Deloitte.

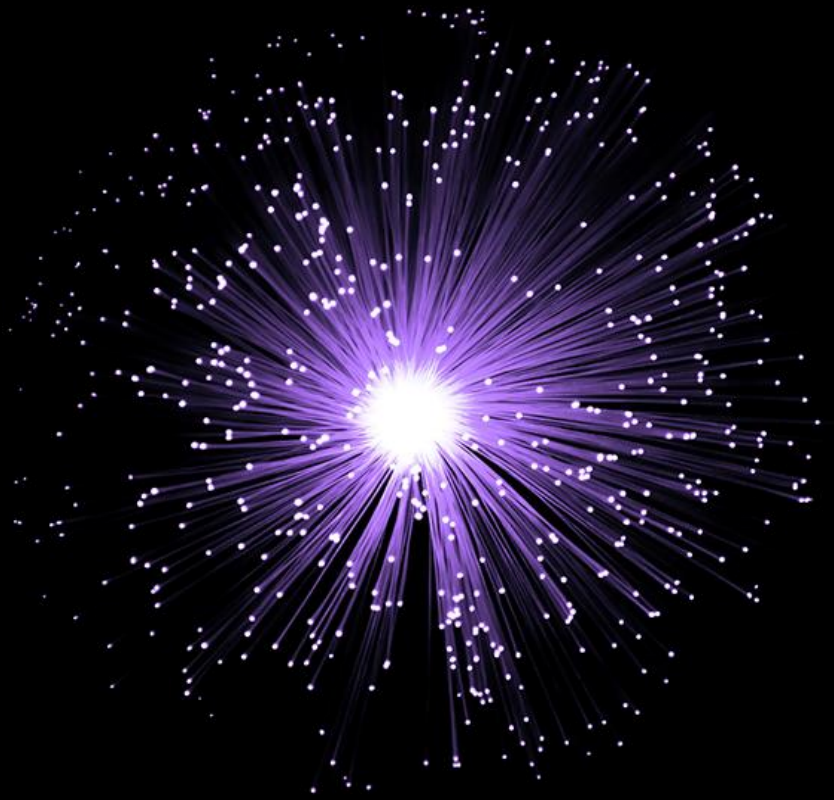

Openreach Quality of Service modelling

5th July 2017

Important notice from Deloitte

Our consent to disclosure of this report is on the following basis:

The reader acknowledges and agrees that:

(a) The report was prepared solely for the use of Openreach;

(b) The report was prepared from information and explanations provided to us by the management of Openreach;

(c) Matters may exist in the information that might have been assessed differently by you;

(d) The information contained herein is not designed to form the basis of any decisions made by you;

(e) We have not updated the report for any events or transactions which may have occurred subsequent to the date of the report

(f) Our consent that the report is made available to you does not establish any client relationship or any other contractual or other relationship between us;

(g) You will not use the report for any purpose other than to supplant other enquiries or procedures you might undertake for your purpose;

(h) We do not warrant the suitability or sufficiency of the report for your purpose;

(i) We have no duty to, and will not:

(i) Monitor the subject matter of the report;

(ii) Update the Information in respect of any events or transactions that (a) might have occurred subsequent to the Report's completion or (b) may in future occur;

(iii) Provide you with any other or additional information, or;

(iv) Correct any inaccuracies in the report that might become apparent to us in the future.

(j) The DTTL Parties neither owe nor accept any duty of care to you in connection with the report and none of them shall be liable to you for any loss, damage, cost or expense of whatsoever nature which is caused by your use of or reliance on the report. Should you choose to rely upon the report, you do so entirely at your own risk and you are responsible for carrying out your own independent investigations;

(k) You agree not to bring, or threaten to bring, any actions, proceedings or claims against any of the DTTL Parties in any way arising out of or in connection with the report or its use; and

(l) You will not obtain any copyright or other intellectual property rights in the information made available to you.

Table of contents

Key findings

Deloitte have undertaken a review of the methodology, inputs and assumptions of the Openreach Allocation model (developed with support from Ernst and Young) and the Ofcom Resource and Performance model built by Analysys Mason. Both of these models aim to estimate the uplift in resources necessary to achieve the revised performance standards proposed by Ofcom.

However, the two models yield very different results, with Ofcom arguing (based on the Analysys Mason model) that the proposed performance standards can be met with a resource uplift of 4.7% while Openreach argue that these targets may not be achievable, even with a large number of additional resources.

Overall, the review indicates that there are a number of features of the Analysys Mason model that means it is more likely to understate the costs associated with the stipulated performance:

- The Analysys Mason model has less functionality to capture variations in task duration, travel time and engineer skills requirements. As such, the model requires a number of simplifying assumptions to be made including assuming a flat task completion time and universal multiskilling of engineers. These assumptions will tend to understate the additional costs associated with dealing with volatility in the taskload. In contrast, the Openreach model uses real distributional data on task times, travel times and task type to capture this variation.

- The Analysys Mason model allows for resources to be shared across regions in order to address variance in demand. By including this functionality, the Analysys Mason model goes beyond the Openreach Allocation model in capturing the realities of task allocation. However, it is not clear from the evidence available whether the level of resource loans between regions implied within the Analysys Mason model is consistent with Openreach practices. In particular, the Analysys Mason model may understate the challenges and productivity costs associated with resource loans; in conjunction with the smoothing of task and travel times, this could lead to underestimation of the costs of managing peaks and troughs in demand and staff availability.

- The two models also make significantly different assumptions about the on-the-day failure rate, which determines the glass ceiling on potential performance. While there is an ongoing debate between Ofcom and Openreach about the extent to which this on-the-day failure rate may be reduced over time, Analysys Mason acknowledge that initially the on-the-day failure rate assumptions had to be reduced below the rate indicated in Ofcom data in order to meet Ofcom's proposed performance standards. While this has been addressed through changes in the methodology for modelling on-the-day failures, Analysys Mason's sensitivity analysis indicates that their glass ceiling assumption remains an important driver of model outputs.

The Openreach model also has its limitations, for example in not endogenously capturing the fact that resources can be shared across regions. However, given the greater sophistication of the model and its ability to reflect the flow of tasks at a greater level of granularity, it is believed to better represent the current reality.

Executive summary

Executive summary

This report presents the findings of a review of the differences between Analysys Mason's Resource and Performance model and Openreach's Allocation model

Openreach is the largest owner and operator of local access fixed telecommunication networks in Great Britain. The network operator is in charge of line installations, disconnections and fault repairs from homes to exchanges.

Since the 2014 Fixed Access Market Review (FAMR), Ofcom has imposed Quality of Service (QoS) remedies, articulated in two steps. First, Ofcom assigns each regulated product a time limit by which Openreach is required to complete its repairs or offer a provision appointment. This time limit was set 1 to 2 days after the fault for repairs (depending on care level) and 12 days for provisions.

The regulator then prescribes Openreach a performance target in terms of the proportion of repairs and provisions which are to be completed by the set deadlines. For 2016-17, this performance target was established at 80% of repairs and 90% of provisions.

Ofcom is now considering increasing Openreach QoS standards (see Table 1 and 2). The suggested new standards are informed by a model of task allocation and delivery developed by Analysys Mason for Ofcom; Openreach have also developed their own model with EY to assess the resource requirements necessary to deliver these standards.

The two models point at materially different costs of performance (see Table 3).

In order to help address these differences, Openreach has commissioned Deloitte to conduct an independent review of the two models. This review is to consider the two models' methodology, inputs, assumptions and testing.

The findings of this review are presented in this draft report.

Table 1: Suggested new provision standards

Provision standards	2016-17	2017-18	2018-19	2020-21
% completed by committed date (force majeure)	90% (89%)	92% (91%)	92% (91%)	95% (94%)
Working days within which first available date (FAD) offered	12	12	12	10
% with regulated appointment date (force majeure)	80% (79%)	90% (89%)	90% (89%)	90% (89%)

Table 2: Suggested new repairs standards

Repair standards	2016-17	2017-18	2018-19	2020-21
% of repairs to be completed within SLA timescales (force majeure)	80% *(77%)*	83% *(80%)*	90% *(87%)*	93% *(90%)*
% of repairs to be completed within SLA +5 days	N/A	95%	96%	97%

Table 3: Cost of performance (repairs and provision)

Model output	Ofcom*	Openreach
FTE increase required to achieve 95% provision and 90% repair performance	3.8% - 6.8%	24.9%
FTE increase required to achieve 95% provision and 93% repair performance	4.7% - 8.8%	47.6%

* Ofcom's 93% repair performance scenario is based on actual care level mix and leaves care level 2 performance at 90%. Its 90% repair performance scenario is based on a 50-50 care level mix.

Executive summary

The two models have a similar underlying approach but the Analysys Mason model developed for Ofcom makes more simplifying assumptions

Openreach Allocation Model		Ofcom Resource and Performance Model
• **Tasks:** Based on probabilistic distributions which are derived from April 2014 to December 2016 data. • **Resourcing:** The user sets the available pool of engineers by skill and SOM (Senior Operating Manager) area. This varies due to shrinkage (including travel to other areas), attrition and hiring.	**Inputs**	• **Tasks:** Historical fault and installation data for 2011-12 to 2015-16. • **Resourcing:** The number of FTEs is determined by the user and remains constant throughout the simulation. FTEs can be transferred from one SOM to another at 75% productivity.
• **Workstack:** All jobs are set in the same queue, with jobs being prioritised according to their care level, appointment status and due date. • **Allocation:** Jobs are allocated to the first available engineers with the least skill; tasks are prioritised so as to balance performance.	**Queuing**	• **Workstack:** Care level 1, Care level 2 and installation jobs are set in separate queues. • **Allocation:** Engineers are allocated in priority to appointed provision jobs, then due repairs jobs, tail jobs and lastly jobs which are not yet due.
• **Completion:** Task times are sampled from frequency distributions based on Openreach historical data. Travel times are sampled from an independent distribution for each visit. • **Glass ceiling:** Major on-the-day failures are assumed to fail automatically; this is based on historical data. Minor on-the-day failures require multiple visits to be completed.	**Completion**	• **Completion:** Any engineer is assumed to be able to complete 4 tasks (of any level) per day; this is implicitly assumed to include travel time. • **Glass ceiling:** Each job completion may result in a major failure, which results in a 5 day delay, or a minor failure, which results in a one day delay, according to a certain probability.
• **Output:** The model's final output is a performance level achieved for a number of resourcing scenarios; although the model can accommodate many other types of scenarios.	**Output**	• **Output:** As in Openreach's model, the final output is a performance level achieved for a number of resourcing scenarios.

Executive summary

There are a number of points of difference between the models, including Analysys Mason assumptions that understate task complexity

As shown in the previous slide, while the underlying methodology of the two models is similar, there are a number of areas in which they differ, namely in the assumptions made and the level of granularity of the analysis. These differences may in turn bias the results in a particular direction.

The table below summarises the key points of discussion, including:

- The key differences.
- Implications for the model outputs and testing of impact.
- View on reasonableness of model approach: whether the evidence suggests that the Openreach or Analysys Mason approach and conclusion are better supported by the data. The position of the arrow indicates the extent to which one model is seen to be a better representation of reality and better supported by the data.

Point of difference	Implications	Analysys Mason	Openreach
Glass ceiling The Analysys Mason model uses a lower on-the-day failure rate assumption than what is supported by Ofcom or Openreach's data. In Openreach's model, major on-the-day failures are assumed to fail automatically while minor on-the-day failures result into multiple visits. The rates of on-the-day failures is based on historical data.	The level of the glass ceiling both caps performance and increases the cost of improvements as they approach this theoretical limit. Analysys Mason's sensitivity testing shows that increasing the repair minor fail rate from 3% to 5% almost doubles resource requirements from a 4.7% uplift to a 8.8% uplift.	 Current data supports the Openreach assumptions. The Analysys Mason model assumptions are not supported by either Ofcom or Openreach data and appear to have been reverse-engineered.	
Engineer skill mix The Analysys Mason model assumes all engineers can complete any task. In practice, some assignments require underground access and more specialised engineers. Upskilling programmes may mitigate the impact of this assumption in the longer term.	This universal multiskilling assumption understates the actual constraints on Openreach's resourcing and as such may lower the cost of the suggested requirements. Testing the assumption of universal multi-skilling within the Openreach model indicates that it raises baseline performance by up to 4 percentage points and slightly reduces the cost of raising performance to the required level, but only by about 1 percentage point.	 The Analysys Mason model does not capture the reality of skill mix, but this does not have a significant impact on resource needs to improve performance.	

Executive summary

The Analysys Mason model makes simplifying assumptions that reduce the variance in tasks and resources, which may lead to underestimation of costs

Point of difference	Implications	Analysys Mason	Openreach
Resource variability The Analysys Mason model assumes resourcing is constant within each SOM. In practice, resourcing varies significantly with training, holidays and sickness; this routinely represents 40-50% of the workforce.	This constant resourcing assumption could lead to an understatement of the true cost of performance, since the variation in staffing increases the risk of a backlog being created. Ofcom recognises that a small decline in resourcing can have a large cumulative impact on performance. Sensitivity testing indicates that assuming flat resourcing reduces the resource uplift required by up to 1.4 percentage points.	 The Analysys Mason model approach represents a simplification that would tend to reduce variance in resources relative to tasks and therefore reduce costs of performance. However, this assumption's impact on outputs is moderate.	
Task and travel time variability The Analysys Mason model assumes a constant task time and does not account for variation in travel time. In practice, there is significant variance and the distribution is highly skewed, with a long tail of more complex and time-consuming tasks.	Ofcom note that "in practice job durations will vary. Not taking this into account could overlook variations that could on some days have a material impact on the overall number of jobs that can be completed during the day."[1] However the shape of the task distribution – in particular the fact that while there is a long upper tail, the majority of tasks take less than the mean time – can mitigate against this effect, making it easier to clear simpler tasks.	 While the Analysys Mason model oversimplifies the variation in task and travel times, scenario testing within the Openreach model in fact suggests that this assumption may understate the cost of delivering stipulated performance by up to 2 percentage points. On aggregate, flattening task times is estimated to slightly reduce the cost of reaching proposed targets, although in some GMs the cost increases.	

1- Ofcom, Quality of Service for WLR, MPF and GEA, 2017 page 207

Executive summary

The ability to transfer resources across regions tends to reduce variance and costs; the models differ in the extent to which this is captured

Point of difference	Implications	Analysys Mason	Openreach
Operating level variance Both the Analysys Mason and Openreach models simulate performance at SOM level. In practice, engineers work within Preferred Working Areas (PWAs) that are much smaller than SOMs.	Ofcom recognise that "By setting regulatory standards at a GM (General Manager) level, measured on an annual basis, on a large volume of work undertaken by a large workforce, the effects of variability in input factors is muted."[1] This would suggest that both models may understate the challenges associated with higher levels of task volume variance, although the Openreach Allocation model may partially capture this through task and travel time variance. However, in practice this may be mitigated by higher levels of staff loans between PWAs as required, allowing some of the impact to be smoothed.	 Neither model allows for task assignment at the PWA level; this suggests that both may understate the costs of performance. The Openreach model may capture some of this effect through task and travel time variance, which increases the variance of the overall workload.	
Patch loans The Analysys Mason model allows for resources to be loaned across SOMs in order to meet targets, based on proximity and/or supporting the worst performing SOM. It is assumed that in this case engineers' productivity falls by 25%. The Openreach model does not incorporate these loans, but uses actual shrinkage data alongside actual task data to capture fluctuating resources across areas.	Allowing for loans across regions enables Openreach to smooth risk across areas and improve performance. This reduces the additional resources required to meet targets. By not allowing workers to transfer across regions, the Openreach model imposes an additional constraint, making it more challenging to meet targets. In practice, evidence suggests that while patch loans occur, they tend to involve few personnel. They are also associated with increases in travel time that reduce productivity. Information on the patch loans distribution generated within the Analysys Mason model has been requested from Ofcom but has not been made available for the purposes of this review.	 Sensitivity testing indicates that the Analysys Mason outputs are sensitive to this assumption. However, the lack of data on the implied frequency of patch loans in the Analysys Mason model means that the accuracy of their approach is unclear. While Openreach may overstate costs by not modelling patch loans, the Analysys Mason model may overstate the frequency and ease of these loans.	

1- Ofcom, Quality of Service for WLR, MPF and GEA, 2017 page 207

Conclusion

Overall, the greater granularity of the Openreach Allocation model comes closer to capturing the reality of resourcing and task allocation, but the glass ceiling inputs drives the majority of the difference in the results

Despite a similar underlying approach, the two models generate very different outputs: the Analysys Mason model estimates an increase in resources of 4.7% is required to achieve the 93% performance target, while Openreach estimates that up to 25% more resources may be needed to achieve 90% performance, while 93% may not be attainable in all areas.

This review of the Openreach and Analysys Mason models identified a number of differences in methodology and assumptions that explain the difference in model outputs. While a number of factors contribute to the difference, the assumptions made regarding the glass ceiling and the failure rate and the number of visits required are estimated to be the most important factor. Given the importance of this assumption and the differing views on the extent to which the glass ceiling may be increased, this requires further consideration.

Area	Key finding	Contribution to difference in model outputs
Glass ceiling	Assuming a lower failure rate increases the maximum performance possible. Analysys Mason acknowledge that it was initially necessary to reduce fail rates in order to meet Ofcom's requested target.	**High**
Engineer skill mix	The Analysys Mason model does not capture the reality of skill mix, but this does not have a significant impact on the additional resources needed to improve performance.	**Low**
Resource variability	The Analysys Mason approach would tend to reduce variance in resources and therefore reduce costs of performance. However, the impact on outputs of this assumption is moderate.	**Low**
Task and travel time variability	While the Analysys Mason model oversimplifies the variation in task and travel times, sensitivity testing indicates that this approach may in fact slightly overstate performance costs.	**None**
Task volume variability	Neither model allows for task assignment at the PWA level; this suggests that both may understate the costs of performance.	**None**
Patch loans	While Openreach may overstate costs by not modelling patch loans, Analysys Mason may overstate the frequency and ease of these loans.	**Medium**

Introduction

Introduction

This draft report presents the findings of a review of the differences between the Ofcom Resource and Performance model and the Openreach Allocation model

Since the 2014 Fixed Access Market Review (FAMR), Ofcom has imposed Quality of Service (QoS) remedies on Openreach, articulated in two steps.

First, Ofcom assigns each regulated product a time limit by which Openreach is required to complete its repairs or offer a provision appointment. This time limit was set at 1 to 2 days after the fault for repairs, depending on the Service Level Agreement (SLA): Care Level 2 (CL2) requires repair the following day and Care Level 1 (CL1) within 2 days. Provisions of new connections should be scheduled within 12 days and delivered during the appointment.

The regulator then prescribes Openreach a performance target in terms of the proportion of repairs and provisions which are to be completed by the set deadlines. For 2016-17, this performance target was established at 80% of repairs and 90% of provisions.

Ofcom is now considering increasing Openreach QoS standards (see Table 4 and 5). The suggested new standards are informed by a model of task allocation and delivery developed by Analysys Mason for Ofcom; Openreach have also developed their own model with EY to assess the resource requirements necessary to deliver these standards.

The two models estimate materially different costs of delivering improved performance (see Table 6).

In order to help understand these differences, Openreach has commissioned Deloitte to conduct an independent review of the two models. This review is to consider the two models methodology, inputs, assumptions and testing.

The findings of this review are presented in this draft report.

Table 4: Suggested new provision standards[1]

Provision standards	2016-17	2017-18	2018-19	2020-21
% of completed by committed date (force majeure)	90% (89%)	92% (91%)	92% (91%)	95% (94%)
Working days within which first date offered	12	12	12	10
% with regulated appointment date (force majeure)	80% (79%)	90% (89%)	90% (89%)	90% (89%)

Table 5: Suggested new repairs standards[2]

Repair standards	2016-17	2017-18	2018-19	2020-21
% of repairs to be completed within SLA timescales (force majeure)	80% *(77%)*	83% *(80%)*	90% *(87%)*	93% *(90%)*
% of repairs to be completed within SLA +5 days	N/A	95%	96%	97%

Table 6: Cost of performance (repairs and provision)[3]

Model output	Ofcom*	Openreach
FTE increase required to achieve 95% provision and 90% repair performance	3.8% - 6.8%	24.9%
FTE increase required to achieve 93% performance	4.7% - 8.8%	47.6%

* Ofcom's 93% repair performance scenario is based on actual care level mix and leaves care level 2 performance at 90%. Its 90% repair performance scenario is based on a 50-50 care level mix.

1- Ofcom, Quality of Service for WLR, MPF and GEA, 2017
2- Ibid.
3- Ofcom, Quality of Service for WLR, MPF and GEA, 2017

Scope of work

The review considered the methodology, input, assumptions and testing of the two models

This independent review of the two models considered the following aspects of the Ofcom and Openreach modelling:

- **Methodology:** The model specification documents were closely analysed to assess the models' overall structure, approach and limitations. Further, the review also relied on other presentations and reports in assessing specific assumptions; the list of these documents is available in the next slide. This analysis highlighted the points of discussion explored in this report.
- **Inputs:** The analysis reviewed the inputs feeding into the models. This input analysis highlighted some additional differences between the two models, such as the on-the-day failure rate inputs.
- **Assumptions:** The review considered the code of the models for a more detailed examination of the model assumptions. This analysis also included a consideration of model calibration.
- **Testing:** In order to test the impact of these differing assumptions and approaches on the model results, existing sensitivity testing has been revised and additional scenario testing has been conducted both within the Analysys Mason model and the Openreach model.

The scope of this review has been limited by time, resource and the information granted to us. As such, the analysis should be viewed with the following limitations:

- **No audit of the models calculations:** The audit of the model calculations and workings was outside the scope of this review. This work focusses on the methodology and overall approach of the models. At the time of the report, a full audit of the Openreach model was being undertaken by Analysys Mason.
- **Reliance on Openreach's data and estimates:** Similarly, reviewing the quality of the data upon which the models inputs are based was outside the scope of this review. This review did not seek to assess the accuracy or overall reasonableness of the data or scenario outputs which were provided to us during this review.
- **Replication of model build and/or outputs:** Further, this review did not seek to modify or update either the Ofcom or the Openreach model nor to benchmark them against independent modelling.

Information and data available

This review relied on presentations, reports, scenario output and data provided to us by Analysys Mason and Openreach

As part of this work, we requested and received documents, presentation and spreadsheets from Openreach, Analysys Mason and Ofcom. The review considered the following versions of the models, shared with us by EY and Analysys Mason, respectively:

- **Allocation model 17032017**
- **Resource performance Model (Stakeholders)**

Table 7: List of datasets used in this review

Filename	Source	Description
Loans in out hours by SOM 1516.xlsx	Openreach	Patch loan daily data for 2016-17, by emitting and receiving SOM
Provision Glass Ceiling 1516.xlsx	Openreach	Data on provision major on-the-day failures
TT drag from loans.xslx	Openreach	Data on the task time drag for loaned resource
AM_JOBDATA_WITH_PWA_V3.csv	Openreach	Job volume data by product care level and PWA
Task times.zip	Openreach	Task and travel time data by product and SOM, including adjustments and furthers
NUM_ENGS_BY_WEEK.csv	Openreach	Data on the number of engineers, by skill class SOM and week, for 2015-16
SHRINKAGE_BY_WEEK.csv	Openreach	Data on the shrinkage by week and SOM, for 2015-16
TBL_AM_ASSIST_DISTNS.csv	Openreach	Data on the proportion of skilled and unskilled assists
Cost of delivering services (…).xlsx	Openreach	Data on the cost of engineers upskilling and hiring
Stuart – Glass ceiling analysis.xlsx	Openreach	Data on the repair on-the-day failure rate
Calibration Approach for Allocation Model v3	Openreach	Description of Openreach's approach in calibrating the Allocation model
Openreach Service Demand Modelling	EY	Model specification document of the 2017 Allocation model
Openreach Discrete Event Simulation Model	EY	Model specification document of the 2013 Distribution model
Overview of the Quality of Service Model and its output for WLR/LLU Charge Control 2017	Analysys Mason	Model specification document of the 2017 Resource and Performance model
SLA and SLG comparisons for WLR and LLU MPF	Analysys Mason	Report benchmarking internationally Openreach's Service Levels Agreement (SLAs) and Service Level Guarantees (SLGs).
Quality of service for WLR, MPF and GEA	Ofcom	Consultation document
Clarifications to Stakeholder	Ofcom	Response to queries about the Resource and Performance model

Model Specifications

Overview of model methodology

Both models use a queuing simulation approach in which tasks are logged and allocated to available resources

Model objectives

The objective of the modelling is to estimate the additional resources that are required in order to deliver the specified level of service. This determines the costs associated with moving to the new MSLs.

Both models simulate performance using a queueing model of task arrival, allocation and completion.

Model methodology

Inputs: In these models, the pool of engineers available to complete tasks is assumed to be pre-determined. The models then simulate the arrival of tasks – repairs and provisions – that vary in their complexity and the level of prioritisation (based on the care level).

Within the Analysys Mason model resources can be transferred across SOMs through patch loans, whereas the Openreach model does not endogenise this in the model but captures resource variability through the shrinkage assumption.

Queuing: As new tasks arrive, they are placed into a queue and then assigned to available engineers for completion; the model design specifies how tasks are assigned and prioritised.

Completion: Once assigned to an engineer, the model simulates the processing and completion of these tasks, taking into account factors such as the probability of additional resources being required, the engineer failing in their assigned task and/or variation in the time to complete. If not successful, tasks may be returned to the queue to be reassigned.

A glass ceiling limits the performance that can be achieved: the Analysys Mason model captures this directly through failure rate assumptions and the associated delay; this is captured more endogenously in the Openreach model where minor on-the-day failure delays depend on resource availability, but major issues (eg the need for civil engineering result in an automatic fail).

Model outputs

Outputs: Once successfully completed, the model records the time taken to complete the task and whether it was compliant with the specified care level.

The output of the models is the percentage of repairs and completions that are delivered within the specified time, for a given level of resources. In order to estimate the number of resources required to deliver a specified level of service, the models can be run a number of times with different levels of resources to find the relationship between performance and resourcing.

Model specifications: Input data

The models use the same data on tax volumes but the Analysys Mason model uses less detailed data on task duration, travel times and resources

The two models inputs differ in three key ways:

- **Historical data versus probability distribution:** Most of the inputs of the Openreach model are drawn from probability distributions based on historical data whilst the Analysys Mason model is directly based on historical data.
- **Input granularity:** The Analysys Mason model inputs tend to be less granular than the Openreach model input. For example, the repairs data does not distinguish across product types (beyond the SLA), task duration or travel time and the resource level is calibrated in the model but does not account for weekly fluctuations or the mix of skills.
- **Parameters:** The Analysys Mason model allows the user more freedom in choosing the value of certain key parameters, such as the on-the-day failure or no-show rate. In the Openreach model, these are also drawn from probability distributions based on historical data.

Openreach Allocation Model

- **Repairs:** Repairs are logged at SOM level based on actual time. Repairs volume, log time, product type and fault location are drawn from probability distributions derived from historical data. The model includes Care Levels 1 to 4, non-MSL and broadband jobs, and accounts for assists.
- **Provisions:** Provision volumes, log times and product types are also drawn from probability distributions based on historical data. The model includes Fibre-to-the-cabinet (FTTC) provisions.
- **Resourcing:** The user inputs the number of engineers by SOM and skill types, with 4 different skill categories allowed. Further, this level of resource varies based on shrinkage, attrition and hiring; based on historic data.
- **On-the day failures and no-shows:** On-the-day failures are simulated through a 'number of visits' statistic which is drawn from a probability distribution based on historical data. Similarly, the rate of no-shows is based on historical data.

Ofcom Resource and Allocation model

- **Repairs:** Repairs are logged at SOM level based on an approximate time, with days divided into 5 periods. Repairs volume and care levels are based on historical data for 2011-12 to 2015-16 at SOM level. The model only includes repairs of Care Levels 1 to 2.
- **Provisions:** Provision volumes and approximated log times are also based on UK level historical data and allocated to SOMs based on number of lines and faults in 2011-12. The model includes FTTC provisions.
- **Resourcing:** The number of FTEs is calibrated in the baseline model and assumed flat across the year. The model then allocates these FTEs to SOMs based on the number of lines and faults in 2011-12. Engineers are assumed to be multiskilled.
- **On-the day failures and no-shows:** The on-the-day failure rate and no-show rate is chosen by the user.

Model specification: Assumptions

The Analysys Mason model is based on some additional simplifying assumptions regarding variability and task completion times.

The Analysys Mason model is a more high-level simulation of Openreach activities based on some simplifying assumptions, such as engineers multi-skilling, an exogenous glass ceiling, and the smoothing out of variability in some parameters such as task time and travel time. However, the version of the Openreach model this review had access to did not include any patch loan feature, but assumed resources were limited to work within one SOM.

Openreach Allocation Model

- **Queuing:** All jobs enter a common workstack once logged. The model prioritises jobs based on the remaining time before SLA expiration. The model tracks jobs from log time to completion and measures actual completion times.
- **Allocation:** Jobs are prioritised according to the following order: Care Levels 3 and 4, appointed jobs, due jobs and jobs not yet due. Tail jobs' priority is based on their delay and a cut-off time based on historical data.
- **Completion:** Jobs are allocated to engineers based on their skill requirements. Travel and task times are drawn from probabilistic distributions which depend on the fault type and SOM. Completion may be interrupted if task time is greater than the engineer's work shift.
- **Loans:** No loans are possible in the version of the model shared with us. In another version, loans were simulated through the calibration of the shrinkage.
- **Glass ceiling:** The glass ceiling is mainly driven by the proportion of jobs inputted as Care Level 6 to 10, which the model automatically considers as fails since more complex work is required. This captures the major on-the-day failures (e.g. civil engineering, complex faults). Further, the model assigns to each job a number of visits statistic, which determines the number of visits necessary for a jobs to be completed. This may also result in a SLA failure, although some of the Care Level 1 and Care Level 2 jobs may still be completed within SLA deadlines even if they require more than one visit.

Ofcom Resource and Performance model

- **Queuing:** Provision, Care Level 1 and Care Level 2 jobs from the fault and provisioning data are each logged into a different queue. The model tracks jobs from log time to completion, and measures actual completion times.
- **Allocation:** The model is designed to balance performance across all job types. Resources are pro-rated by the length of the queue in each workstack, divided by the target days for each type of job. Analysys Mason however recognises that the model may not always achieve an optimal allocation of resource between provision and repair jobs.
- **Completion:** All engineers are assumed to be able to complete any job. All jobs' completion time is set to 2.5 hours, so that engineers complete 4 jobs a day.
- **Loans:** Resources may be loaned across SOMs, at 75% productivity. The model generates loans based on a stress factor.
- **Glass ceiling:** The glass ceiling is driven by repairs and provisions minor and major on-the-day failure rates. Minor failures delay completion by 1 day while major failures delay it by 5 days.

Model specification: Outputs

The two models estimate significantly different costs of performance

In the Ofcom base scenario,[1] reducing the target days for provision appointments from 12 to 10 days would require a 0.6% increase in FTEs. Increasing the performance target 95%, 93% and 90% for provision, Care Level 1 and 2, respectively, would require a 4.7% increase in FTEs.

The Openreach model points at significantly higher results. A 93% repair and 95% provision performance for 10 days FAD is estimated to require an increase of close to 50%in FTEs relative to 80% repair 90% provision and 12 days FAD performance.

Overall, both models' cost performance curve suggests that the incremental cost of performance increases sharply as performance approaches the glass ceiling (see Graph 1). The Openreach model suggests a lower glass ceiling than what is assumed in the Analysys Mason model.

Graph 1: Performance versus FTEs in the Analysys Mason model, base care level mix[1]

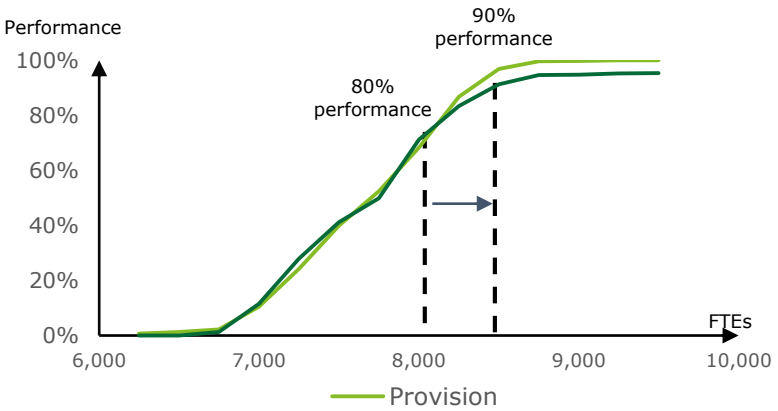

Model calibration: Both models were calibrated to adjust their output to match Openreach historical performance. In the Analysis Mason model, the calibration was based on:

- Adjustments of the total number of FTEs. The number of FTE used in the final model is significantly below actuals.
- Modification of the prioritisation rule through 'bias' factors.

The Openreach model's calibration was based on:

- Improved granularity of the job input data (breakdown between MSL and non MSL, SIM provide jobs).
- Improved job prioritisation rule (better performance balance, stress factor).
- Smoothing of task times and FTE variations by SOM.
- Introduction of the glass ceiling on performance through automatically failed faults class.

Table 8: Cost of performance (repairs and provision)[2]

Model output	Ofcom	Openreach
FTE increase required to achieve 95% provision and 90% performance	3.8% - 6.8%	24.9%
FTE increase required to achieve 93% performance	4.7% - 8.8%	47.6%

* Ofcom's 93% repair performance scenario is based on actual care level mix and leaves care level 2 performance at 90%. Its 90% repair performance scenario is based on a 50-50 care level mix.

1- Analysys Mason, Overview of the Quality of Service Model and its outputs, 2017
2- Ofcom, Quality of Service for WLR, MPF and GEA, 2017 and Openreach data.

Model specification: Sensitivity testing

The Analysys Mason model outputs are sensitive to the on-the-day failure rate and care level mix

The sensitivity testing scenarios ran by Analysys Mason suggested that the Analysys Mason model is sensitive to some of its assumptions, notably its on-the-day failure parameter. Indeed, the consultants find that under a 5% repair minor fail rate, the required increase in FTEs for an increase in CL1, CL2 and provision performance to 93%, 90% and 95% rising to 8.8% from 4.7%. Further, the model was also found to be sensitive to the care level mix, with the cost of performance increasing to 8.0% in a scenario in which all tasks are CL2.

Openreach tested the model's sensitivity to the patch loans assumption through a calibration of the shrinkage rate and found this assumption to have little impact on the model output.

Table 9: Sensitivity to fault rates and Fibre to the Cabinet (FTTC) provision (% increase in FTEs) in the Analysys Mason model[1]

Scenarios	Provision 90%, CL1 and CL2 80%	Provision 95%, CL1 93%, CL2 90%
Base parameters	0.6%	4.7%
10% reduction in FTTC provision	0.7%	5.0%
20% reduction in FTTC provision	0.5%	4.8%
10% reduction in fault volumes	0.8%	4.8%
5% repair minor fail rate	1.1%	8.8%

Table 10: % increase in FTE to reach performance target, under different care level mix, relative to 80% repairs and provision target and 12 days FAD[2]

Scenarios	Provision 90%, CL1 and CL2 80%, 10 day FAD	Provision 95%, CL1 93%, and CL2 90%, 10 day FAD
Base care mix	0.6%	4.7%
100% CL1, 0% CL2	0.0%	2.2%
60% CL1, 40% CL2	0.4%	4.8%
50% CL1, 50% CL2	0.4%	5.8%
40% CL1, 60% CL2	0.4%	5.9%
0% CL1, 100% CL2	2.4%	8.0%

1- Analysys Mason, Overview of the Quality of Service Model and its outputs, 2017
2- Ibid.

Points of difference

Points of difference

The comparison of the two models identified key points of difference which may impact the models' estimated cost of performance.

The review considered the methodology, input, assumptions and testing of the two models. As shown in the above slides, while the underlying methodology of the two models is similar, there are a number of areas in which they differ in the assumptions made and the level of granularity of the analysis.

Methodology: The underlying methodology is similar in terms of the allocation of tasks to the available pool of resources. However, the Analysys Mason model has less functionality to capture variance in task duration, skill requirements, travel times and staff availability.

The models also differ in their treatment of resources: the Analysys Mason model endogenously captures the transfer of resources across SOMs whereas the Openreach Allocation model captures this via shrinkages, without having it respond to demand. The models also differ in the extent to which the glass ceiling is endogenised within the model (and in the specific assumptions made).

Inputs: The Openreach model uses historic data as the basis for task volumes and complexity, task duration and travel times. The model also uses actual data on resource availability to capture fluctuations in engineer availability and the skill mix.

While Analysys Mason use actual task volume data, the functionality of the model requires a constant task time, workforce size and skill mix to be imposed. These inputs are not based on actual data but are calibrated to match Openreach historic performance.

Assumptions: As noted above, the Analysys Mason model makes a number of simplifying assumptions that tend to smooth out the variance in the model.

The two models also differ in the approach taken to failure rates and the glass ceiling: the Openreach Allocation model is based on historic on-the-day failures whereas Analysys Mason make assumptions about the probability of minor and major fails and the associated delay.

This section of the report focusses on the key points of difference between the models and the implications of these differences for the model results. The following slides provide detail on:

- **The role these differences play in the model;**
- **How the approaches align with available data; and**
- **The potential implications for the model outputs of differences in approach and assumptions.**

Points of difference

The following slides investigate these key points of difference and the implications for model outputs

The table below summarises the key points of difference from the previous slide, which are considered in detail across the subsequent slides.

Type	Area	Key finding	Testing conducted
Assumption	Glass ceiling	The Analysys Mason model uses a lower on-the-day failure rate assumption than what is supported by current Ofcom or Openreach data. The model also does not allow for the additional costs associated with the need for multiple visits or assists.	Review of Analysys Mason sensitivity testing on repair on-the-day failure rates and further testing on provision on-the-day failures.
Methodology/ inputs	Engineer skill mix	The Analysys Mason model assumes all engineers can complete any task. In practice, some assignments require underground access and more specialised engineers.	Sensitivity testing based on increasing UG's share of total engineer workforce to 85% within the Openreach model.
Methodology/ inputs	Resource variability	The Analysys Mason model assumes resourcing is constant within each SOM; in practice, resourcing varies significantly with training, holidays and sickness.	Sensitivity testing based on a making shrinkage constant in the Openreach model.
Methodology/ inputs	Task and travel time variability	The Analysys Mason model assumes a constant task time and does not account for variation in travel time. In practice, there is significant variance and the distribution is highly skewed, with a long tail of more complex and time-consuming tasks.	Sensitivity testing based on assuming a constant task completion time within the Openreach model.
Methodology	Task volume variability	Both the Analysys Mason and Openreach models simulate performance at Senior Operating Manager (SOM) level. In practice, engineers work within Preferred Working Areas (PWAs) that are much smaller than SOMs.	Not tested as neither model has the capability to forecast performance at the PWA level.
Methodology	Patch loans	The Analysys Mason model allows for resources to be loaned across SOMs in order to meet targets; the Openreach model does not incorporate these loans, but uses shrinkage data to capture fluctuating resources.	Sensitivity testing based on switching off the patch loans functionality in the Analysys Mason model.

Glass ceiling

The glass ceiling assumptions determine maximal performance within the model, given that some tasks will not be successfully completed

In order to capture the fact that not all repairs and installations will be successful on the first attempt, both models introduce a glass ceiling on performance. This is represented through a certain probability that a task results in failure, thus requiring a subsequent visit or an additional engineer.

- In the Analysys Mason model, this glass ceiling is reflected through the proportion of tasks that incur either a minor fail – which can be addressed the following day – or a major fail, which takes five days to resolve.
- In the Openreach model, this is represented through a proportion of jobs which are assumed to automatically fail (major on-the-day failures) or require additional visits (minor on-the-day failures). This is based on historical data.

A fail or assist will not automatically result in the MSL being failed in all cases since it may be possible to schedule an additional visit within the specified time.

- Within the Analysys Mason model, major fails will always result in the MSL being missed, since a 5 day delay is assumed.
- Provision minor fails will also fail the MSL, which requires on-the-day completion.
- Repair minor fails may be successfully completed, if not originally attempted on the day of the deadline.

The Openreach model does not impose a specific delay associated with fails; rather this is determined endogenously based on the time to schedule an engineer with the required skills to do the additional visits.

The following slides set out how the glass ceiling is modelled, the assumptions made and implications for the results, and the competing views on the extent to which the glass ceiling may be raised.

Impact of the glass ceiling assumption

The fact that some tasks will not be successful therefore imposes an upper bound on performance. For example, if 5% of tasks face a major fail (such as a need for civil engineering) that results in the MSL being missed, then performance against the MSL cannot exceed 95%, regardless of the level of resourcing.

In addition, for levels of performance close to the glass ceiling, the resources required to deliver a marginal improvement in performance will increase rapidly. This effectively shifts the curve capturing the relationship between resources and performance to the left.

Analysys Mason view on the glass ceiling

The role played by the glass ceiling assumption in the model is noted by Analysys Mason, who acknowledge that the 95% target for provisioning *"is only achievable if fails do not exceed 5%."*[1]

In addition to placing an absolute upper bound on performance, Analysys Mason also note the effect that the glass ceiling has on resourcing costs: *"as the target approaches the "glass ceiling" there is less and less ability for any other job to miss the target in order for the required performance level to be achieved, which leads to very large increases in the required number of resources."*[2]

1- Analysys Mason, Overview of the Quality of Service model and its outputs, 2017 page 20

2- Ibid.

Glass ceiling

The Analysys Mason model assumes a reduced minor fail rate in order to ensure that the targets specified by Ofcom are achievable

The repair glass ceiling assumptions made in the model differ significantly from the raw data

The table below sets Ofcom's model assumptions against the historical fault rates under Analysys Mason's definition and Openreach's definition.

Table 11: Repairs on-the-day failures in FY 2017

Fault rates	Provision		Repairs	
	Minor	Major	Minor	Major
Analysys Mason model	3.6%	1.4%	3%	3%
Raw data (Analysys Mason's definition)	6%	1.4%	5%	3%
Raw data (Openreach's definition)	N/A	2.5%	2.6%	7.4%

The Analysys Mason model relies on on-the-day failure rates which are unsupported by the available data. As a result:

- The Analysys Mason model overestimates the feasibility of the suggested requirements. Indeed, under Analysys Mason's glass ceiling analysis, the suggested provision target would not be achievable.
- Further, under Openreach's definition of the glass ceiling, the repair target would also not be achievable.

In its clarification to stakeholders document, the regulator suggested that these modifications were motivated by Ofcom's willingness to account for potential future improvement in the glass ceiling. This argument is examined in the following slide.

Motivation for Analysys Mason assumptions

By Analysys Mason's own account, the on-the-day failure rates initially had to be reduced in the Openreach Model to ensure the feasibility of Ofcom's targets.

- "Ofcom has requested a scenario with 95% [provision] performance against this target, which is only achievable if fails do not exceed 5%. […] We have reduced the proportion of minor fails to 3.6% to ensure that Ofcom's preferred CDD scenario is possible to achieve"[1]
- "We have made a similar assumption relating to the number of minor fails. At one stage of the model construction, all minor fails would fail to meet the SLA. This meant that it was necessary to reduce the proportion of repair minor fails from 5% to 3% such that the target of 93% repair performance was possible to achieve."

The model output is particularly sensitive to this assumption, as shown in Table 12.

Table 12: Analysys Mason sensitivity testing results; required increase in FTE relatively to baseline

Minimum standard (provision/repair)	3% repair minor failure rate	5% repair minor failure rate
80% / 80%	0.6%	1.1%
90% / 93%	4.7%	8.8%

1- Analysys Mason, Overview of the Quality of Service model and its outputs, 2017 page 20
2-Ibid

Glass ceiling

The Analysys Mason and Openreach models have different on-the-day failure rates. This has implications for the feasibility of the suggested standards

A comparative analysis of Analysys Mason and Openreach input data highlights discrepancies between Ofcom's and Openreach's glass ceiling assumptions.

Analysys Mason excludes 3 faults categories from the glass ceiling which Openreach included in the Allocation model glass ceiling. These three fault types represent in total 2.13% of on-the-day minor faults. This includes:

- Assists: Tasks not completed because of the need for assistance from an additional engineer.
- Different skilled engineer: The job was not completed because the fault or provision required a different class of engineer.
- Specialist skill: The repair was not completed because the fault required specialist skills.

Further, Analysys Mason seem to consider as "Minor" many of the faults that Openreach considers as "Major". Indeed, according to Openreach's classification, 0.5% of the faults that Analysys Mason includes in the glass ceiling are minor and 7.4% major (see Table 13).

In the Analysys Mason model, major faults cause delays of 5 days. According to Openreach data, this represents 77% of the faults identified as glass ceiling by Analysys Mason. This would suggest that major on-the-day failures would represent 6% of all activity, against 2% for minor failures.

This would have implications for the feasibility of the the 97% SLA +5 days target which may not be achievable under the current major failures rate.

Table 13: Repairs on-the-day failures in FY 2017[1]

Fault	Type	Included by Ofcom?	% in FY17
Assists	Minor	Not included	0.21%
Different skilled engineer	Minor	Not included	1.86%
Hoist required	Minor	Yes	0.32%
Exchange equipment	Minor	Yes	0.14%
Customer Damage	Minor	Yes	0%
Specialist skill required	Minor	Not included	0.06%
Specialist tools/stores	Minor	Yes	0.04%
Common to complex fault	Major	Yes	1.10%
Reject clear	Major	Yes	0.73%
Long duration work	Major	Yes	0.76%
CP access - readiness - non appointed	Major	Yes	1.06%
Unexpected hoist	Major	Yes	0.17%
Engineer access obstructed	Major	Yes	0.20%
Network damage	Major	Yes	0.13%
Safety hazard	Major	Yes	0.31%
Needs civil engineering	Major	Yes	2.97%
Total, included by Ofcom			**7.93%**
Total, included by Openreach			**10.06%**

1- Openreach Glass ceiling analysis.

Glass ceiling

An Openreach review of the fault rate data indicates that even under Analysys Mason's approach, the glass ceiling could be below 93%

Openreach has suggested that it could reduce minor on-the-day failures by 1.29 percentage points relative to 2016-17 levels. This is mainly achieved through reducing on-the-day failure types which are not included in Analysys Mason's glass ceiling (1.07% out of 1.29%).

Openreach argues that it could only reduce major failures by 0.03% relative to 2016-17 levels. It argues that the level of on-the-day major failures experienced in 2016-17 was already low thanks to favourable weather conditions. Further reductions in major on-the-day failures (1.8%) would only be possible on the basis of industry wide engagements.

The glass ceiling would therefore be raised by 1.29 percentage points on Openreach's definition and 0.19 percentage points relative to Analysys Mason's definition. This could improve by a further 1.8 percentage points should Openreach's proposals on clear reject and customer access be accepted be accepted by the wider industry.

Table 14: Planned reduction in minor on-the-day failures relatively to 2016-17

Fault	Included by Ofcom?	Reduction
Assists	Not included	-0.11%
Different skilled engineer	Not included	-0.93%
Specialist skilled required	Not included	-0.03%
Hoist required	Included	-0.10%
Exchange equipment	Included	-0.07%
Specialist tools / stores	Included	-0.02%
Reject clear	Included	-0.03%
Total reduction		**-1.29%**

Graph 2: Openreach estimation of the repair glass ceiling

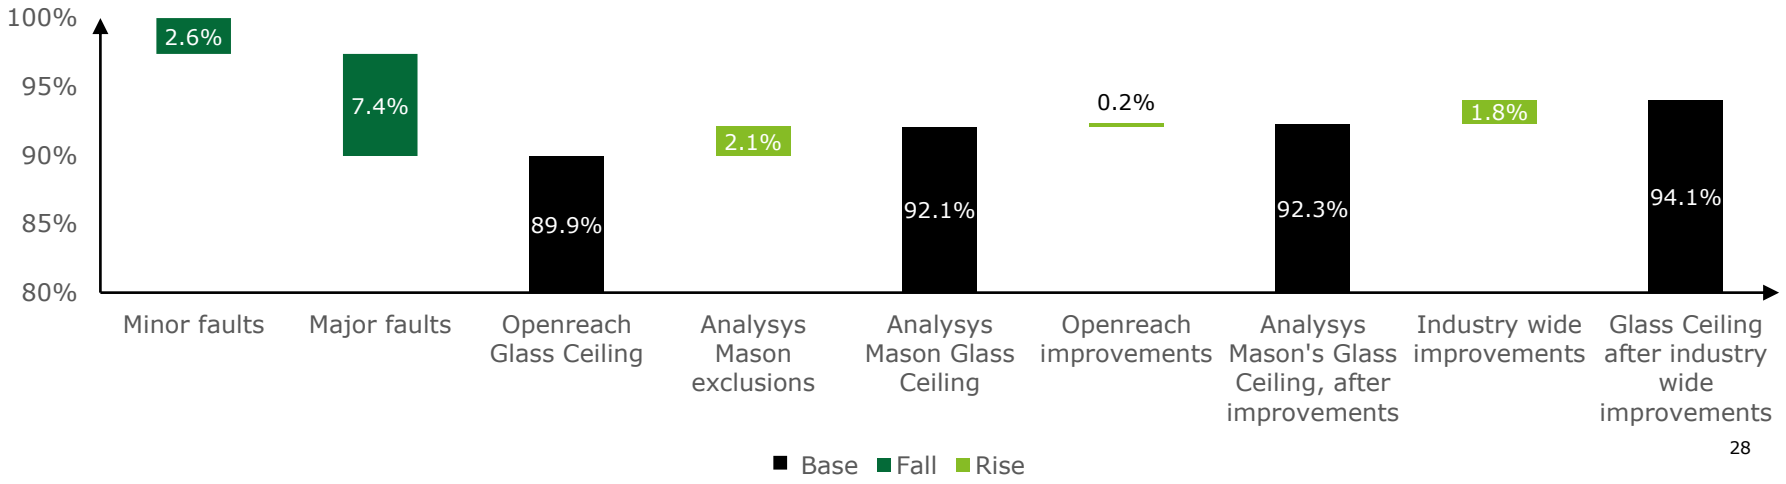

Glass ceiling

However, Ofcom argues that Openreach's SLA failure rate may be lower than the on-the-day failure rate, implying a higher glass ceiling

Ofcom argues that Openreach may be able to raise the performance significantly above what would be suggested by the 2015-16 8.6% on-the-day failure rate.

- Out of these 8.6%, Ofcom excludes 2.1% which are classified as Customer Premise Access. The regulator argues they do not count as SLA faults.
- Further, Ofcom argues that with enough resources Openreach may reattempt a job after an initial on-the-day failure and still meet the SLA. As such Ofcom considers actual SLA glass ceiling to be 95.3% and 93.8% for Care Levels 1 and 2, respectively.
- Further, Ofcom considers that there is scope to reduce the current on-the-day failure rate, notably failures due to the need for a different skill engineer (2.3%), hoist or assist (1.1%), access obstruction and safety hazard (0.4%) and 0.1% for non-standard tools.

This would feed into a final 97.5% and 95.8% glass ceiling for CL1 and CL2, for a weighted average of 96.6%. However, the costs and implications for resourcing associated with these changes are not assessed.

Openreach has identified three challenges to Ofcom's view on the SLA Glass Ceiling.

- Openreach is able to restart the SLA clock for Customer Premise Access only in the case of appointments. As such, a significant share of the 2.1% would remain part of the SLA glass ceiling.
- The categorisation of the nature of the on-the-day failure should not overemphasise the first visit. In practice, there can be multiple events which lead to the failure. Openreach revisited its glass ceiling data on this basis.
- Openreach identified new failure scenarios such as clear rejects, long duration visits, common faults and network damage

Graph 3: Ofcom's estimation of the repair glass ceiling (CL2)

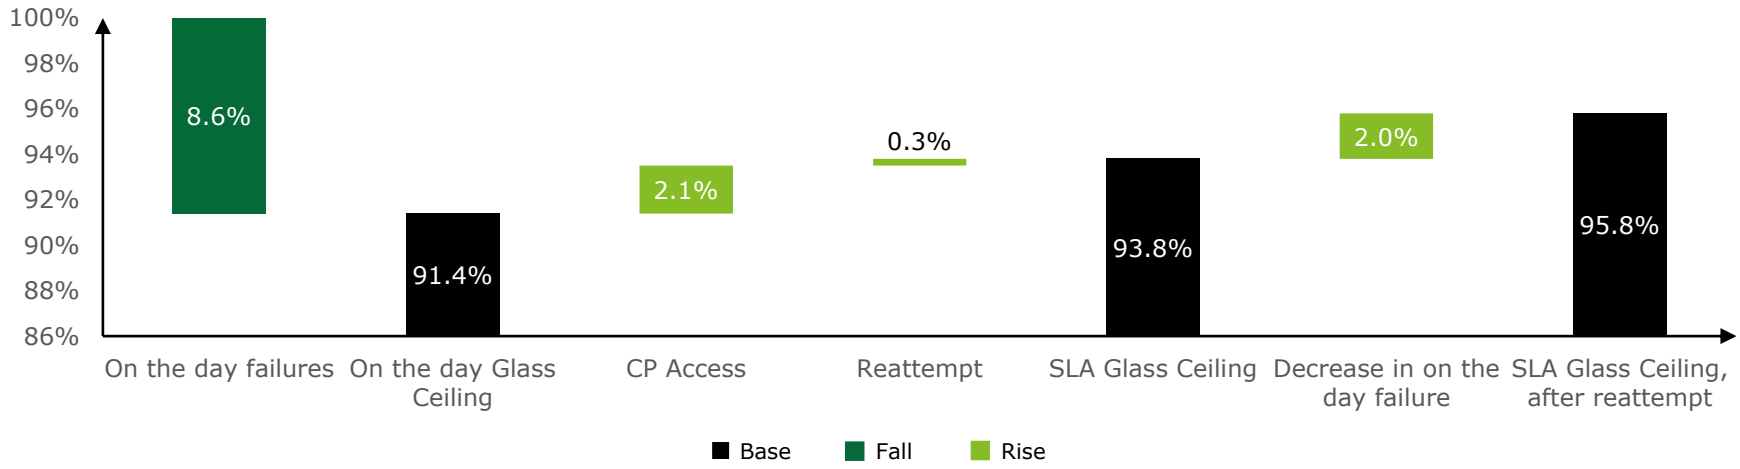

Engineering skill requirements

The Analysys Mason model is based on the assumption that tasks can be assigned to any engineer, rather than some requiring specialised skills

The Analysys Mason model assumes that every Openreach engineers can complete any provision or repair job and so new tasks can be allocated to any available engineer.

In practice, Openreach's engineer workforce is divided into a hierarchy of four types of engineers, as shown in Table 15. Within this hierarchy, lower-skilled engineers can only deliver more basic tasks while more highly skilled engineers can complete both basic and more complex tasks.

This is reflected in the Openreach model by having tasks assigned to the available engineer with the minimal skills required for completion; this is intended to leave more highly-skilled engineers available to complete the tasks that require specialist skills.

The fact that certain tasks require more specialised skills to be completed effectively adds an additional constraint in the model and limits the extent to which Openreach can manage the queue of tasks by pooling engineer resources.

While this constraint is partly mitigated by the fact that Openreach's engineer population is considerably more skilled than its average task distribution, variance in task arrival and the locality of underground work could potentially make this constraint binding.

As a consequence, the universal multiskilling assumption understates the constraints Openreach faces in allocating engineers to jobs.

Table 15: Openreach engineer skill classes[1]

Engineer skill class	Installation	% of engineers
One Man Installer (OMI)	Provision only	9%
Customer Access Line (CAL)	All the above and basic repair jobs	27%
Underground (UG)	All the above and underground repair jobs	33%
Broadband Underground (BBUG)	All the above and broadband jobs.	31%

Graph 4: Openreach engineers distribution and task skill requirement distribution (exc. BBUG), 2015-16[2]

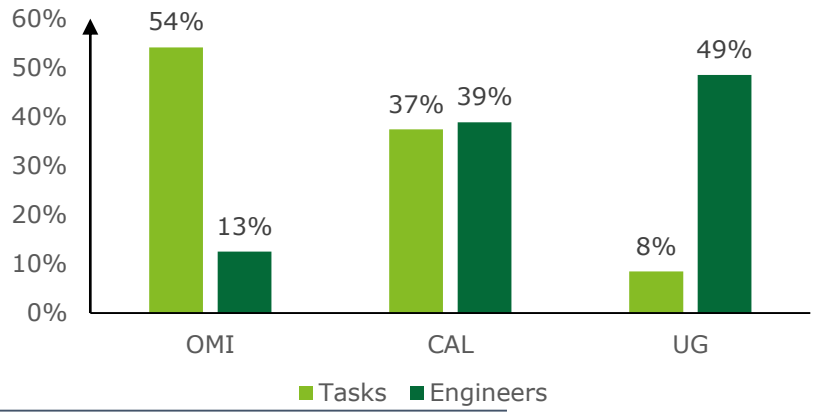

1- Deloitte analysis based on Openreach data.
2- Ibid.

Engineering skill requirements

Sensitivity analysis within the Openreach Allocation model shows that assuming multi-skilling makes it easier to reach performance targets.

In order to test the implications of the assumption of universal multi-skilling, Openreach ran a modified version of the Openreach model in which all CAL and OMI engineers are upskilled to UG.

In this scenario, for a constant number of FTEs, baseline performance increases by 3 to 4 percentage points for repairs and 1 percentage point for provision. Further, this assumption decreases the additional cost of meeting the performance target, albeit by a small amount.

It is recognised that Openreach could take additional steps to upskill their workforce and increase the proportion of multi-skilled engineers. This suggests that in the future the Analysys Mason assumption that all engineers can complete any task may be more appropriate. This limitation was noted by Analysys Mason in their model specification:

"In reality some technicians will not be capable of carrying out all types of job and this will place some additional constraint on Openreach's resourcing. As a result the model may underestimate the number of resources required to achieve a given performance level against targets."[1]

However, increasing the skill level across the workforce will come at additional cost, estimated by Openreach at £42.3 million over the next five years. To provide context for this figure, this is equivalent to increasing resources by about 1.2% relatively to the total expected labour costs over these five years. If this upskilling assumption is a key driver of the estimated resourcing needs, then these additional costs should also be taken into account.

Graph 5: Performance under actual mix skill and upskilling scenario.[3]

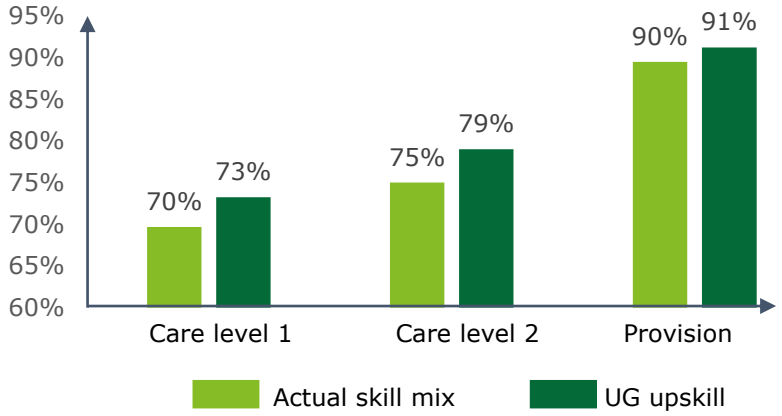

Table 16: Required increase in FTE to meet performance targets relatively to baseline, by skilling assumption[2]

95% provision and 93% repair performance	Actual	UG multi skilling	Difference
2018-19	0.7%	0.2%	-0.5%
2019-20	9.3%	8.3%	-1.0%
2020-21	23.1%	22.9%	-0.2%

1- Analysys Mason, Overview of the Quality of Service model and its outputs, 2017 page 24 3- Source: Openreach
2- Source: Openreach. Note that the cost of performance is here reported post MBORC, as such they cannot be compared against pre MBORC numbers.

Variation in tasks and resources

Variance in task volumes, completion times and resource availability is a key driver of performance in queuing models

The ultimate determinants of performance in queueing models are the availability of resources relative to the flow of tasks; however, the variability of the volume of tasks, the duration of tasks and the resources available are also key factors.

This can clearly be seen by considering a simplified case in which the volume of tasks, their duration and the workforce are constant. In this case, it is straightforward to forecast the resources available and schedule them at 100% utilisation and efficiency. However, as the variance of task volumes change – even if the average volume is the same – then additional resources will be required in order to address spikes in demand. If the level of resources does not increase, then spikes in demand will not be addressed and will add to a backlog, having a persistent long-run impact on performance.

The greater the level of variability in performance, the more resources are needed to ensure that spikes in resources can be managed, and the greater the risk of under-utilisation in periods of low demand (as shown in the figure opposite). This therefore increases resource costs, even if the average flow of tasks remains the same.

A similar argument applies to variations in task times and travel times and variation in resource availability, since all of these factors have the effect of increasing the variance around the number of tasks that can be completed within a day. While both models allow for variance in the volume of tasks, they differ in the extent to which task times and resource variability are captured.

The following slides discuss how assumptions about variance impact the model:

- Resource availability: accounting for variation in staff available.
- Task time variability: accounting for differences in task duration.
- Travel time variability: accounting for variation in travel times, which increase the overall time taken to complete a task.
- Task volume variability: while this is reflected in both models, the task flows are aggregated at the SOM level, which may obscure some of the practical variation at the ground level.

Graph 6: Illustration of the impact of variance on performance

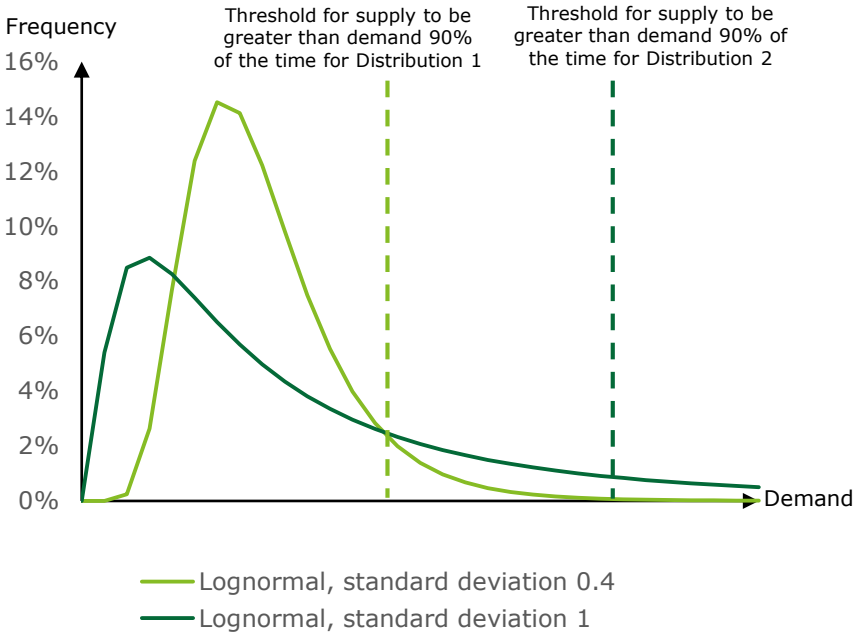

Variation in tasks and resources: staffing levels

The Analysys Mason model assumes a constant resourcing level. Labour force variability may increase the cost of delivering higher performance

The Analysys Mason model assumes that Openreach's total resourcing is constant across the year. In practice, Openreach faces both attrition and shrinkage (for example holidays, training and sickness) which may temporarily reduce the number of engineers available in a certain patch and therefore affect job completion.

While the variance in the total number of FTEs tends to be relatively small within each SOM, the variance in shrinkage rates significantly impacts the resources available (see Graph 7).

Resourcing variability may increase the cost of performance as it may require Openreach to have extra capacity so as to face any unexpected decrease in resourcing. In its consultation document, Ofcom recognises that even a small variation in resourcing can have a significant impact on performance.

Graph 7: Weekly shrinkage rate, SOM 1511, FY 2015-16[1]

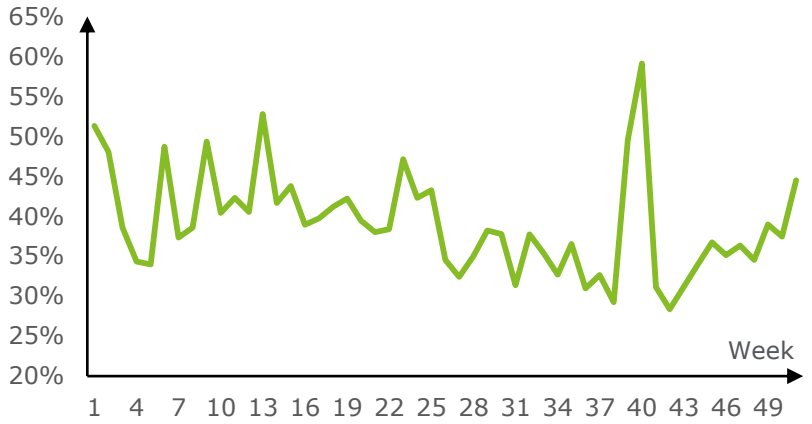

This assumption was tested through a modification to EY's model to set shrinkage to a constant rate. The sensitivity scenario suggests that assuming no variation in resourcing will tend to lead to underestimates of the additional resources required to achieve high standards of performance.

However, it is noted that this effect does not appear to hold for the Year 1 requirements, which specify a lower service requirement. This may reflect the fact that there is some correlation between shrinkage and task volumes in the model, which relies on actuals for both (for example, training being scheduled when demand is lower). As performance requirements increase and approach the glass ceiling the impact of increased variance may outweigh the benefits of better resource scheduling.

Table 17: Required increase in FTE to meet performance targets relative to baseline, by shrinkage assumption[2]

95% provision and 93% repair performance	Actual	Flat shrinkage	Difference
Year 1	0.7%	1.1%	0.4%
Year 2	9.3%	9.1%	-0.2%
Year 3	23.1%	21.7%	-1.4%

1- Deloitte analysis on Openreach data.
2- Source: Openreach. Note that the cost of performance is here reported post MBORC, as such they cannot be compared against pre MBORC numbers.

Variation in tasks and resources: task and travel times

The Analysys Mason model assumes a constant task and travel time, whereas Openreach data shows a long tail of lengthy tasks

The Analysys Mason model assumes a constant task and travel time of 2.5 hours, so that the number of jobs completed daily equals four times the number of FTEs.

In reality however, the task time distribution is significantly skewed (see Graph 11), implying that there is a significant probability of certain jobs taking two or three times longer than usual. Further, average travel time varies significantly across General Manager areas, with the average travel time in Scotland (GM1) being twice as long as that in London (GM9).

As noted above, higher levels of variance will tend to make it more challenging to manage the queue of tasks, thus increasing the resources needed to deliver specified standards. The regulator also acknowledged the potential impact of supply variability:

"In practice job durations will vary. Not taking this into account could overlook variations that could on some days have a material impact on the overall number of jobs that can be completed during the day."

On the other hand, the skewness of the distribution means that the majority of tasks take less than the mean time to complete.

As such, the omission of task time variability may have an ambiguous overall impact, with the effect varying significantly across regions. The Analysys Mason model may underestimate the cost of performance in GMs with the longest travel time, with implications for the cost of meeting the performance target in each of the 9 GMs.

Graph 8: Task time distribution for the provision of a WLR-MSL line in the SOM BVH33, in minutes[2]

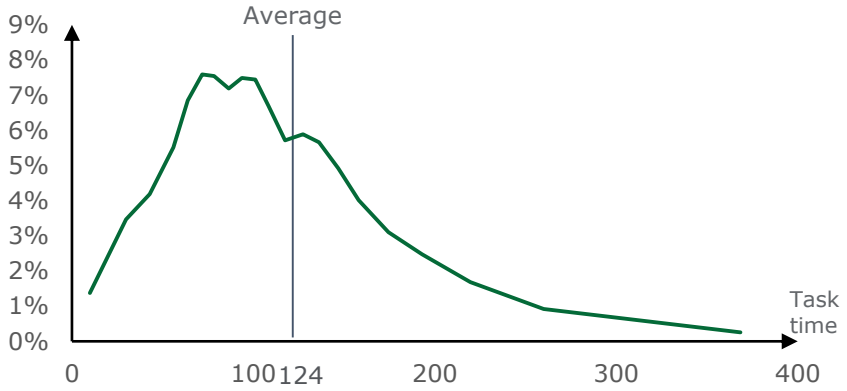

Graph 9: Travel time distribution in SOM BVG11, in minutes[2]

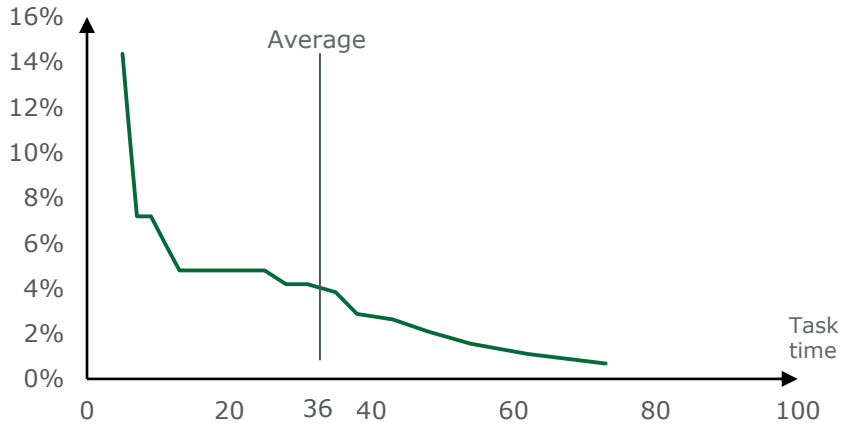

1- Deloitte analysis based on Openreach data. Note that the task time distribution is here smoothed and approximated; the actual distribution used in the model is discrete.
2- Deloitte analysis based on Openreach data. Note that the travel time distribution is here smoothed and approximated; the actual distribution used in the model is discrete.

Variation in tasks and resources: task and travel times

Scenario testing shows that task time alone does not increase the overall cost of achieving the MSL, but impacts may vary by region

To test the impact of task and travel time variability on performance, the Openreach model was run under a constant task and travel time assumption.

Overall, this suggests that task and travel time variability may decrease the cost of achieving a higher level of performance by up to 2%. As suggested on the previous slide, this decrease in performance costs may be a function of the skewness of the distribution and the larger number of shorter than average tasks.

However, the scenario results at GM level show that the impact of this flat task and travel time assumption varies significantly across regions. For example, the estimated cost of achieving higher performance appears to be under-estimated in the South of England (GM10) and Scotland (GM1) within the Analysys Mason model.

These results do not show a clear pattern across GMs.

- For example, in Scotland – where the Analysys Mason assumptions appear to underestimate the cost of performance, both the average task time and the variation in task time are greater than in other GMs (Graph 10). A flat assumption that does not capture the challenges associated with higher and less predictable completion times would be expected to understate costs in this case.

- However, this would not explain the impact seen in the South of England, which generally has shorter task times and less variation.

Graph 10: Average completion time and completion time standard deviation across GM[2]

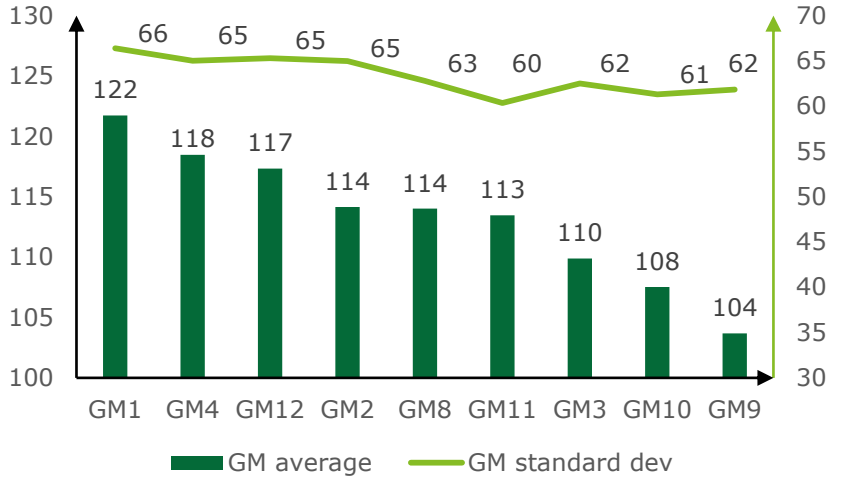

Graph 11: Impact of the flat travel and task time assumption on the cost of performance across GM[2]

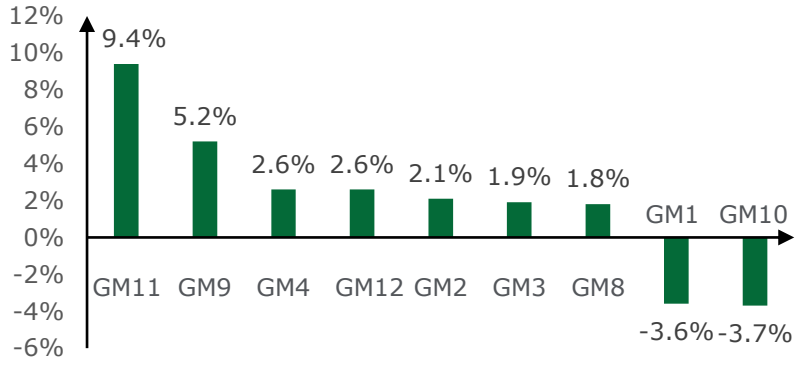

Variation in tasks and resources: job volume

Both models simulate repairs and provisions at SOM level while Openreach operates at OM/PWA level, which face higher variance

Both the Ofcom and Openreach model simulate repairs and provision jobs at Senior Operating Manager area level, while Openreach rather operates at Operational Manager (OM) or Preferred Work Area (PWA) level.

While Openreach did suggest that some recent initiatives encouraged engineers mobility across PWAs of the same SOM, the models' implicit assumption of perfect mobility within a SOM potentially understates the operational constraints faced by Openreach.

This aggregation may lead to underestimate the variability actually faced by Openreach, with job volume variance being on average 30% higher at PWA level than at SOM level. This in turn could lead the models to underestimate the cost of performance.

Ofcom has argued that:

"By setting regulatory standards at a GM level, measured on an annual basis, on a large volume of work undertaken by a large workforce, the effects of variability in input factors is muted".[3]

This assumption was not tested through a modification of the Openreach model. Indeed, assuming the same variance at SOM level as in PWA level would underestimate Openreach's ability to move resource across PWAs. To be fully tested, this assumption would require additional modelling at PWA level, which is outside the scope of this review.

Graph 12: Job volume distribution at SOM and PWA level[1]

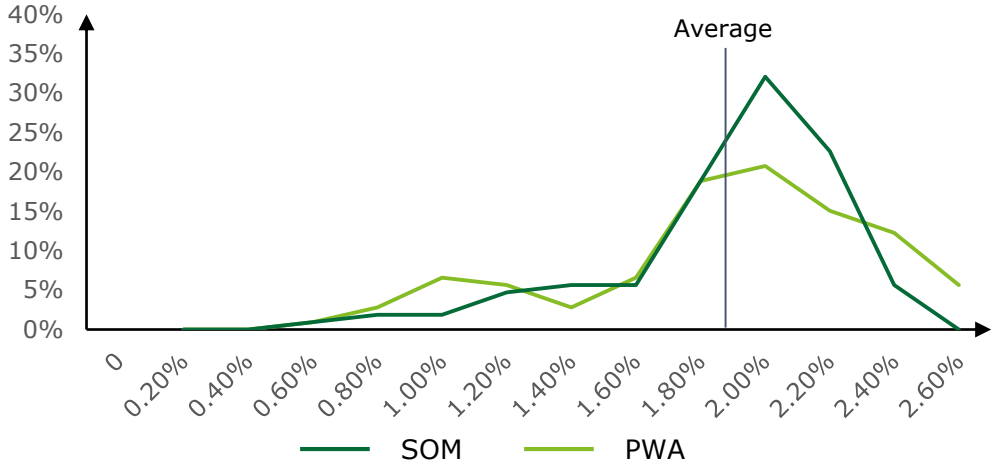

Note: The SOM and PWA represented here are Colchester and 'AT-SAXMUNDHM', respectively, the PWA being part the SOM. Weekly job volumes are expressed as a percentage of 2016-17 total; both distribution's average is 1.9%. Note that the distributions are both moving average approximations of the original, discrete distributions.

Graph 13: Job volume variance by level of aggregation[2]

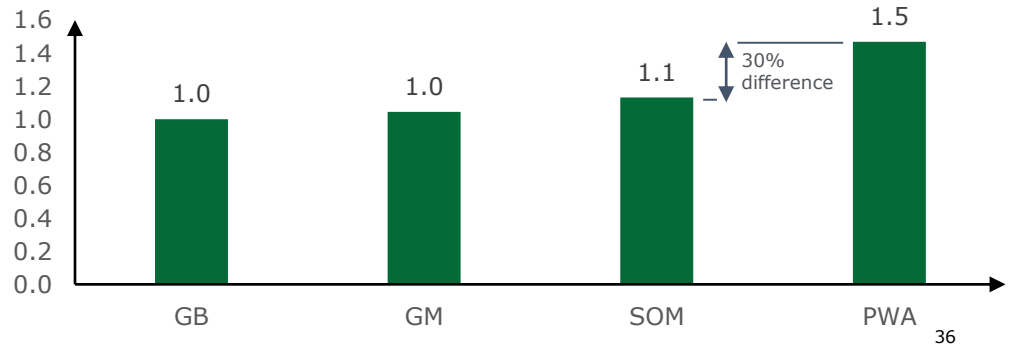

1- Deloitte analysis based on Openreach data.
2- Deloitte analysis based on Openreach data. Great Britain's job volume variance was normalised to 1.
3- Ofcom, Clarifications on the Ofcom Resource Performance Model, 2017 page 2

Patch loans

Some variance in demand can be addressed through patch loans modelled by Analysys Mason

The effects of inputs and demand variance may be mitigated by pooling resources across working areas through "patch loans". Openreach does move its engineers across patches, both in the short run and long term. Ofcom and Openreach both modelled patch loans.

The Analysys Mason model triggers loans based on the spare capacity indicator calculated as the ratio of peak available resource over queue length. The following situations may trigger a patch loan:

1. A SOM's spare capacity indicator is above the 'has spare' threshold and at least one of the neighbouring SOM spare capacity indicator is below the 'no issue' threshold.
2. A SOM's spare capacity indicator is below the 'has issue' threshold and at least one of the neighbouring SOM spare capacity indicator is above the 'has no spare' threshold.
3. A SOM's spare capacity indicator is above the 'crisis spare' threshold and at least one SOM spare capacity indicator is below the 'crisis no issue' threshold. No need for geographical proximity.
4. A SOM's spare capacity indicator is below the 'has crisis' threshold and at least one SOM spare capacity indicator is above the 'crisis no spare' threshold. No need for geographical proximity.

The Openreach model also has a capacity to model loans through an adjustment to SOM's shrinkage rate (which may turn negative if needed).

Openreach patch loans data suggests that loans are highly frequent, with most SOMs receiving a loan over 180 days a year. However, the loans are generally very small, with close to half of all loans representing less than 3 FTEs. This suggests that a significant share of these loans may be generated by Openreach initiatives to make its structure more flexible and enable engineers to work across patch borders (fuzzy patches).

Graph 14: Distribution of SOMs, by number of days in which they received a loan in 2015-16[1]

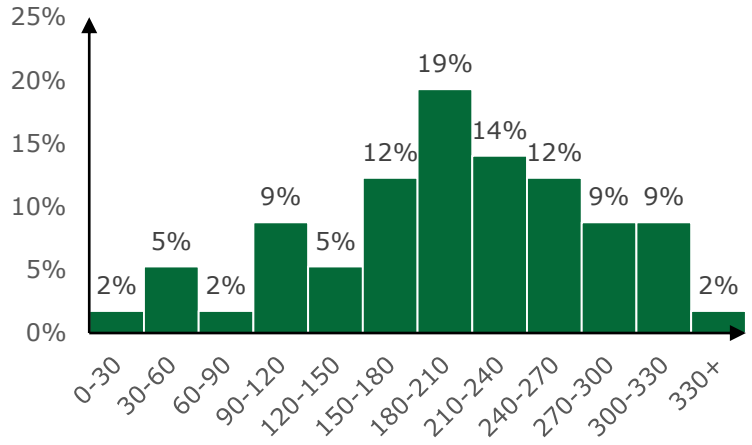

Graph 15: Patch loans distribution, by FTEs loaned[2]

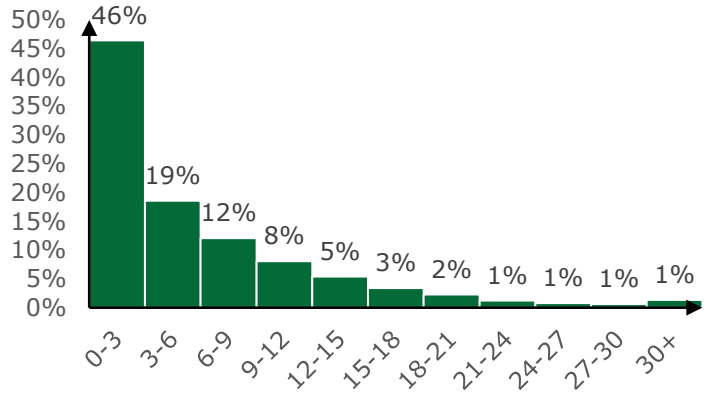

1- Deloitte analysis based on Openreach data
2- Ibid.

Patch loans

Sensitivity testing indicates that the model outputs are sensitive to assumptions made on patch loans, but there is insufficient data to form a view on accuracy

The sensitivity of the cost of performance estimate to the patch loans assumption was tested by rerunning the Analysys Mason model with patch loans disabled.

This analysis suggests that the AM model is highly sensitive to this assumption (Table 19). The required increase in FTEs to reach 95% provision and 93% repair performance in all 9 GMs increases from 4.6% to 9.8%.

The cost performance curves described in Graph 16 show that the patch loan assumption is particularly important for performance level included between 70% and 95%. This may be driven by the fact that at these levels of resourcing some GMs may have spare resource at the same time as others lack resource, enabling loan opportunities. Interestingly, there seems to be some interactions between the patch loan and glass ceiling assumption, with the glass ceiling being slightly lower when no loans are allowed.

While this analysis indicates that the model is sensitive to the assumptions made about the feasibility of patch loans, there is insufficient information available to form a view on whether this approach is realistic.

In particular, there is no information on the frequency and scale of loans generated within the Analysys Mason model. This data could then be compared against the Openreach data presented on the previous slide to determine how well the Analysys Mason model captures reality. Ofcom acknowledge that they have not sought to conduct this calibration, or to test the sensitivity of the model to the assumptions made about productivity.

Table 18: Required increase in FTE to meet performance targets relatively to baseline, by patch loan assumption[1]

95% provision and 93% repair performance	Base loans	No patch loans
National outcome	1.9%	8.0%
All GM	4.6%	9.8%

Graph 16: CL2 performance by number of FTEs and patch loan assumption[2]

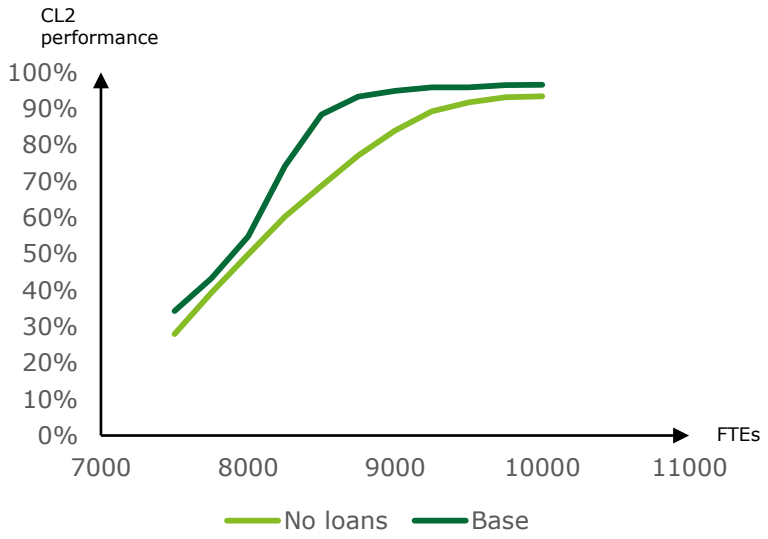

1- Source: Openreach
2- Deloitte analysis on Analysys Mason model and Ofcom data

Glossary

Abbreviation Glossary

The following abbreviations are used in this report.

Abbreviation	Approach
BBUG	Broadband Underground. The highest skill category for Openreach Engineers.
CAL	Customer Access Line. A skill category for Openreach Engineers.
CL	Care Level.
CP	Customer Premise.
FAD	First Available Date.
FAMR	Fixed Access Market Review.
FTE	Full Time Equivalent.
FTTC	Fibre-To-The-Cabinet
GEA	Generic Ethernet Access.
GM	General Manager. Openreach is divided into 9 GMs in Great Britain
OMI	One Man Installer. The lowest skill category for Openreach Engineers.
LLU	Local Loop Unbundling.
MBORC	Matter Beyond Openreach's Reasonable Control. Provision in Openreach's contracts which releases the company from liability under the relevant product terms and contracts.
MSL	Minimum Service Level.
MPF	Metallic Path Facility.
PWA	Preferred Work Area
QoS	Quality of Service.
SLA	Service Level Agreements. A contractual commitment provided by Openreach to telecoms providers about service standard.
SML	Service Maintenance Level. A repair service contract offered by Openreach for fault repair. Specifies a timeline by which faults must be addressed. See MSL.
SOM	Senior Operating Manager. Openreach is divided in 57 SOMs in Great Britain.
UG	Underground. A skill category for Openreach Engineers.
WLR	Wholesale Line Rental.

Deloitte.

This report has been prepared on the basis of the limitations set out in the engagement letter and the matters noted in the Important Notice From Deloitte on page 1.