Openreach Discrete Event Simulation Model

Methodology document

November 2013

Ernst & Young LLP

Building a better working world

DISCLAIMER NOTICE

This report ('Report') was prepared by Ernst & Young LLP for BT Group plc, under BT Group plc's instruction.

Ernst & Young LLP does not accept or assume any responsibility in respect of the Report to any readers of the Report, other than BT Group plc ('Third Parties', which, for the avoidance of doubt, includes Ofcom). To the fullest extent permitted by law, Ernst & Young LLP will accept no liability in respect of the Report to any Third Parties. Should any Third Parties choose to rely on the Report, then they do so at their own risk.

Ernst & Young LLP has not been instructed by its client, BT Group plc, to respond to queries or requests for information from any Third Party and Ernst & Young LLP shall not respond to such queries or requests for information. Further Ernst & Young LLP is not instructed by BT Group plc to update the Report for subsequent events or additional work (if any) performed by Ernst & Young LLP. Accordingly, without prejudice to the generality of the foregoing, Ernst & Young LLP accepts no responsibility to any Third Party to update the Report for such matters.

Contents

1. Introduction

1.1 Background

Ofcom is currently conducting the Fixed Access Market Review (FAMR)[1]. As part of the FAMR, it is considering a number of aspects relating to Openreach's service quality and cost base, including Service Level Agreement (SLA) and Service Level Guarantee (SLG) arrangements[2], the potential set of minimum service standards for Repair and Provision jobs, and the different levels of resource to provide differential Repair performance (referred to as 'Care Levels' by Openreach).

In this context, Openreach is seeking to understand the costs of providing service at different levels of performance, and to use this analysis as an input to: (i) Ofcom's latest charge control modelling for Local Loop Unbundling (LLU) and Wholesale Line Rental (WLR) services[3], which is expected to come into force from 1 April 2014; and (ii) its own internal business planning and field force management.

1.2 Openreach Discrete Event Simulation Model

Openreach sought to develop a model that would estimate the level and mix of resource required for Openreach to achieve: (i) existing 'base line' service levels; and (ii) alternative (i.e., higher) levels of minimum average service standard for both LLU and WLR.

The result is the Openreach Discrete Event Simulation Model (referred to as 'the Model') which models relationships between service level agreements, service level performance and the required number of Openreach engineers to deliver determined service levels, using Openreach source data from the two most recent financial years, i.e., 2011/12 and 2012/13.

This document summarises the Model scope, input data and Model calculation steps, and provides a summary of the key modelling assumptions. The remainder of the document is structured as follows:

- Section 2: provides the Model overview, which covers the Model scope, modelling approach, and key inputs;
- Section 3: sets out the Model calculation steps, explaining how the Model calculates the required number of engineers; and
- Section 4: provides a summary of some identified modelling issues and explains the approach taken to addressing these.

[1] http://stakeholders.ofcom.org.uk/consultations/fixed-access-market-reviews/
[2] SLAs cover the maximum time available to Openreach to complete a Repair or Provision activity. SLGs represent the charges that Openreach is liable to pay to other Communications Providers in the event it fails to meet an SLA for a particular Repair or Provision job.
[3] http://stakeholders.ofcom.org.uk/binaries/consultations/llu-wlr-cc-13/summary/LLU_WLR_CC_2014.pdf

2. Model overview

2.1 Model scope

The Model considers Repair and Provision jobs relating to Openreach products[4] undertaken by field engineers in Openreach's Service Delivery department for the financial years 2011/12 and 2012/13, and seeks to model the number of engineers that would have been required to complete those jobs under different SLA targets (i.e., the target time for job completion associated with the SLA), and for different levels of performance against those SLAs (i.e., the percentage of jobs completed within the SLA target). The Model assumes that Openreach had sufficient ability and incentive to optimise performance given the resource it had available in 2011/12 and 2012/13. It therefore does not seek to 'optimise' performance. The objective of the modelling has informed the approach taken, whereby the input to the Model is a distribution of performance against an SLA target that is then 'flexed' for different SLA target scenarios to derive the required number of engineers. An alternative would have been to use an 'allocation rules' approach, whereby the Model calculates the performance that a fixed number of engineers can complete given a set of resourcing rules as to how incoming jobs are prioritised.

The analysis was conducted for individual 'General Manager' (GM) areas, which are organisational areas defined by Openreach, each covering a specific geographic area in the United Kingdom. The Service Delivery field force is separated into ten GM areas[5], but for the purposes of our analysis we have excluded Northern Ireland to be consistent with the scope of Ofcom's charge control modelling. The Model calculates resource requirements and performance within each GM area, whilst assuming free movement of resource within each of the GM areas. This is a simplification as, in practice, the area a Service Delivery engineer normally works in is limited to a much smaller area within reach of their normal headquarters location[6]. For example, within a GM area (e.g., Scotland), an individual engineer may have a normal working area of a single town or surrounding area, or just a part of one of the major cities. Thus the Model allows for greater freedom of movement amongst engineers than is actually the case, with the impact being, other things equal, a possible underestimation of required resource. This simplification, which reduces model computation time significantly, is used consistently in base line and scenario runs, and so is not expected to materially distort the modelled results.

Openreach's engineering field force can be trained to different levels, and hence not all engineers can complete all Repair and Provision tasks. The Model takes into account the type of skill required to complete each Repair and Provision job, and reflects this constraint on engineering requirements when calculating the required number of engineers by skill type for each day. Skill types are assumed to follow a hierarchy in that an engineer can undertake jobs of their own skill level as well as lower skill level jobs, and the Model allows for available engineers to undertake jobs of a lower skill level.

The focus of the exercise has been on modelling the impact of changing performances on the copper product portfolio, but all jobs performed by the Service Delivery workforce, including those relating to fibre access (NGA), ISDN2 and ISDN30 are included in the data set used in the modelling, in order to be consistent with the breadth of Ofcom's charge control modelling.

The Model considers the 'Care Level' assigned to each product, which specifies a limit as to the number of days Openreach has to complete a job as follows:

[4] The products included in the modelling are Metallic Path Facility (MPF), Shared Metallic Path Facility (SMPF), Wholesale Line Rental (WLR), ISDN2, ISDN30, Generic Ethernet Access (GEA) and Broadband boost (BBB).

[5] These are South Wales and South Midlands, North West, Wessex, North East, North Wales and North Midlands, Scotland, East Anglia, South East, London and Northern Ireland.

[6] Each GM area comprises a number of 'Senior Operations Manager' (SOM) patches. Reporting to the GMs, the 58 SOMs each typically have line management responsibility for the Service Delivery engineers who are headquartered within the SOM patch. There are a total of 430 Preferred Work Areas (PWA) across the UK, to which individual engineers are assigned. Each PWA typically covers a fraction of the SOM patch. An engineer will occasionally undertake work outside his/her normal PWA, and indeed occasionally would need to be assigned to jobs outside the home SOM patch, either on a job-by-job basis or for periods of one or more days.

- Care Level 1: end of next working day plus one;
- Care Level 2: end of next working day, irrespective of the time of day the fault is received;
- Care Level 3: end of next half working day; and
- Care Level 4: within 6 hours of the call being logged with Openreach.

The Model considers only Service Delivery resource, and does not model the impact of changed workload profiles on other supporting divisions such as Network Investment (for jobs requiring civil engineering or other complex work) or Service Management (for jobs requiring interaction with Openreach customers to progress). Further, the Model assumes that demand for increased performance levels on Repair and Provision activity can be met immediately by employing and deploying additional engineers. Therefore, the time and cost incurred to recruit, train and equip additional engineers, and the additional cost incurred during the period where newly recruited engineers increase productivity to the existing average level, are not reflected in the modelling.

2.2 Introduction to the Model

The Model uses an approach known as Discrete Event Simulation[7] (DES). A simulation is a computer model that mimics the operation of a real or proposed system, in this case being the day-to-day operation of engineers in Openreach.

The Model simulates the flow of jobs from customers logging calls regarding Repair and Provision jobs through to job completion by modelling the outcome as a result of changes to the distribution of completion times (i.e., an increase or decrease in the proportion of jobs completed within the SLA target), or as a change to the SLA target (i.e., the time available for Openreach engineers to complete a Repair or Provision job). This enables the Model to reflect the impact that changes in response times and performance levels have on the required number of engineers of different skill levels on a daily basis through the period of the simulation. The Model is also capable of assessing the impact of variations in forecast demand on the required number of engineers.

2.3 Model inputs

The Model draws from a range of data sources from Openreach's operational systems. This data includes information on the jobs completed by Service Delivery, the task times associated with each job, job prioritisation rules, shift patterns, engineer skill types, and skill matrices (which define which jobs different skilled engineers can complete).

2.3.1 Service Delivery volumes

The information in respect of Service Delivery volumes is extracted from Openreach's source operational systems. The data extract contains the following detail for every Repair and Provision job logged for the financial years 2011/12 and 2012/13:

- Log date and time: the date and time that Openreach was notified of the fault or of a request for a new Provision;
- Completion date and time: when the Repair or Provision job was completed;
- Product: the product the job relates to, e.g., WLR, MPF;
- Care Level: the Care Level of the product to which the Repair job relates[8];

[7] Discrete event simulation (DES) models the operation of a system as a discrete sequence of events in time. Each event (such as an Openreach field activity) is assumed to occur at a particular instant in time and marks a change of state in the system. Between consecutive events, no change in the system is assumed to occur; thus the simulation can directly 'jump' in time from one event to the next.

[8] Openreach does not assign a Care Level to Provision jobs.

- Appointment date: the date on which an appointment (if required) was arranged; and
- GM area: the GM area to which the job is logged.

2.3.2 Task time information

Openreach collates information on the time it takes to complete each job, and in its 'Tracker' database[9] calculates an average task time for each job it undertakes, by dividing through total 'kilo man hours' (KMH) per job type (including travel time), by the volume of that job type[10]. This data is recorded by Openreach on a weekly basis by job for each GM area. The Model uses this average task time data as an input to derive the number of engineers required at any point of the day to complete active jobs.

2.3.3 Manpower information

The data used in the Model in respect of Openreach's workforce includes:

- A hierarchy of engineer skill levels, and the jobs which engineers of each skill level can complete, which is used to determine the skill level of the engineer required for each specific job;
- Job allocation rules, which define how resources are allocated by Openreach during busy periods;
- The standard days and hours of work of Openreach engineers across a week;
- Skill types of the engineering force, comprising:
    - OMI (One Man Installer): the least skilled engineers, who are trained solely for copper Provision jobs;
    - Broadband OMI: engineers trained for fibre (as well as copper) Provision;
    - CAL (Customer Apparatus and Line): engineers trained for both Provision jobs and simple Repairs;
    - Broadband CAL: engineers trained to complete fibre (and copper) Repairs and Broadband boost jobs;
    - CSE (Customer Service Engineer): highest skilled engineers, trained to the level of a Broadband CAL and, in addition, capable of completing certain complex Repair jobs.

[9] Openreach's Tracker database summarises by GM area the total kilo man hours logged against Repair and Provision jobs on a weekly basis.
[10] Certain jobs may require a number of separate visits, with the hours for each visit being booked against that job.

3. Model calculation steps

To understand how resource requirements may change in line with different SLA targets, and different performances against the SLA targets, the Model is first run using an actual (i.e., historic) distribution of performance against SLA targets (i.e., the percentage of jobs completed within actual SLA targets) to define a 'base line' number of engineers. This distribution is then 'flexed' to run scenarios reflecting different performances and/or different SLA targets and the Model is re-run to calculate the required number of engineers for each scenario. The percentage difference ('delta') between the base line number of engineers and those required in the specific scenario is then calculated. The calculation steps in the Model are summarised in Figure 1 below.

Figure 1: Model calculation steps

Step 1
- **Service Delivery volumes:** jobs are logged in the system at a constant rate every half day, based on actual job log times from 2011/12 and 2012/13

Step 2
- **Care Level:** each job is randomly assigned a Care Level, based on monthly statistical distributions of Care Levels for that particular type of job

Step 3
- **SLA Target:** on the basis of the assigned job Care Level the relevant SLA target for job completion is calculated

Step 4
- **Job completion:** the job is assigned a job completion time using a statistical distribution of completion times for Care Level, job type, time of year, and GM area. The distribution is updated weekly based on Care Level and performance for that week and for that particular GM area

Step 5
- **Task times:** an average task time is assigned to the job based on product type, month, and GM area and the start time is calculated using the completion time minus the average task time

Step 6
- **Jobs queuing:** the job is then placed in a queue until assignation of a start time

Step 7
- **Daily engineer demand:** the job is released from the queue, and either an available engineer (of the same or higher skill level) is assigned that job, or an additional engineer is added to the system. At the end of the simulated day, the total number of engineers required by skill level is recorded and then reset to zero for the next day

Step 8
- **Output:** a required number of engineers for the year is derived

Step 9
- **Flexing performance:** by amending the job completion distributions, the model can be re-run, and the change against the base line measured

These steps are explained in more detail below.

3.1 Step 1: Service Delivery volumes

The input data on job volumes provides information on the date and time of each job logged for financial years 2011/12 and 2012/13. The data is first adjusted to remove all jobs that do not require an engineer appointment (i.e., jobs classified under 'Other' in the database, and all 'Provision Frame' jobs). This is to ensure that only field jobs that require an engineering appointment are included in the Model.

The job is entered into the Model on the date it was logged, and at the appropriate geographic location, defined in terms of (i) GM area and (ii) MDF location. However, due to the size of the raw data input (approximately 20 million rows for each year), the Model adopts a simplifying assumption in respect of the time a specific job is input into the Model. Specifically the Model, instead of adopting the actual time the job was logged, splits the day into two half days, with the job assigned to a half day depending on the time it was logged. All jobs logged in a half day are entered into the system at a constant rate; for example, if across a 12 hour period, 48 jobs were logged, then the jobs enter the Model at a rate of one every 15 minutes. Jobs are input into the Model in the form of the number of jobs for each product by GM area by half day, which produces a total number of records in the input data set of approximately 36,000 per GM area per year.

3.2 Step 2: Care Level

The original data set specifies the product Care Level for each Repair job logged. However, rather than using the actual product Care Level, the Model creates a distribution of Care Level by product by summing the actual number of jobs in each Care Level by product, by month and by GM area, and then randomly allocates a Care Level to each job (in line with the calculated proportions). This approach was taken to minimise the volume of data (i.e., number of input rows) which require loading into the model, and does not have an impact on the accuracy of the calculation[11]. The total number of Care Level / product combinations is consistent with actual historic data.

The Care Level is used for defining the SLA target for Repair jobs only. Openreach's SLA target for Provision jobs does not vary by Care Level.

3.3 Step 3: SLA target

The SLA target for Repair jobs is calculated with reference to the time the job was logged by Openreach and the SLA target associated with the Care Level. For example, if a job is logged at 9:00am on Monday 25 August and is assigned Care Level 1, the SLA target would be Wednesday 27 August (i.e., end of next working day +1).

For Provision jobs, the SLA target is the same irrespective of Care Level, at 13 days from the day the call was logged. There was no formal SLA target for Provision jobs for the financial year 2011/2012. Therefore, 13 days represents a notional SLA target, against which notional performance levels can be derived for the base line. This has implications for interpreting the results for Provision jobs using 2011/12 data, as the lack of SLA target for Provision jobs relaxed the constraint on Openreach resourcing decisions that a formal SLA target for Provision jobs would have represented (e.g., during periods of high Repair activity, Openreach could take the decision to delay appointing Provision jobs without penalty). Therefore, the results derived using 2011/12 data should be regarded as providing only an indicative view of the relationships between different scenarios.

[11] Using a sampled random number to determine the Care Level to assign to each job does not affect the SLA target for Care Levels 1 to 3, as the SLA target is based on half or full working days. For Care Level 4, where the SLA is 6 hours, the random number approach may result in some misallocation of Care Level 4 jobs across the week. However, as Care Level 4 only accounts for 0.3% of the total jobs (i.e., 2-3 jobs per day per GM area), the impact is likely to be minimal. Further, there is no obvious pattern observable from the historic data as to when Care Level 4 jobs are logged, so using a random distribution is likely to represent a reasonable approximation. Note that the same sampled random numbers are used for each scenario – where Care Level 4 volumes and performances are kept constant – such that this approximation does not affect the comparison of results between scenarios.

The Provision SLA target is based on an 'Earliest Engineering Appointment' SLA, i.e., the 13 day SLA target refers to the ability of Openreach to offer an appointment for a Provision job within 13 days, rather than to complete the job. The input data used for the modelling did not capture the earliest available appointment offered, and so the calculated performances against the SLA target are the proportion of jobs actually completed within the SLA target. This is on the basis that if Openreach has the ability to offer an appointment within 13 days, it is assumed to have the ability to complete the job, irrespective of whether the customer subsequently requests an alternative date outside the 13 day period.

3.4 Step 4: Job completion

A distribution of the time taken to complete a job from when it was logged is derived for each product by Care Level and for each GM area. It is then possible to calculate Openreach's performance against its SLA targets by dividing the sum of jobs completed within the SLA target by all jobs completed. The SLA target is defined as zero and then completion time is calculated as a variance from this reference point, as shown in Figure 2 below:

Figure 2: Example distribution of variances in completion times from the SLA target

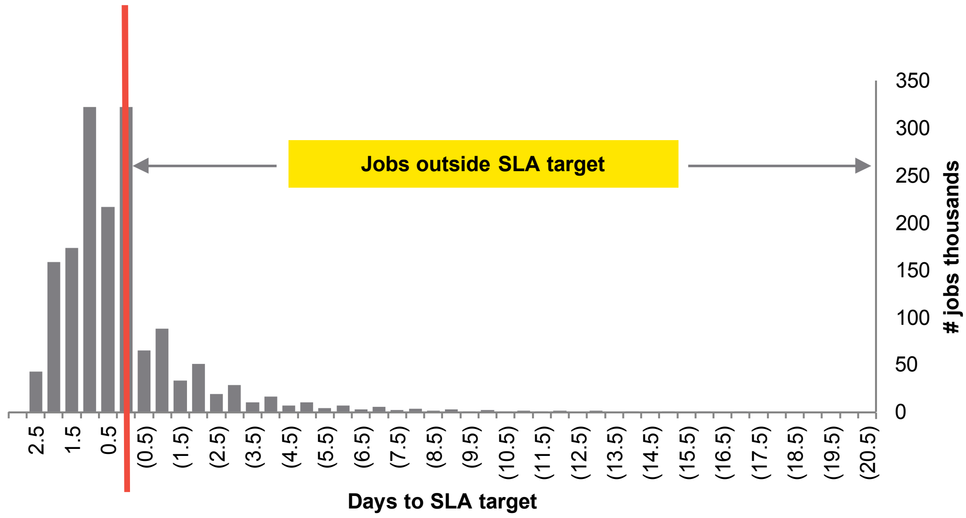

The derived distributions were compared with standard statistical distributions, and the gamma distribution (a distribution commonly used to model waiting times) was selected as the most appropriate distribution to use when the performance against the SLA target is adjusted in a Model scenario run.

A gamma distribution is generally defined by two parameters: (i) a shape parameter **α**; and (ii) a rate parameter **β**, which are both positive, real numbers and enable the Model to fully characterise the statistical properties of the distribution including mode and cumulative distribution. The mode of a continuous probability distribution is the value at which its probability density function has its maximum value (i.e., the mode is at the peak), while the cumulative distribution gives the area under the probability density function from zero to x_i, with x_i being the number of days to complete from the day the job was logged.

The observed actual distribution is used to identify the peak (i.e., mode) for each Care Level and the SLA target is used as the reference point to measure the cumulative distribution (i.e., levels of performance). For example, Figure 3 below shows an actual and a fitted gamma distribution for Care Level 1 jobs in a single GM area with a weekly performance in respect to the SLA target of 78%.

Figure 3: Example of actual and fitted gamma distributions of job completion at a given level of performance

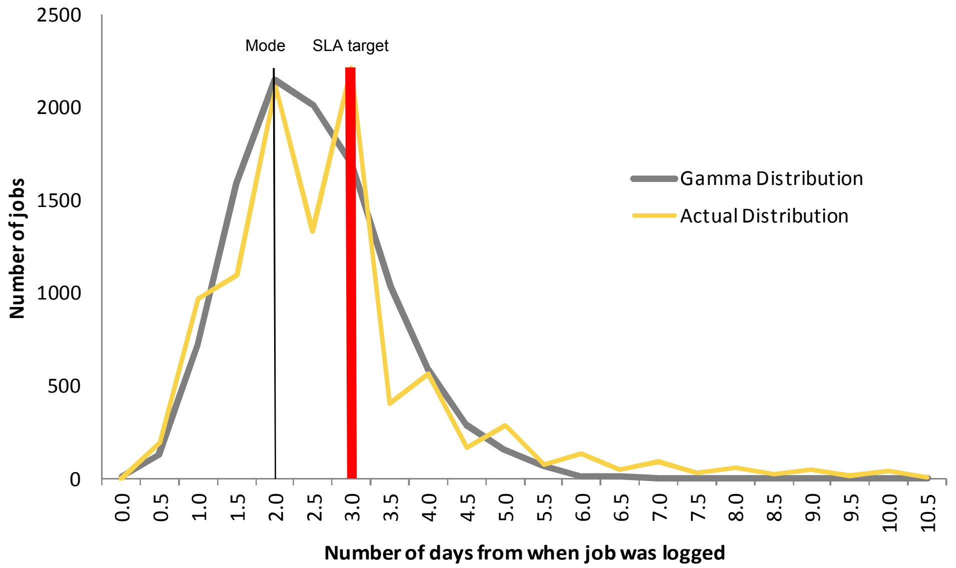

In order to run the Model for different performance levels by Care Level, a separate gamma distribution is derived by Care Level by adjusting the gamma distribution parameters. This is achieved by solving numerically the system of two equations for the two unknown parameters α and β and by assuming the mode (*m*) remains constant and is equal to $(\alpha\text{-}1)/\beta = x_m$ (where x_m is the known position of the mode) and setting the cumulative distribution calculated between zero and the SLA target equal to the desired performance level. The numerical tolerance on the performances is 0.1%. Figure 4 below shows an example of adjusted gamma distributions at different levels of performance.

Figure 4: Gamma distributions under different performance levels

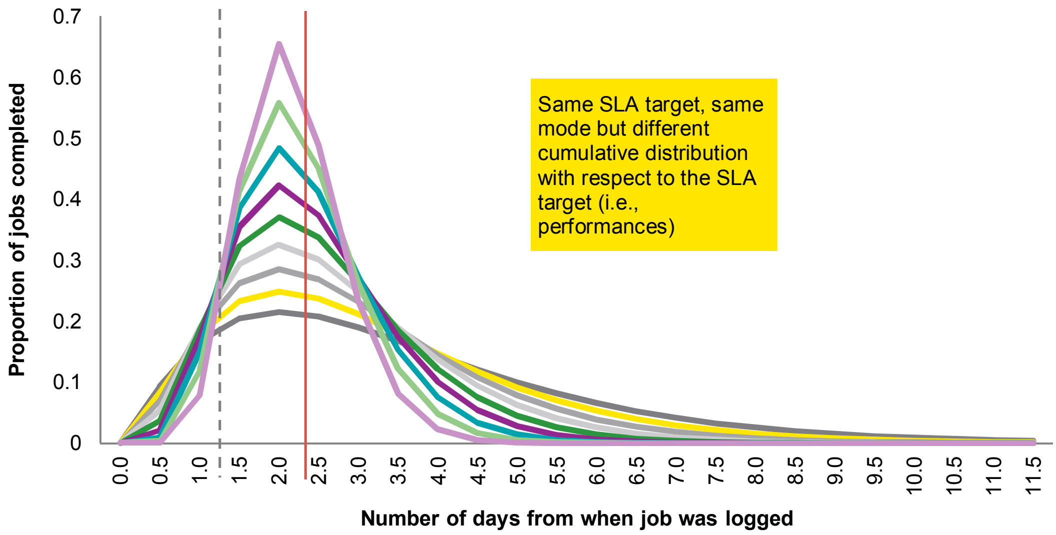

The gamma parameters used in the Model depend on the SLA target and the performance levels. Weekly performances are rounded to the nearest integer to limit the number of gamma distribution parameters required. The gamma distribution is then used to assign a completion time to each job in the simulation, based on its Care Level.

3.5 Step 5: Task times

Each job is assigned a task time, i.e., the time it takes an engineer to complete the job. Openreach calculates an average task time by product for Provision jobs, and by 'network clear code' in the case of Repairs, referring to the section of the network in which the fault was identified and subsequently cleared, by month and by GM area.

A proportion of jobs require multiple visits but, as a simplifying assumption to reduce the Model complexity, a single visit per job is modelled, with the applied average task times reflecting all the time booked against jobs, including the time spread across multiple visits.

The task time is subtracted from the job completion time (derived from the completion distribution) to derive a job start time. The period between the start time and the completion time is the time that the job is 'active', i.e., how long a resource will be 'busy' for on that particular job. For example:

- Completion Time = 04/08/2011 14:23:30
- Average Task Time (mapped to the job) = 2.5 hours
- Start Time = 04/08/2011 14:23:30 *minus* 2.5 hours = 04/08/2011 11:53:30

Figure 5: Example derivation of job start time

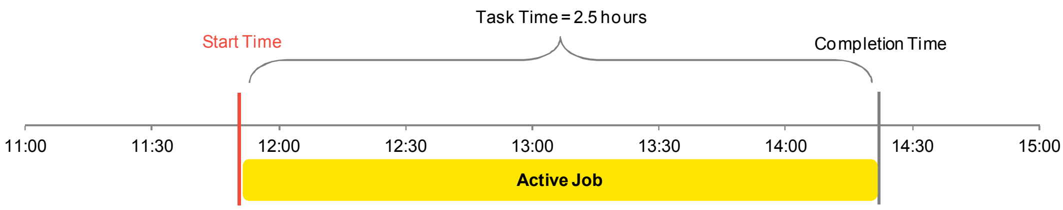

3.6 Step 6: Job queuing

Each job is assigned a 'waiting time', which is calculated as the difference between the time the job was logged and the calculated start time of the particular job. Jobs are held in a queue for the waiting time, i.e., until the Model clock reaches the calculated start time.

3.7 Step 7: Daily engineer demand

A job leaves the queue and becomes active at its calculated start time. An engineer of the appropriate skill type is assigned to the job, and is assumed to be available for the remainder of the shift. Active jobs are stacked up over the day, and in the event that no engineer with an appropriate skill level is available to complete the job, a new engineer is added to the system. An example is shown in Figure 6 below which shows that, although a maximum of four jobs overlap at any one time, different skill sets are required at different points in the day and, as a consequence, a maximum of two engineers of Skill Type 1, one engineer of Skill Type 2 and three engineers of Skill Type 3 are required.

Figure 6: Example of job stacking by performance type

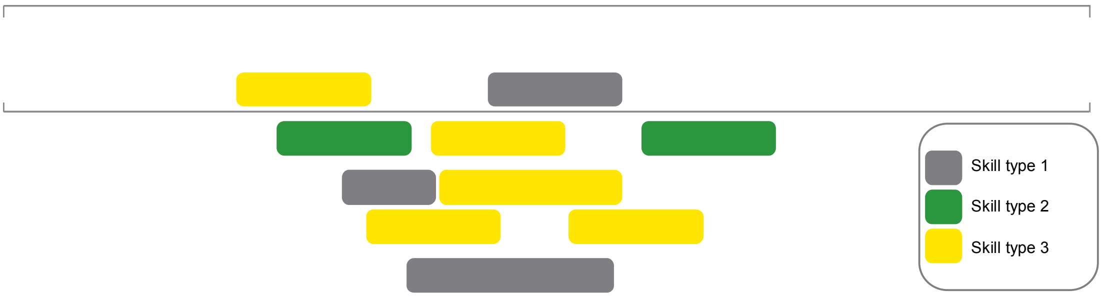

Over the course of the day the Model assumes two shifts for engineers: 08:00 – 16:00 and 11:00 – 19:00, consistent with Openreach's daily shift patterns. The modelled total number of required engineers for a given day is the higher of: (i) the sum of the maxima of each shift and (ii) the maximum in the period when the two shifts overlap. An example is shown in Figure 7, where the maximum number of engineers required is 1,177, which is the sum of the individual shift maximums (which in this example occur outside the overlap period), as opposed to the maximum during the overlap period (736).

Figure 7: Deriving the maximum number of engineers in a day

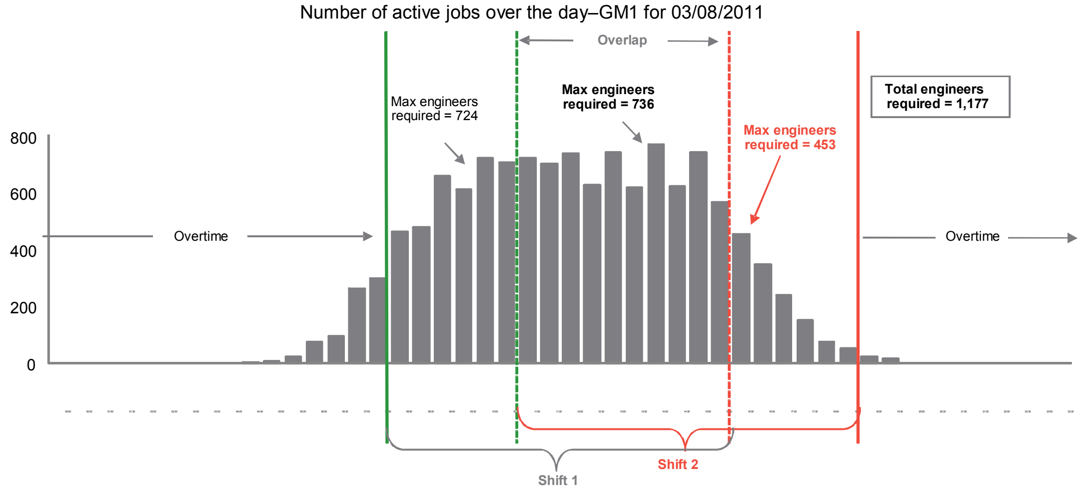

The same calculation is carried out for each day of the year, and this generates the number of engineers required each day to complete the required number of jobs at the given performance level, as shown in Figure 8 below.

Figure 8: Example of the required number of engineers by day

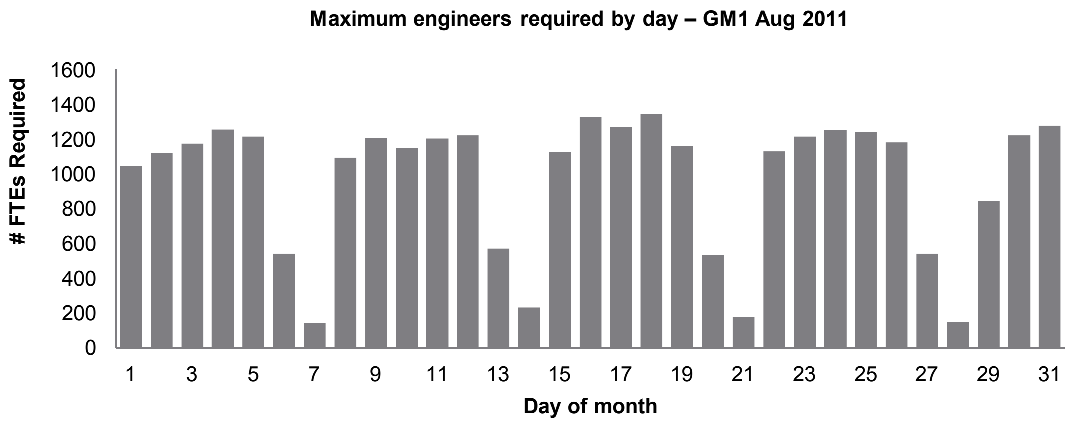

3.8 Step 8: Output

The daily profile of required engineers by skill type is used to derive a single figure for total engineers required by skill type for the year to meet the level of performance and SLA targets implied by the distributions. The Model calculates the maximum number of engineers over the year by skill type, taking into account the fact that highly skilled engineers, when available, can perform jobs of a lower skill type. This approach is referred to as the 'redistributed max'. Figure 9 below shows a summary example of the redistributed max calculation for two skill types. The left hand table shows example output for the daily demand for engineers of two skill types, CSE and CAL, where CSE engineers are more skilled than CAL engineers. For example, on 26/03/12 the Model calculated the need for 163 CSE engineers and 279 CAL engineers. The right hand table shows the conversion of this daily demand for each skill type into a total demand by skill type for the modelled period.

The first step of the redistributed max calculation is to identify the day with the highest demand for the most skilled resource (in this example CSE engineers). In the example, the maximum requirement was 216 (03/04/2012). The second step is to identify the day with the highest requirement of the next skilled resource (CAL engineers). To model this on each day the number of unused CSE resource is subtracted from the required number of CAL resource. For example, on 26/03/2012 there were 163 CSE jobs, and 279 CAL jobs. Therefore the required number of CAL resource was 279 – (216-163) = 226. This calculation is repeated for every day of the simulation to determine the highest number of CAL engineers required. In the Model this calculation is conducted for the whole year for all individual skill types. Once the above steps have been completed, the total number of engineers required is calculated by summing engineers by skill type.

Figure 9: Deriving the number of engineers required by skill type

Simulation output				Adjusted output				
Date	CSE	CAL	Total	Used CSE	Available CSE	Required CAL	Used CSE for CAL	Used CAL
26/03/2012	163	279	442	163	53	226	53	226
27/03/2012	177	307	484	177	39	268	39	268
28/03/2012	214	289	503	214	2	287	2	287
29/03/2012	177	285	462	177	39	246	39	246
30/03/2012	191	275	466	191	25	250	25	250
31/03/2012	16	20	36	16	200	(180)	20	0
01/04/2012	2	7	9	2	214	(207)	7	0
02/04/2012	184	284	468	184	32	252	32	252
03/04/2012	216	284	500	216	-	284	0	284
04/04/2012	186	322	508	186	30	292	30	292
05/04/2012	194	308	502	194	22	286	22	286
06/04/2012	182	256	438	182	34	222	34	222
Required engineers				**216**				**292**

Alternative approaches to determining a single required number of engineers derived from the Model output were considered, and their merits and demerits are discussed in Appendix A.

3.9 Step 9: Flexing performance

The Model is run initially to create a base line with completion distributions that match actual performance against SLA targets. To understand the change in resource requirements for performances and/or SLA targets, the Model is then re-run using revised distributions, but with the same total volume of jobs. The distributions can be flexed in a number of different ways including:

► Revised SLA target but constant performance: if the SLA target (defined in terms of the allowed time to complete a job) is amended for a specific Care Level, then a revised gamma distribution is derived such that the proportion of jobs completed within the new SLA target is the same as under the base line SLA target. For example, if the SLA target is reduced, but performance is assumed to remain constant, then the mode in the individual distribution for a specific job types increases in order to maintain the area under the graph (which represents the performance against the SLA target). Conversely, increasing the SLA target means more time to meet the target and therefore the mode decreases in order to maintain the area under the graph.

- Revised performance but a constant SLA target: for scenarios where the modelled level of performance is changed, the distribution of job completions is amended, with the SLA target remaining constant. For example, for scenarios where performance is increased from the base line, the distribution is adjusted such that a greater proportion of jobs are completed within the SLA target.

4. Detailed modelling issues

A number of specific modelling issues arose during the course of model construction, and the approaches taken to reflect these issues are discussed in the following section.

4.1 Glass ceiling

Openreach has identified a 'glass ceiling' to the successful completion of Provision or Repair tasks on the day the task is issued to the Openreach engineer, which is considered to arise due to a range of factors outside the immediate control of Openreach. Examples include engineers being unable to gain access to premises, tasks requiring civil engineering and tasks requiring specialist tools or skills not available to the engineer at first visit. Based on previous analysis, Openreach estimates that these issues effectively place a limit (or 'glass ceiling') on the ability to successfully complete tasks on a given day, with Openreach's assessment of the limits being 79.5% for Repair activity and 83.7% for Provision activity.

Whilst the above analysis does not imply that there is the same degree of limitation for performance against SLA target (for example as there will be some cases where a follow-up task can be successfully issued and completed within the SLA target), Openreach considers there to be a maximum level of completion performance that could be achieved within the SLA target, as a consequence of the glass ceiling on daily task completion.

To reflect this, functionality has been added to the Model to allow a constraint to be placed on weekly performance, such that it cannot rise above a defined maximum percentage performance against SLA targets. As a consequence, to meet an overall annual performance target, higher performances need to be achieved in lower performing weeks to compensate for this constraint. Openreach conducted its own analysis, and provided a figure of 90% as the level at which the constraint should apply in the Model. The rationale for this figure is provided by Openreach in a separate document entitled *"Openreach analysis of additional factors impacting service costs in very high performance scenarios."*

Openreach also considers that, as the level at which these practical limitations impact is approached (i.e., as performance against SLA targets approaches the maximum that is practically possible), it becomes more time-consuming to complete the marginal job. Further, under periods of high demand for resource in a particular area, there is likely to be a greater proportion of jobs performed by engineers working outside their normal geographic working area, and that this would tend to increase task times (both due to the relative unfamiliarity of the engineer with the location as well as increased travelling between jobs and to/from the engineer's normal home location). The occurrences of such jobs performed by 'off-patch' engineers will include jobs performed by engineers outside their normal working area but within their 'home' SOM patch, as well as jobs performed outside the home SOM patch.

These impacts are reflected in the Model by uplifting task times under scenarios where performance approaches 90%. In these scenarios, uplifts of between 5% and 10 have been applied to Repair task times and of between 0% and 5% to Provision task times. The rationale for the task time adjustments is provided by Openreach in the document referred to above.

As discussed in Section 3, to model changes in performance it is necessary to amend the distribution of performance against the SLA target, such that the required level of performance is achieved when the Model is run. However, as there are significant variations in performance by week, there exist a number of approaches that could be adopted such that the desired average annual level of performance is achieved. The modelled approach is to maintain the weekly performance profile, but increase (or decrease) weekly performance by the same required percentage change in annual performance. For example, if the base line annual performance was 60% of jobs completed within the SLA target, and the target performance for the scenario to be run is 80% then the performance in each week and across each GM area is uplifted by an increase of 20 percentage points.

As discussed above, a constraint on weekly performance of 90% is imposed in the Model to reflect the glass ceiling identified by Openreach. Where required performance in a week would exceed this constraint, performance is limited to 90%, and higher performance (to account for the shortfall resulting from the glass ceiling) is required in other weeks. The approach taken in the Model is to uplift performance in the lowest performing ten weeks of the year first. An example of how this would impact weekly performance is shown in Figure 10 below.

Figure 10: Impact on weekly performances from applying the 'glass ceiling'

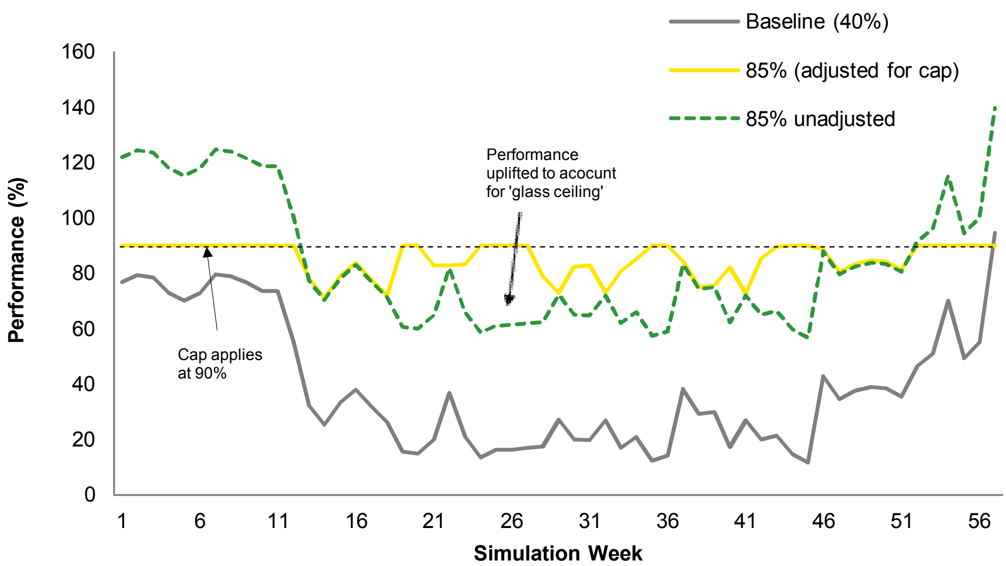

In 2012/13, a different approach was adopted for measuring changes in performance levels against SLA targets to reflect the observation from actual data that there was considerable variability in performance against SLA targets over the course of the year. As a consequence, Openreach observed that the need to complete jobs that had already failed SLA targets (the 'backlog') was a key driver of the required number of engineers in the base line run, particularly towards the end of the financial year. As a result, Openreach considered that the most appropriate means of measuring a delta between base line and scenarios was to adjust the base line such that performance against SLA targets was held constant in every week at the annual average. The impact of higher performance was then derived using the same approach as for 2011/12, i.e., a constant percentage point increase in weekly performance.

4.2 Application of the Model to different Provision SLA target scenarios

The historic data on the completion of Repair and Provision jobs demonstrates cyclic weekly patterns. These patterns can be explained largely by the underlying Care Levels and SLA targets associated with Openreach products[12]. As a result, in the historic data set, both Provision and Repair jobs tend to have a weekly cycle, characterised by a typical day of the week where job volumes are highest, with the peaks aligned during the middle of the working week. As a consequence, different scenarios for SLA targets could lead to results which fail to appropriately articulate the impact on the Openreach business of, for example, increasing the stringency of SLAs, on resource requirements.

In particular, an 11 day Provision SLA target will result in the modelled peak day of activity falling on a Monday, which would result in the Model calculating a lower required number of engineers than for a 12 day Provision SLA target, because it would no longer align with the

[12] For example, weekend calls from customers regarding line faults for Care Level 1 and Care Level 2 products are not logged by Openreach as a Repair job until Monday, as Sunday is not considered a working day for these Care Levels. The same is true of requests for Provision jobs received over the weekend.

peak for Repair jobs. Openreach considers that such an outcome would not reflect the actual operational situation, where there exist practical limitations to completing Provision jobs within a shorter SLA target, where a shorter SLA target for Provisions would place a constraint on the completion of Repair jobs, and where Openreach resource management decisions would be revised reflecting both the supply of engineers as well as the demand.

To reflect this, the modelling approach is adjusted when scenarios are run for an 11 day Provision SLA target. The first step is to exclude Repair jobs from the Model such that the effect on required resource for Provision activity can be isolated. The gamma distributions for 13, 12, and 11 day SLA targets are then adjusted such that the mode of completions is fixed on the same day regardless of the SLA target. The Model is run for the base line, for each of the different SLA targets and for different scenarios on performance against those targets. The resulting deltas are used to extrapolate the results from the 12 day Provision SLA target scenarios from the unadjusted Model runs to derive a delta for an 11 day SLA Provision target.

Openreach considers that adjustment of the gamma distributions and the use of results derived from extrapolation will add to the degree of modelling error. Therefore, it considers that any results for 11 day or other Provision lead time scenarios should be regarded as only being indicative estimates of cost implications.

4.3 Application of the Model to different Care Level volume mix scenarios

The Model contains the functionality to amend the proportion of jobs between Care Levels (holding Care Levels and job volumes constant) and to assess the impact of a change in the 'Care Level mix' on the number of required engineers. However, in order to perform this analysis, Openreach considers that a number of considerations need to be taken into account:

- Working week: Care Level 2 jobs include Saturday as a working day for the purpose of calculating the SLA target and for resourcing. However, Openreach employs a much smaller Saturday shift (i.e., a lower number of available engineers) compared with a normal weekday. As a consequence, shifting all Care Level 1 jobs to Care Level 2 would result in a higher number of jobs assigned to a Saturday than could be performed in the real world due to engineer availability constraints. Therefore, when running scenarios relating to a changing Care Level mix in the Model, the Care Level working week is adjusted to be consistent with Care Level 1 (i.e., Saturdays have been excluded[13]);

- Performances: the base line performances (against SLA targets) differ across Care Levels, and this difference in performance therefore requires recognition in the modelling in the event that the Care Level mix is changed. To reflect this, when jobs are 'moved' between Care Levels 1 and 2, the average base line performance for each Care Level is adjusted to reflect the shift of jobs, i.e., performance is typically higher for Care Level 1 jobs, and so where a Care Level 1 job is re-designated as a Care Level 2 job, the higher level of performance is reflected in the base line Care Level 2 performance distribution;

- Interaction with Provision jobs: changing the Care Level mix can result in significant changes in the profile of Repair work over the course of a week. For example, the SLA target for Care Level 2 is one day shorter than for Care Level 1 (i.e., end of next working day rather than end of next working day +1), and so an increase in the number of Care Level 2 jobs may lead to a material change in daily jobs compared with the base line to deliver a given performance level. The Model also includes Provision jobs, and it is the sum of the Provision and Repair jobs that drives daily - and indeed weekly - peaks. Therefore, in order to isolate the effects of Care Level mix changes on resource requirements for Repair jobs, Provision jobs are excluded from the Model; and

[13] An alternative would be to hold the proportion of work completed on a Saturday constant, but this approach added additional complexity and the approach taken is considered a reasonable approximation.

► Peak day: The mode of the distribution of performance against SLA targets differs for Care Level 1 and Care Level 2 as a result of the different SLA targets. For example, in Figure 11 below, the dotted lines represent the distributions of performance for Care Levels 1 and 2, showing that the mode for Care Level 2 is one day after the call was logged, and the mode for Care Level 1 is two days after the call was logged. As a result, as the Care Level mix approaches the extremes (i.e., all jobs on Care Level 1 or all on Care Level 2), the modelled scenario may overstate the requirement for engineers compared with the base line due to the correlation of the modes. To account for this, the gamma distribution for Care Level 1 has been smoothed such the mode aligns with the mode for Care Level 2. A new base line is then run which reflects this new amended distribution, and scenarios are run against this base line to derive the resource requirement deltas.

Figure 11: Care Level performance distributions

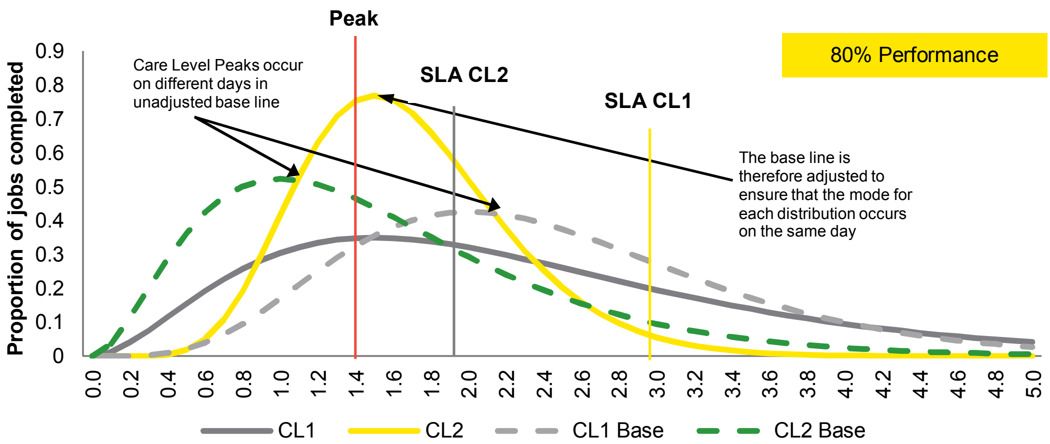

Appendix A

The modelling approach used to derive the annual resource requirement is set out in Section 3 above. Specifically, the approach that has been adopted is referred to as the 'redistributed max' which reflects the fact that highly skilled engineers can undertake lower skilled jobs when available. Alternatives considered were:

- ► Highest absolute number of engineers required on any day in the year, i.e., the highest number of engineers required on a single day over the course of the year irrespective of skill type;
- ► Average of the number of engineers required on the *x*[14] busiest days of the year irrespective of skill type;
- ► Average number of engineers required over the year, i.e., simple average of the daily engineer requirements; and
- ► Redistribution approach (as set out in Section 3) but taking average of *x* busiest days rather than absolute busiest day, i.e., the approach used in the Model, but rather than basing the calculations on the absolute busiest day for any skill type, the average number of engineers required on the busiest *x* number of days is considered.

The advantages and disadvantages of each approach are set out in Figure 12 below.

Figure 12: Different approaches to calculating engineer requirements and their associated advantages and disadvantages.

No.	Approach	Advantages	Disadvantages
1	**Redistributed max**	► Accounts for the different skill types required ► Minimises resources required by allowing for cross flexing of skill types	► More complex to calculate
2	**Redistributed X**	► Accounts for the different skill types required ► Allows for cross flexing of skill types	► More complex to calculate ► Results in higher level of low skilled resource required which is not reflective of Openreach's workforce
3	**Max number of engineers**	► Simple to calculate	► Reflects extreme event, e.g., single outlying data points ► Ignores interaction between skill types
4	**Average number of engineers on busiest x days**	► Simple to calculate	► Ignores interaction between skill types ► The bigger the value of x, the greater the method would systematically understate the effect the Model is being used to investigate, i.e. the increased volume of local work peaks by location, by skill, etc.
5	**Average number of engineers required over the year**	► Simple to calculate	► In the approach total number of jobs is the same, the Model just redistributes to different time slots. Hence average number of engineers is unlikely to understate impact of different scenarios ► Ignores skill types

On the basis of these relative advantages and disadvantages, the redistributed max approach has been adopted using the highest engineer requirement by skill type, as it was considered by Openreach to most accurately reflect their resource availability and scheduling, and in particular the ability of engineers of high skill levels to complete jobs with lower skill requirements.

[14] *x* could be any whole number between 1 and 365.

EY | Assurance | Tax | Transactions | Advisory

Ernst & Young LLP

Ernst & Young LLP, 1 More London Place, London, SE1 2AF.

ey.com