonably either decrease, or increase, the peak of the distribution and therefore the resource requirements which are calculated. Importantly it is unclear to us what magnitude of effect this choice of fixing the mode of the gamma distribution has on the calculated resource requirements, particularly for provision jobs, which is an important point.

There are also several other important issues with the model which we have identified:

- The use of a 'glass ceiling' to ensure that there is a cap on the performance which can be achieved in any week, and the associated parameter of the number of weeks over which jobs are redistributed, can lead to notable decreases in the resource requirement. If these parameters were not present then the resource requirement deltas for increased performance would be higher. However, we do consider their inclusion to be sensible in principle, and the chosen parameter values appear to be within reasonable bounds.
- Job durations within the model are fixed, for a given month, location and site type. It would likely have been more accurate to consider a distribution of job durations, although we

acknowledge that this would have added significant complexity. The impact of this choice is not measurable in the model.

- There is a potentially significant issue relating to the approach taken to redistribution of resources with higher skill levels to lower skill level jobs within the model. There are two important components to this issue.
  - First, the way in which resources are allocated and subsequently redistributed leads to an amount of idle resource being available on the majority of days. In reality this resource could be used to carry out jobs that are in the system and, as such, increase the performance achieved on those days. As a whole this would have the effect of increasing the overall performance of a given scenario above the level which is specified in the scenario and therefore over-stating the resources required to actually achieve a given target performance. Though it is not possible to calculate the exact impact of this effect, the magnitude is potentially significant.
  - Second, we consider the Top N redistribution approach, as implemented in the model, not to be a useful measure. The approach will drive resource requirements from the day with the maximum requirement for each skill level (as with the Maximum day redistribution approach), but will only allow a subset of resources to be redistributed to lower skilled jobs when there are not enough higher-skill jobs to occupy them on days other than the maximum. This results in an increase in resource requirement relative to the Maximum day redistribution approach and results in an even higher level of idle resource (or equivalently a higher level of over-achievement against the specified target). In our view a better approach would be to drive the requirement at each skill level from the Top N days average requirement rather than from the peak requirement, reflecting the fact that the resource pool may not be sized to handle the very worst day of the year and therefore allowing for a lower number of total resources to be calculated.

6.2 Scenario results

The results reported by Openreach seem to accurately reflect the results that we have been able to produce using the Openreach model.

However, the values for certain input parameters used to generate these results are vitally important and in some cases are open to question.

For the scenarios that focus on revised performance targets there are three points of particular note:

- First, there is a big difference between results for 2011/12 and those for 2012/13. We are not really in a position to say which year is more representative of future years, although we note that Openreach advocates 2012/13. However, we can say that the adjustment to performance data for 2012/13[25] would have the effect of pushing up resource requirement deltas when trying to improve performance, relative to a normal performance distribution with higher 'peaks' and, crucially, deeper troughs.

- Second, the job duration inflation used by Openreach (0–5% for provision jobs and 5–10% for repair jobs) for the 85% performance target scenarios lacks clear justification. The justification we have been provided with suggests that a much lower job duration inflation, if any, should be used. This has a major impact on results. For example, for an 85% performance scenario (both repair and provision) Openreach suggests a range of results of between 24.9% and 29.5% resource requirement delta in 2012/13 with a 13-day provisioning SLA. If job duration inflation was not used then a delta of only 23.3% would be produced, which we estimate would rise to up to around 23.6% if a more realistic job duration inflation of up to 1% were used. This parameter therefore has a significant effect on results, and the top end of the range reported by Openreach seems particularly aggressive.

- Third, the choice of redistribution approach is very important and has a significant impact on results. As shown in the example of Figure 3.8 (85% repair and provision targets with a 12 day provision SLA in 2011/12), the resource delta relative to the baseline is 8.11% in the Maximum day redistributed case. However, for a Top 10 day AM redistribution approach this delta falls to 4.49% and for Top 25 day redistribution it falls to 0.55%. It should be noted that both of these Top N day values use the AM approach described in Section 3.3 – the approach implemented in the model would suggest values of 5.10% and 3.78% respectively, although as explained above we do not consider these calculations to be meaningful. In any event, there is a very large change in the results and the associated conclusion by varying the choice of redistribution approach and we have not seen any clear justification for which value is most appropriate to use for N. In particular, taking the maximum redistributed approach (i.e. N=1) may result in an overstatement of resource deltas for a change in SLA or target performance.

For scenarios looking at a revised mix of jobs by care level we have identified some important uncertainties relating to the overall implementation:

- The gamma distribution parameters used to create job time to complete distributions for a changed mix of care levels are manipulated. These parameters were adjusted in order to align the peak of the CL1 and CL2 time to complete distributions, which are usually kept separate. We understand that this change was intended to allow for a better comparison between the two extreme scenarios of all jobs being designated CL1 versus all jobs being designated CL2. However, this is a departure from the original methodology used to ascertain the resource requirement deltas for changing performance targets, and has an unknown and potentially significant effect on the results.

---

25 The 2012/13 provision performance was, for a few weeks in the year, very poor. This led to a backlog of provision jobs within the system, which in turn caused an unrepresentative spike in resource requirement later in the year. To avoid this being modelled and resulting in an unrealistically high resource requirement, the provision performance distribution was (completely) flattened.

- The exclusion of provision jobs from the model, whilst isolating the effects caused by changes in care level mix, is a departure from reality and reduces any potential economies of scope that may be attainable across repair and provision jobs. This departure has an unknown effect on the results and any effect would most likely be caused by a statistically significant few days which would be driven by the raw input data itself.

- Saturday jobs are excluded, which for CL2 jobs is a departure from reality. We understand that the modelling change is made to enable what Openreach consider to be a more appropriate comparison in FTE requirements between scenarios with varying proportions of jobs. However, it will have the effect of increasing the weekday peak and hence the resource requirements. Also, if the resource requirement delta were to be used as a proxy for cost differences between different care-level mix scenarios then in our view it would be preferable for Saturday jobs to be included (to account for the increased cost that Saturday working entails) and for the cost metric to reflect that additional cost.

In addition, the care-level results are heavily dependent on some relatively unexplained parameter choices, in particular the choice of the Top 10 redistribution approach. For example, for 2012/13 the model calculates a resource requirement delta of 19% for all jobs to be classified at CL2 instead of CL1 using the Top 10 redistributed approach as implemented. Using what we consider to be a more meaningful approach to the Top 10 redistribution method (i.e. the AM approach) yields a similar gradient of 18.6%. However, the Maximum day redistribution approach, which was used consistently by Openreach for reporting all results other than for changing the care-level mix yields a gradient of only 14%, significantly lower than the Top 10 approach.

The choice of redistribution approach can therefore make a significant difference to the results, and although we do not particularly recommend one (correctly implemented) redistribution approach over another, we do not see any clear reason for the approach to be different in the care-level mix scenarios from all the other scenarios reported on by Openreach.

6.3 Summary

Overall, we think a useful and productive effort has been made to significantly improve the understanding of the relationship between resource requirements and QoS. Prior to this model, these parameters were simply not addressable in a systematic way. The model does not seem to be biased towards results favourable to Openreach (e.g. it includes the glass ceiling and 10-week job redistribution parameters, both of which serve to reduce resource requirement deltas compared to possible alternative choices).

Nonetheless, there are some material issues with the model. In particular, these relate to the (inconsistent) choice of input parameter values between scenarios and the very significant impact of changing certain parameters, specifically the job durations and the redistribution approach. More fundamentally, the way in which the gamma distribution is manipulated to try to test scenarios with different performance, SLAs or job mixes from those observed historically is not based on any testable rationale, and the reporting of the peak resource implies a level of idle resource which may lead to a systematic overestimate of resource requirements. Both of these reduce our level of confidence in the model outputs.