WITN00460200

Witness Name: GARETH IDRIS JENKINS
Statement No.: WITN00460200
Dated: 1 JUNE 2023

POST OFFICE HORIZON IT INQUIRY

SECOND WITNESS STATEMENT OF GARETH IDRIS JENKINS

I, Gareth Idris Jenkins, will say as follows:

1. I make this second witness statement in response to a request dated 31 August 2022 made under Rule 9 of the Inquiry Rules 2006 regarding Phase 3 of the Post Office Horizon IT Inquiry. This Rule 9 request asked me 56 questions and provided me with 93 documents (over 3,000 pages in total) upon which to comment or to inform my answers.

Introduction

2. On 6 February 2023 I signed a witness statement (WITN00460100) that answered questions 1-15 and 53-56 in this Rule 9 request. I relied on my rights in respect of self-incrimination to decline to answer questions 16-52. This was not an easy decision. I took it upon the advice of my lawyers. I have always said that I wanted to assist the Inquiry but I was concerned, at that stage, that I had little understanding of the documentary material held by the Inquiry about phases 3 and 4 or the scope of the parallel police investigation. I said that I would keep this decision under review with the assistance of my lawyers.

3. Since my first witness statement was filed, I understand from my lawyers that many thousands of documents have been disclosed on the Inquiry's database. I understand that, at the time of my finalising this statement, of the approximately 60,000 documents currently on this database, a very considerable proportion are relevant to the period of time that I worked on Horizon. Being able to see some of this material has proved

WITN00460200

important as many documents I have seen relate to matters that I had forgotten. I understand that disclosure to the database, particularly in relation to phase 4, is ongoing. My lawyers and I have tried to keep on top of this huge and evolving documentary picture, as well as the written and oral evidence given to the Inquiry. As a result, I reached the point whereby I felt that I had a better understanding of the scope of phase 3 and the underlying documents and could better answer the questions asked. I have therefore decided that I will waive my rights in relation to self-incrimination for the purpose of answering the remaining questions in the Rule 9 request (questions 16 to 52). That is the purpose of this second witness statement.

4. But I also wish to make clear that an important factor in my decision to waive my rights and give this second witness statement is because I have read and, in a number of cases, watched the evidence that the SPMs have given to the Inquiry. I cannot begin to understand the terrible things they suffered. They deserve a full explanation for what went so disastrously wrong. For my part, I spent the best part of my day to day working life focusing on the technical minutiae of a system which I considered, overall, to have worked well and to have supported important public services. It has been very difficult to have to confront the experiences which the SPMs have conveyed to the Inquiry. I have been struck by the lack of effective support and assistance that they were given.

5. I had originally been scheduled to appear before phase 3 of the Inquiry on 4 and 5 May 2023. On 20 April 2023, the Inquiry provided my lawyers with 27 further documents and, on 27 April 2023, with 226 further documents. Together this amounted to over 4,000 pages of new evidence (making over 7,000 pages of evidence in total) that I might be asked about on 4 and 5 May 2023. My lawyers told me that a large number of these documents had not previously been uploaded to the Inquiry's Relativity database. Because these documents were additional to those in the Rule 9 request, their relevance was not, in many instances, clear-cut or it was not clear what in a very lengthy document I might need to specifically respond to. However, the Inquiry did identify the page numbers in about 50 of those documents to direct me to some of the relevant extracts. I am mentioning all of this because whilst I have done my best to review this material, there are documents (for example, lengthy technical documents that are relevant to many different issues) the exact relevance of which are unclear. I have not tried in this statement to pre-empt or guess that but rather responded to questions 16 to 52 in the Rule 9 request by reference to the 93

WITN00460200

documents that accompanied this request. Where I think I can, and it might be useful, I have incorporated references to documents within the wider Rule 10 documents given to me on 20 and 27 April 2023.

6. I am grateful to the Inquiry for giving me time to provide this statement with the assistance of my lawyers and agreeing to postpone my phase 3 oral evidence so that the Inquiry and Core Participants have sufficient time to consider it. I have used this time to try to familiarise myself with the new documents I have been asked to consider and I continue to do so. Given the volume of those new documents, I have also had to rely on my lawyers to direct my attention to particular documents.

7. In my first statement at paragraphs 8 to 11, I expressed certain concerns about the limitations of my memory and the incomplete nature of the available documentary record. That I had forgotten certain issues or events has been borne out by my consideration of the new documents. For example, I explained in my first statement that my knowledge of bugs, errors and defects would have been based upon my being allocated specific issues to respond to by third line support (or by my routing to them to others in fourth line support). However, the new documents suggest that I had a greater overview of the issues that arose during the pilot of Horizon Online than I had remembered. I remain concerned that the complete documentary record is not always available and refer to this below.

8. Before turning to questions 16-52 in the Rule 9 request, I would like to expand on two matters raised in my first statement: (a) the EPOSS counter system, and (b) Project Impact and the removal of the suspense account.

The EPOSS counter system

9. In my first statement at paragraphs 28 to 31, I explained that I did not have a direct role in relation to the EPOSS counter system, that in 1998 I had no responsibility for EPOSS and that I could not recall having any involvement in EPOSS at that time.

10. I would like to expand on these comments. The Riposte software that Escher provided to Fujitsu on the Post Office account comprised two main components. The first component was the message store. The second component was the desktop counter application, which was the basis of EPOSS. I was involved in the technical operation

WITN00460200

of the message store from 1996 to around 2003. This component had been designed by Drew Sutherland at Escher, and his main interface at Fujitsu was Mark Jarosz, with whom I worked fairly closely. I was not involved in the technical operation of the desktop counter application until around 2003, when Project Impact commenced. Prior to then, I occasionally became involved in discussions about the counter application, but only when it affected my work on the Agent layer.

11. I recall that a number of Fujitsu employees and contractors attended a three-week induction course in Boston (where the headquarters of Escher are located) in 1996. I attended the first week in order to get an understanding of how the EPOSS counter would interact with the Agents. I then returned home as this first week of training was sufficient to brief me for my role in the Agents team. Others stayed on and attended the full three weeks so that they could learn in more detail how to write counter applications. One of these people was Brian Orzel, whom I came to regard as the main interface between Escher and the counter development teams at Fujitsu on EPOSS matters. Brian worked as a contractor at Fujitsu. I recall that he subsequently spent a year on secondment at Escher's offices in Boston in the late 1990s so that he could work more closely with Escher. Apart from Brian, I regarded Alan Ward and Steve Warwick as the main EPOSS experts within Fujitsu. For a period of around six months in the late 1990s, I recall Alan travelling to Escher's offices in Boston once every week. I believe that Alan and Steve would have the principal authors of the main Fujitsu design documents for EPOSS.

12. The new material provided to me by the Inquiry includes documents FUJ00079301 (E8), FUJ00079303 (E9) and FUJ00079304 (E10), which are emails from 1999. These emails were not available to me when I signed my first statement. They have jogged my memory about certain details but they do not change my fundamental recollection as set out above or what I said in my first statement. These emails relate to the design that was being done by Steve Warwick to add in further reconciliation checks to EPOSS to ensure that the cash account was accurate. I think this was required due to one of the Acceptance Incidents on Legacy Horizon (but I cannot be certain of that). My involvement was primarily to be aware of the new reconciliation totals that were being generated so that they could be harvested by the TPS Harvester Agent at the end of each day, thus allowing the TPS host to carry out the reconciliation checks. In order to do that I was reviewing Steve's design paper and I probably made those

WITN00460200

comments in these emails to show how the reconciliation process could be improved, although I was not responsible for the detail of the design.

13. My main concern was to ensure that I understood the new reconciliation messages that were being produced by EPOSS so that the TPS Harvester Agent could harvest them correctly. As a member of the Agents team at that time, it was my responsibility to ensure that the TPS Harvester Agent did its job properly during the end of day reconciliations, and that no other part of the Horizon system put its operation at risk. I believe that I was also concerned about whether EPOSS was sufficiently resilient to crashes of the PC (for example as a result of a power cut) during these reconciliations. In other words, these emails are consistent with my recollection that, at this time, I was considering EPOSS from the perspective of my responsibility for the Agent layer.

14. The Inquiry has also provided me with [FUJ00079488] (E13), which is a change proposal (*"1. Description of Change Proposed: The Cash Account message attribute grammar is changed to provide the cash account line number and CAP as used by EPOSS in the Cash Account report. The new attribute grammar is defined below"*). Included in this document is an Action Number of 19 February 1998 which states:

> *"TSC IMPACT (G JENKINS) We're currently redesigning TPS and reissuing it for a different reason. We have decided to take the opportunity (sic) of incorporating this CP (together with CP1041 and 1039 and another as yet unnumbered CP on TPS from Stephen Channell) at the same time. As such there will be zero cost, and the handover will be done 27/2/98.19/2/98."*

15. This was a request to change the way in which the Counter recorded a Cash Account Line in the message store. This affected the TPS Harvester Agent, for which (as noted above) I was responsible. I was expected to provide a cost for making the change. My comment indicates that as we were already making other changes to the Agent for unrelated reasons, the incremental cost of this change was zero. Whilst I do not recall this specific change, it is another example of where I was involved because EPOSS impacted on the Agent layer.

16. I have reviewed the evidence that Anthony Oppenheim, Terence Austin and Jan Holmes gave to the Inquiry concerning EPOSS, in particular the decision-making process in 1999-2001 about whether EPOSS should be rewritten, the audit of EPOSS

WITN00460200

and the activities of the EPOSS taskforce. They do not suggest that I had any role – and I do not recall playing one – in responding to concerns about the stability of EPOSS or possible degradation in its coding. I note that Mr Austin, in particular, states that he had a hands-on role in commissioning the EPOSS audit. He named persons who formed the team who responded on EPOSS issues, including Alan Ward (the chief architect), John Hunt, Steve Warwick (whom he described as the expert in EPOSS) and Pete Jeram. This is consistent with my recollection that Alan and Steve were the main experts in EPOSS. I note that Mr Austin told the Inquiry that he didn't know who I was.

17. The Inquiry has also recently provided me with an email dated 10 May 2000 (E11) [FUJ00079333]. I was copied into this email chain, which had originated from an email from Stephen Muchow about concerns about risks of degraded counter and cash account performance and of possible code regression between CI3 and CI4. I believe that I was copied into this email because I had responsibility for the migration of the Data Centre between CI3 and CI4. I note that in the email after the one into which I was copied, Mike Coombs asked that the issue be added to the migration meeting that was being called for. Whilst this was something I would have needed to know about from the perspective of migration, I wouldn't have been involved in the detailed counter changes. This is another example where I became involved in discussions about EPOSS because it was relevant to my other responsibilities.

18. My technical knowledge of EPOSS began to develop in around 2003, when I recall attending a training workshop to show me how to use the counter and, in particular, to understand how the balancing process worked. This was necessary because, by this stage, I was moving away from the Agents team and starting my new role on Project Impact.

19. Prior to this stage, and whilst I had no responsibility for EPOSS, I was aware that it had suffered from problems which I assumed were being worked on and resolved. This knowledge was gleaned from general office discussion rather than any direct involvement. I think it is worth repeating that my base was Bracknell and the EPOSS development was being done in Feltham. I was a frequent visitor to Feltham, but the office I tended to visit was not the one in which the EPOSS developers worked, so I didn't have much contact with them.

WITN00460200

20. I have read the evidence that Anne Chambers gave to the Inquiry concerning EPOSS. She named Matt Arris as the EPOSS team leader when she joined Fujitsu. That name initially rang no bells with me at all, but since hearing it I have a vague memory of someone with that name working on the counter development team. I do not recall having any dealings with him. Anne also told the Inquiry that she understood me to be the principal Fujitsu expert on the counter application, although I note that she said she wasn't sure if I was based in Feltham and that she didn't meet me for two, three or four years after she joined the SSC. I can't remember when I first came into contact with Anne but I think that she is right that we didn't meet for a number of years and that this fits with my role in the early 2000s. In those early years of Horizon when I worked on the Agents, my main contacts in the SSC were John Simpkins, Pat Carroll and Mik Peach. I think that it is likely that I started to work with Anne (or got to know her better) when Project Impact went live in 2004 or 2005 but I can't be certain. I acquired more knowledge and understanding of the counter through my work on Project Impact and I think that I would only have been regarded as having the sort of expertise she attributes to me after that point.

21. In summary, until around 2003, I do not recall having anything to do with EPOSS other than indirectly, when it had an impact on my work with the agents, e.g. when messages were not being written properly which then affected the agent layer. To the best of my recollection, this happened very infrequently. From around 2003 onwards, I began to gain technical expertise and a practical understanding of how EPOSS operated.

Project Impact and the removal of the suspense account

22. In my first statement at paragraphs 41 to 46, I summarised my involvement in Project Impact and, at paragraphs 43 and 44, I explained my understanding of POL's decision to remove the suspense account.

23. Since signing my first statement, I have read the evidence of Susan Harding, Phil Boardman and Steven Grayston. I have also reviewed documents relating to Project Impact that I had not reviewed when I signed my first statement. I agree entirely with the evidence given to the Inquiry that one of POL's objectives for Project Impact was to remove the suspense account, or more accurately to remove all suspense accounts into which SPMs could post discrepancies. I understood that by doing so POL

WITN00460200

ultimately intended to save costs. I also retain a vague memory of Project Impact workshops in which POL employees (I cannot remember who) suggested that SPMs had, historically, used the suspense accounts for money they had stolen and justified their removal on that basis. I was not involved in any POL work or analysis that demonstrated this concern. Rather, it was my job to assist in the design and provision of the system to meet POL's objectives.

24. My recollection was that whilst many suspense accounts were removed, some of them remained in place and that there were still suspense accounts in place at the point at which I retired. I have looked through a number of technical documents to see if they are consistent with my memory. The Inquiry has drawn my attention to an email dated 12 February 2004 (E63) [FUJ00126038] that I and many others received. It refers to POL's decision the day before to remove the suspense account altogether and states that this would have the effect of requiring branches to make good all losses immediately. However, (D90) [POL00038916], dated 16 September 2004, is what I believe to be the final design document that sets out the technical changes which were actually implemented by Project Impact. Pages 147 and 148 list all suspense account products that were present before Project Impact and identify what was to happen to each of them afterwards. This list indicates that some of these suspense account products were to be retained. There is further relevant information in section 2.5.1.1.4 of this document on page 47: this says that the Housekeeping menu (which is where the suspense account products are displayed) is to be restricted to Managers and Supervisors and that the set of products is to be rationalised, with some suspense account products being retained. This design document is consistent with my memory that there was removal of some suspense accounts but retention of other suspense accounts. The documents at [FUJ00090393] (F78) (see the list at 9.3) and [FUJ00090060] (F77) seem consistent with this.

25. I understand that POL's requirements about SPMs settling disputed sums before rolling over raises an entirely different, practical issue. However, from my perspective as a designer of the system concerned with *how* to change the software, I don't think that I appreciated or thought through the potential consequences of these changes.

Bugs, errors and defects: general comments

[Rule 9 request: questions 23 to 49]

WITN00460200

26. The Inquiry has asked me a number of questions about bugs, errors and defects (**BEDs**), both in Legacy Horizon and Horizon Online. I hope it is helpful if I make some general comments at the outset.

27. First, both Legacy Horizon and Horizon Online contained or gave rise to BEDs. It would have been impossible for a computer system, especially one that operated on the scale of Horizon, to be free of BEDs. The fact that BEDs were both anticipated and occurred is the reason why Fujitsu employed large teams of people in four lines of support whose job was to investigate and to attempt to fix them. This also reflects the fact that BEDs cover a multitude of issues. Many will give rise to issues of performance or operational issues. The majority will not impact on the branch accounts.

28. Second, my memory is that BEDs were more prevalent during the testing and pilot phases of Legacy Horizon in 1997-1999 and Horizon Online in 2009-2010. That did not surprise me. It is one of the purposes of having a test phase and then a pilot phase to detect BEDs and to fix as many of them as possible before (a) the system goes live post-testing, and (b) the live system is rolled out to the whole estate post-pilot. As the Inquiry will have seen, there are sequences of actions that occur in real life which may reveal the existence of a BED and which would not have been picked up in testing. The process of testing but then piloting on live data ought to reveal many such sequences. Fujitsu would conduct pilots incrementally across the estate. For example, depending on the complexity of the software change, we would start by introducing it to two or three branches, then anything between 50 and 100 branches, then 1000 branches, and then roll it out to the whole estate. My belief at the time was that both Legacy Horizon and Horizon Online stabilised as they transitioned through these testing and pilot stages. However, there were ongoing issues in Legacy Horizon after its pilot phase ended, which I understood, at the time, to have persisted until around 2003 or 2004, and which (for the most part) related to the Riposte software supplied by Escher. Other specific issues could and did present themselves from time to time.

29. Third, the nature of my role in fourth line support meant that, apart from the pilot phase of Horizon Online between January and June 2010, I only saw part of the picture about the BEDs that affected Legacy Horizon and Horizon Online. I was an experienced architect but I did not have a detailed knowledge of each moving part of Horizon, which was a vast and complex system, nor did I have oversight over the whole system. My

WITN00460200

work on the Agents layer gave me a good overview of the expected functioning of the system but not the detail of all BEDs as and when they arose. What was referred to fourth line support only represented a sub-set of the issues affecting Horizon at any one time (although this ought to have included the most serious issues that third line support could not fix, or issues that required a code fix). It also seems to have been the position (from my understanding of the evidence given to the Inquiry) that PEAKs which reached the SSC may not necessarily have been joined up with each other. Only a small proportion of BEDs were allocated to me in my role in fourth line support, and with some limited exceptions, I was not involved in fixing them. I only learned about certain BEDs many months or years after they had been allocated to and fixed by others.

30. Fourth, of the BEDs that I knew about whilst I worked at Fujitsu, those that I remember as being of particular significance were the Callendar Square / Falkirk bug (in Legacy Horizon), the Receipts and Payments Mismatch bug (in Horizon Online), the Suspense Account bug (in Horizon Online) and the Dalmellington bug (in Horizon Online). I have read Mr Justice Fraser's Technical Appendix to his Horizon Issues judgment, in which he identifies 29 BEDs in total, including the four mentioned above. Having looked at the evidence given to me by the Inquiry, it appears that I knew about the existence of some of the other 25 BEDs mentioned in the Technical Appendix (I expand on this in more detail below). My contemporaneous knowledge of some of these other BEDs was sometimes limited. The underlying evidence suggests that on occasion I was copied into a single email or that someone spoke to me about an issue, which is then reflected in a single entry in a PinICL or PEAK. There are other BEDs that I had a greater degree of involvement in responding to. I think it likely that once some of these other BEDs were fixed, they receded in my memory or their significance diminished. My memory of these other BEDs has been jogged by seeing documents in the civil proceedings and the Inquiry.

31. Fifth, both Legacy Horizon and Horizon Online were configured to create an audit trail and I believed at the time, and I continue to believe now, that this audit trail had forensic integrity. It recorded the transactions as inputted by Branch staff (which might not reflect what they had in fact done) but, in particular, recorded any erroneous transactions that arose as a result of a BED (thus assisting in their identification).

WITN00460200

32. Sixth, Legacy Horizon and Horizon Online were completely different systems as far as the architecture of Horizon within branches was concerned. Horizon Online retained many of the back-end systems and interfaces from Legacy Horizon but was otherwise an entirely new counter system. Amongst other transformations, it removed all use of Riposte and replaced it with a simple Java-based application, and it stored all data in data centres, rather than locally at the counter. As such, I am not aware of any BED that affected branch accounts in Legacy Horizon that also affected branch accounts in Horizon Online (and vice versa). To this day I do not understand, on a technical level, how the same BED could have an impact on branch accounts in both systems.

Data Tree Build Failure Discrepancies

[Rule 9: questions 24 to 29]

33. The Inquiry has asked me a number of questions about data tree build failure discrepancies. These only occurred in Legacy Horizon (not Horizon Online).

34. I have read the second corporate statement of Fujitsu dated 29 December 2022 [FUJ00126035]. This divides data tree build failure discrepancies into two separate categories. Looking at the documents available to me, this division makes sense, although as explained below, I would add a third category.

35. All three categories of this problem concerned the data server component in Riposte but I believe each problem was completely different from the other. The first category seems to have occurred in 1999-2000. I do not believe that I was aware of this first category at the time. The second category occurred in 2005-2006. I believe that I was aware of this second category at the time. The third category occurred in 2007. I believe that I was aware of this third category at the time. I hope it makes sense if I address each category in turn.

36. Before doing so, however, I will provide an overview of what the data server is and how it operated on the counter. This is only an overview because I claim no particular technical expertise on the data server. In simple terms, the data server was a key component of the Legacy Horizon counter system and was used for generating reports and preparing accounts. It did this by scanning the Riposte message store (which was where the Riposte product held its data) searching for “relevant” messages. For this

WITN00460200

purpose, “relevant” messages were all transactions for a given period, usually the current accounting period. This meant that, prior to Project Impact, it would be scanning for up to a week’s worth of messages, but after Project Impact, it could be up to a month’s worth. Each message found would be added to various nodes of a tree of accumulators that reflected the accounting hierarchy of products. For example, a stamp was considered to be, first, a stamp, but also a postage product and ultimately a sale made at the branch. This hierarchy was defined in reference data also retrieved from the local message store. Given the volume of such transactions, this process could be quite slow: it could take up to an hour to build the transactions for a month in a busy branch. Because of this, once the data tree had been built, it was re-used for subsequent processes during balancing, for example recording discrepancies, which also needed to be taken into account in later phases of the balancing. In order to support this functionality without rebuilding the tree, a “notification” mechanism was used whereby a background process would be monitoring the message store for new relevant messages that were written to it, and these new messages were then added to the tree.

First category: 1999-2000

37. My understanding of the first category of the data tree build problem is not contemporaneous: it is based on my reading of the documents given to me by the Inquiry and from what I recall of my work in relation to the Horizon civil litigation in 2018, which is when I believe I first became aware of it. The fact that I did not know about this problem at the time does not surprise me. This is because the data server is a component on the counter EPOSS system which, as explained above, I had no responsibility for at that time.

38. From reading the documents, I believe that what happened is that a bug in the counter EPOSS system resulted in a failure to add relevant messages to the data tree if the counter system was running too slowly or additional messages were arriving too quickly.

39. When I first became aware of this bug in 2018, I agreed that it had the potential to cause discrepancies to branch accounts and would therefore have affected the reliability of Legacy Horizon. That remains my view.

WITN00460200

40. The only caveat I would add is that the bug did not necessarily cause a discrepancy. This is because the shortfalls caused by the bug appeared in an "office snapshot" and not the Cash Account Report, meaning that there was no problem with the underlying transaction data. I agree with the analysis set out in Mr Justice Fraser's Technical Appendix to his Horizon Issues judgment that the bug would not affect branch accounts, that there was no issue with the underlying transaction data and, if the office snapshot was re-run, it would very likely provide the correct information, because the data reading issue was temporary. I also agree with Mr Justice Fraser that a discrepancy would only arise if the branch ran the office snapshot, got an inaccurate report and then rolled over, making good any discrepancies in the process, leading to a shortfall in the branch accounts.

41. Apart from what I can read in the documents, I do not know who or which teams within Fujitsu were aware of this problem. I do not know what steps they took to investigate or fix the problem, other than to say that I believe that it was fixed in late 2000 and that it did not re-occur in any branch. As noted above, the second and third categories of the data tree build problem are different from this first category and so I do not regard them as a re-occurrence. I do not know whether and how Fujitsu informed POL about this problem. I do not know whether and how Fujitsu informed SPMs about this problem.

Second category: 2005-2006

42. I did have some contemporaneous knowledge of the second category of data tree build problems, although I confess that I had forgotten many of the details until I reviewed the documents given to me in the civil proceedings in 2018 and the Inquiry.

43. I believe that the simplest way of explaining this second category is by looking at two PEAKs, both of which were raised on test rigs (i.e. they were not raised from problems concerning live data in branches). The first PEAK is PC0121925 (D77) [POL00028867], raised on 13 June 2005. The second PEAK is PC0123319 (D79) [POL00030523], which was cloned from it on 19 July 2005. The reason that test rigs are relevant to the problem is because in 2005, Fujitsu was carrying out testing following the change in accounting, implemented by Project Impact, from weekly cash accounts to monthly branch trading statements. As a result, the data server needed to be reconfigured to change how transactions were scanned and added to the data

WITN00460200

tree during the balancing process. It was in the context of testing the reconfiguration of the data server that these two PEAKS were raised. The testing identified that some of the (hypothetical) transactions put through the test rigs, such as cash discrepancies, were not being added to the data trees.

44. In relation to the first PEAK (PC0121925), I can see that the problem identified in testing was not reliably reproducible, despite the development team spending a lot of effort trying to do so. Martin McConnell recorded at 12.50pm on 29 June 2005:

> *"After several days of attempting to recreate this problem with a keyboard controlling test program which has generated tens of thousands of EPOSS transactions, thousands of print preview cut off reports I have only seen one instance of this problem whereby a message does not get passed to EPOSS counter code via the message port interface. Am attempting a fresh run as of this morning (29th June) and will leave for a further 24 hours or so to see if I can see the problem at least once more."*

45. Further down this first PEAK, it records that the team were only able to reproduce the error about once in approximately 150,000 attempts (see the entry by Martin McConnell at 10:58am on 4 July 2005), on a process that would only be conducted on live data in actual branches once every few weeks.

46. The analysis in this first PEAK is reflected in the Test Report for S80 dated 3 August 2005 (D65) [FUJ00086360]. This report stated, at pages 40-41, that automated testing on the test rig had repeated the test scenario more than 10,000 times and the problem had not occurred once. It also stated that it was suspected that the test rig *"contributed to the problem occurring and that there was no reason to believe that live would see any more incidences".*

47. The Inquiry has asked me to look at a testing plan concerning the S81 release (D68) [FUJ00086363] and to explain my comment: *"I'm not convinced any fix has been applied, so I'm not sure what you can actually test".* I made this comment on 27 July 2005 in track changes. The full text of my comment reads as follows:

WITN00460200

> "*I'm not sure what you can test here. The original problem was a "one off" which occurred once in about 2000 attempts. Martin had a lot of problems reproducing and I'm not convinced any fix has been applied, so I'm not sure what you can actually test.*"

48. As this full text indicates, I was referring to the first PEAK mentioned above (PC0121925) and reflecting what had been documented on the PEAK by Martin McConnell on 29 June 2005, i.e. that no one had been able to reproduce the sequence of actions at the counter which had brought this problem about and therefore to reproduce it. The fact that it was not reproducible demonstrated either that the sequence of events was so rare as to make it next to impossible to recreate them or that there was some form of fault on the test rig.

49. The second PEAK (PC0123319) was a clone of the first PEAK and provided some further information about a specific data tree problem that was easier to reproduce reliably. The PEAK records that Fujitsu developed a fix for this specific problem by 19 August 2005 and that the PEAK was closed on 5 September 2005.

50. Until I reviewed the documents provided to me by the Inquiry, I had believed that the problem identified in this second PEAK had ultimately been fixed after testing and that it had not re-occurred, either in testing or in the live estate. However, from reading (D69) [FUJ00086368], which is a note of an internal Fujitsu morning prayers meeting on 16 August 2005 that I did not attend, I can see that this PEAK is mentioned next to the comment *"problem recurs after application of fix".* However, as noted above, the PEAK itself states that a fix was developed by 19 August 2015 and that the PEAK was closed on 5 September 2005. On this basis, I assume that the problem in the second PEAK was fixed but I cannot be completely certain about that.

51. In relation to both PEAKs, I do not know whether the problems identified in the test environment ever emerged and affected live data in actual branches, or whether the fix rolled out for the second PEAK proved to be effective. In the absence of further documents, I cannot say whether these problems had an actual impact on branch accounts. Reflecting on the two PEAKs now, and with the benefit of hindsight, I believe that there were probably other processes on the test rig counter that slowed down Riposte and made it very unpredictable. As mentioned in the S80 Test Report, it is

WITN00460200

possible that these processes were caused by the testing environment and would not have occurred in actual branches.

Third category: 2007

52. The Inquiry has asked me to consider PEAK PC0146170 (D74) [FUJ00086490]. This was a PEAK raised on 18 May 2007 and concerns what I would describe as the third category of data tree build problems, which was completely different from the first two categories described above. In other words, this was not an issue that had persisted from the earlier data tree issues but had (from memory) arisen from a new use of the data tree mechanism. I did not make any entries in this PEAK and I doubt that I would have read it at the time. However, I can see that I am referred to in the PEAK so must have had some knowledge of the problem. I retain no memory of the matters raised in this PEAK, but reading it now and trying to reconstruct what happened, I believe that the underlying problem, in simple terms, was that the counter was running too slow and/or Riposte was running too fast, and as a result the data tree was not picking up relevant transactions. I can see from the PEAK that I talked to Gerald Barnes, who was a developer, about why this problem was happening, and that Gerald came up with a fix, working in consultation with Chris Bailey, who was a software designer working in fourth line support.

Callendar Square/Falkirk bug

[Rule 9 request: questions 29 to 35]

53. The Inquiry has asked me a number of questions about what it calls the Callendar Square/Falkirk bug. This occurred only in Legacy Horizon (not Horizon Online).

54. In the Technical Appendix to the Horizon Issues judgment, Mr Justice Fraser grouped together a number of problems under a single heading called 'Callendar Square/Falkirk'. I understand why the Court and now the Inquiry has taken this approach but it risks creating the erroneous impression that the bug was homogeneous in the sense that it was caused by the same factors, and manifested itself in the same way, in every branch it affected. This was not the case. As I explain below, what happened at Callendar Square branch in Falkirk in 2005 was caused by an underlying bug in the Riposte software that manifested itself when a combination of unpredictable factors, including (in that particular case) the steps taken by the SPM

WITN00460200

to rectify the issue, came together. In other words, there was an underlying software bug, but how it manifested itself and whether it caused harm was contingent on the presence of other circumstances. These circumstances were not fixed. That is why the problem was hard to reproduce.

55. This is at the root of the distinction between the limited "time out waiting for lock" issues that I observed between 2000 and 2003 and which I ultimately concluded were benign, and the "Callendar Square" issue in 2005, which involved a storm of multiple events and which was not benign.

56. I draw this distinction to also make the point that the lock error was the symptom of an underlying bug in the Riposte software, rather than the cause of it. Diagnosing the cause of this lock error was so difficult because its emergence was contingent upon a coincidence of unpredictable factors that varied from one case to another. And, depending on precisely which factors precipitated the bug in any given case, the effects of the resulting lock error varied in their nature and seriousness.

57. I should also reiterate that, in the period between 2000 and 2003, my focus was the behaviour of the Agents, so I was mainly interested in occurrences of the "time out waiting for lock" events that were recorded in the Data Centre (i.e. on the Correspondence Servers or Agent Servers), rather than occurrences of these events on the counter. I was confident at the time (and remain so) that where the events occurred in the Data Centre, they were benign, because the Agents (for which I was responsible) were designed so that if such an event was raised, they would handle the issue and carry on with the overnight processing successfully. My understanding, based on my involvement in locking issues between 2000 and 2003, was that where these events occurred on the counter, they were also benign.

58. My understanding of the lock problem developed in three main stages. The first stage was in 2000-2003, when I became involved in a small number of lock problems experienced at different branches and on the correspondence servers. The second stage was in 2010, when I learned more of the detail of the specific lock problem experienced at Callendar Square branch in 2005. The third stage was in 2018-2019, when I learned more about the scope of the lock problems throughout the lifetime of Legacy Horizon as a result of my involvement in the civil litigation against POL. In

WITN00460200

other words, my understanding of the lock problem evolved over nearly 20 years. I hope it makes sense if I deal with each of these three stages in turn.

First stage: 2000-2003

59. The Inquiry has given me six PinICLs (and associated emails) dating from 2000 to 2003 which analysed lock errors experienced at three different branches and on two different correspondence servers. I don't think that all of these lock errors can be explained in general terms. I think it is better if I explain the issues that arose, on a case by case basis, by reference to the specific PinICL. I will address the PinICLs in chronological order.

60. PinICL PC0057478 was raised on 9 September 2000 in relation to branch number 260801. It states that at this branch, the lock error occurred during the overnight migration of counters from one version of Riposte to another. I'm not sure what the reference to "migration" means in this context. It may mean migration from the previous manual system to Horizon. Looking at this date, the national rollout of Horizon was still happening and so this could be a counter operating on its first day running Horizon (but I cannot be sure either way). There were no transactions occurring during this period since the branch was closed. The email refers to a Riposte Index rebuild. This was something that was fairly processor intensive and was scheduled to run out of hours, so that was probably the root cause. As a result, the lock error did not have any impact on the business or its accounts. The Inquiry has asked me why I said, in my email on 21 November 2000 at 4.30pm (D36) [FUJ00083564] that the issue raised in this PinICL was *"not a serious problem (and so low priority)."* Although this issue would have had no impact on the branch accounts, I believe that when I made this comment, I was not referring to any issue in this PinICL (i.e. the underlying software bug or the lock error). I believe I was referring instead to a specific event generated on the counter. To illustrate this, I go back to my earlier email on the same day of 8.20am (D31) [FUJ00083548]. In that email, I referred to a number of errors that I had seen on the event log for the branch, including the "Therefore we killed the process directly" event. As I noted at the time, this event was being generated as a result of the cleardesk function, which was an overnight process that ran around 3am every night to restart the counter to pick up new Reference Data. That process had worked correctly at this branch, so it puzzled me why the event "Therefore we killed the process directly" was appearing. My reference to "3p each for a phone call" referred

WITN00460200

to the automatic call that this event would generate, even though it was not indicative of something harmful occurring. It was this event that I was referring to in my email of 4.30pm later that same day (D36) [FUJ00083564], when I said that it was *"not a serious problem, and so low priority."* Immediately after these words, I added: *"[…] but we would like to get rid of all counter error events"*. I believe I added this caveat because, even though the "Therefore we killed the process directly" event did not seem to be caused by anything harmful, I was still keen to resolve any events that might adversely affect the operational performance of the counter.

61. PinICL PC0056922 was raised on 2 November 2000 in relation to branch number 367642. I forwarded it by email to Mark Jarosz on 3 November 2000 (D29) [FUJ00083544]. The Inquiry has not asked me specific questions about this PinICL but I hope it assists if I make some general observations about it. From reading the PinICL now, it seems that the lock had caused problems at the branch until the counter was re-booted at 3am the following day. It probably arose because end of day processes were running at the same time as the SPM was trying to balance and this overloaded the PC, although I cannot be certain of this. The counter would have restarted at 3am the next morning and cleared the issue. As I understand it, the SPM didn't attempt to carry on using another counter and her accounts wouldn't have been impacted. I agree, however, that Riposte was 'rather sick' in that it made the counter unusable until it was restarted.

62. PinICL PC0057957 was raised on 16 November 2000 in relation to branch number 260801. This is the same branch as for PC0057478 (see above). Both calls were raised by SMC from their event monitoring. The issue occurred at midnight. The Inquiry has referred me to an email I received from Mark Jarosz on 1 December 2000 (D42) [FUJ00083582] about this PinICL and has asked me what Mark meant when he said: *"a timeout of this sort [was] likely to be benign in the sense that it should not result in a message store corruption."* To attempt to explain what Mark meant, I need to go back to my email to Mark at 11.48am on 24 November 2000 (D37) [FUJ00083568]. This email forwarded the PinICL and explained to Mark that I had looked through the message store and the event log and noticed that at the time of this failure (just after midnight) there was an LFS (Logistics Feeder System or Logistics Feeder Service) background task running. I explained that the LFS had written a BLOB to the message store at 0.01am and then a further BLOB around one minute later. I suggested that there were probably message store scans occurring between the two

WITN00460200

BLOBs and queried whether that had caused the time out. I observed that since the failure occurred just after midnight, Riposte would be reloaded soon. Despite the indications that this was not a serious issue, I said that we needed a *"definitive statement"* from Drew Sutherland (the Escher designer of the message store) as to whether this event was benign, or what problems we could have when it happened. It was this request for a definitive statement that I understood that Mark (having spoken to Drew) was responding to in his email to me on 1 December 2000 (D42) [FUJ00083582]. I have read the evidence that Mark gave to the Inquiry about this email although I did not really understand his explanation. When I read his email at the time, I believe that I understood him to mean that, because this particular incidence of the lock error was unlikely to impact the message store, it would not have any impact on branch accounts. I respected Mark, and had I disagreed with this opinion, I would probably have said so in my response. In the same email, Mark told me that:

> *"[…] had the operation which was affected by this timeout been a message server internal operation […] then an additional error event should have been logged. Therefore a possibility is that an API call has time [sic] out and the application is not checking for error events".*

The Inquiry has asked me to explain this comment. I believe that I probably understood Mark to mean that, if an API call timed out, no error code would be returned, so the counter would be unable to detect a problem and would carry on regardless. However, in this case, it was an LFS background task which I think was using the "C" interfaces to Riposte and so it could be checked for response codes. In other words, Mark was alerting me to a potential problem and was saying that whilst the problem was likely to be benign, he could not be definitive about it. This is why he went on in his email to make two recommendations for progressing the matter. His first recommendation was to *"get the LFS agent code checked to confirm that all API calls have error checking".* I passed the matter to the LFS team regardless of whether the LFS Agents were picking the issue up [PC0058994/ FUJ00075544 entry of 18 December 2000]. The PinICL confirms that the LFS team, in turn, discussed the matter with Mark Jarosz. Regardless of whether the applications were checking for errors, the SMC were continuing to monitor events and should have detected if these events were happening (and then raised a call to SSC to investigate or contacted the branch). In terms of this specific PinICL, I can confirm that the timeout which occurred would not have impacted the branch accounts.

WITN00460200

63. PinICL PC0065665 was raised on 3 May 2001 on the Correspondence Servers, and is not therefore related to any particular Branch and would not have impacted on any branch's accounts. The Inquiry has referred me to an email I sent to Mark Jarosz on 11 May 2001 (D45) [FUJ00083600] about this PinICL and has asked me to explain why I said: "*What I'm really asking is for confirmation that the associated errors are indeed benign"*. The reference to errors was to the three groups of events that I described in the email. At the beginning of this email, I stated that I had previously raised with Mark the "Error 82" events raised historically on counters. I also stated that I was aware that the error itself was benign, though it could result in other errors to agents.

64. PinICL PC0075892 was raised on 2 May 2002 in relation to branch number 312511. The Inquiry has asked me why I said, in my email on 8 May 2002 (D48) [FUJ00083621] that the problem was a *"one-off".* From the email, it is clear that I had examined the logs and determined that all had been fine at this branch, which had closed early at 1.30pm. Some hours later at 4.24pm there were some Error 32 messages "Timeout while waiting for thread completion". These occurred at five second intervals. The machine then appears to have rebooted and initialised correctly. However, there were then further Error 32 messages and at 6.40pm an Error 89 message ("An unexpected error occurred while attempting to insert a message. Timeout occurred waiting for lock"). These messages continued to occur until the "3am bounce" (i.e. as part of the overnight cleardesk process that ran around 3am each day). After that, everything went back to normal. Going by what I have written, I probably described it as a "one-off" because I believed that Fujitsu's event monitoring system had detected the issue. Had the issue recurred, I believe that the event monitoring system would have generated similar events/errors. To the best of my recollection when I used this term, there had been no other events/errors about the problem that had caused this lock error. Again, in this case, there would not have been an impact on branch accounts.

65. PinICL PC0087709 was raised on 27 February 2003 in relation to a correspondence server and so would have no impact on branch accounts. I forwarded this PinICL to Iain Janssens on 28 February 2003 (D50) [FUJ00083634] and sent and received emails about the PinICL on 27 and 28 February 2003 (D51) [FUJ00083640]. The Inquiry has asked me what the effect was of the agents failing in this case and why I said: *"Riposte is clearly behaving badly, in that it won't support agents until it is*

WITN00460200

restarted". To put this comment in its context, I go back to my earlier email of the same day, when I said:

> *"We've seen a number of cases of the event 82 in the past, however usually they have been "one-offs" and benign. In this case the agents are repeatedly failing with the agent timeouts coinciding with the event 82s. As the PinICL says, the first error was an Event 89 which I don't remember seeing before."*

When I wrote these emails, I was explaining that it had now become clearer to me that the message store on one of the 16 correspondence servers was no longer operating correctly for the specific duration of the locking issue due to a problem with Riposte (though I can now see that I was mistaken in thinking that I hadn't seen an Event 89 previously in this context: see PinICL PC0075892 above). Looking back at this, I can see that the issue was similar to the other instances of locks in that it was not harmful save that, in this case, one of the correspondence servers needed to be restarted for normal operations to be restored. This was not a branch issue. I do not believe that this had any impact other than a slight delay in processing, because if one correspondence server failed and needed to be restarted, there were three others that could be used immediately instead (there were generally clusters of four correspondence servers linked to 5,000 branches each, so if one of those servers failed, the three remaining servers in the same cluster continued to service all 5,000 branches). In relation to the effect of agents failing in this case, this would have no impact on the branch because the agents were designed to recover the work elsewhere if they hit unexpected errors that forced them to shut down.

66. When I look at these six PinICLs, the two connections that I see between them now are (a) the existence of the lock error, and (b) an underlying but unknown software problem in Riposte. However, how the underlying software problem manifested itself – and caused a lock error – depended on different coincidences of factors on each occasion. Had the factors been the same on each occasion, it would have been far easier at the time to reproduce a case that Fujitsu could send to Escher for testing, so that the underlying software problem in Riposte could be fixed. Unfortunately, we were unable to do so.

67. During the period 2000-2003, my belief was that all of the lock errors I saw were benign, in the sense that they did not have any impact on branch accounts. I formed this belief because, immediately after the event *"timeout waiting for lock"*, Riposte always generated a second event "*failing to insert a message".* If this second event

WITN00460200

was generated due to replication, there were automatic protocols in the system that re-attempted the operation later on, i.e. there was a delay in processing but no data loss. If this second event was due to an agent process, the agent would fail or shut down (as they always did when they hit unexpected errors) but recover the work elsewhere through a fix built into the agent layer that my team had designed. As explained above, the lock error required operations staff to intervene and restart the correspondence servers, but this had no adverse impact on branches or their accounts: there were duplicate correspondence servers and the system was designed to carry on should one of the correspondence servers fail. It may also assist if I explain that Riposte provided two separate interfaces: a full application interface and a simplified interface. The agents (for which I was responsible at the time) used the full application interface. The counter software (for which I was not responsible at the time) used the simplified interface. In the full application interface, every call on a Riposte function would return an error code and the agents were designed to trap any unexpected errors, handle them appropriately and ensure that there were no negative consequences.

68. Looking back at my emails, I can see that I questioned whether the lock errors were indeed benign, and on occasions, I asked Mr Jarosz to investigate further, including by consulting Escher. I did this because I felt that it was important to test whether my belief – that the lock errors were benign – was correct. I don't recall Mr Jarosz saying anything that changed my belief. That is not to say that I was unconcerned about the lock errors, but I saw nothing at the time to indicate that these issues went beyond affecting operational performance. I wanted Escher to improve the Riposte product so that SPMs encountering lock errors no longer had to endure the hassle and wasted time of a counter reboot, which was the only "workaround" we could suggest given our inability to fix the problem. It seemed to be impossible to get to the bottom of the underlying software problem causing the lock errors, and without a reproducible case, it was going to be very difficult to get a fix from Escher. Moreover, had the underlying bug been causing actual discrepancies in branch accounts, I would (as stated above) have expected this to have been picked up by Fujitsu's event monitoring (such as in PinICL PC0057478) or reported by branches.

69. I do not know what Fujitsu did between 2000 and 2003 to inform POL or the SPMs about any of the six PinICLs discussed above.

WITN00460200

70. Before moving on, I will address three miscellaneous questions the Inquiry has asked me about lock errors in the period 2000-2003:

a. The Inquiry has asked whether I recall a release management forum meeting (**RMF**) on 1 August 2001 and has asked why the RMF resolved at this meeting to close four PinICLs (PC0050916, PC0063692, PC0065665 and PC0062490). I was not a regular attendee at RMF meetings and would only have attended if it was going to examine outstanding PinICLs assigned to me. I do not recall whether I attended this RMF meeting. I do not have a copy of the minutes from this meeting that might help me. However, looking at one of these PinICLs (PC0065665), I can see that it contains an entry at 8.48am on 10 August 2001 by Barbara Longley, which reads: *"After last weeks [sic] RMF meeting on Wednesday the following PinICLs were decided to be closed as fixed at Future Release ("Assumed fixed by SP6 at BI2"): 50916, 63692, 65665, 62490."* I interpret the words *"assumed fixed"* to mean that the RMF assumed that the problems would be fixed by Escher during the next release of Riposte. In other words, the problems in these PinICLs had not been fixed yet but the RMF decided to close the PinICLs because it considered that the problems would be fixed in a coming release. Whilst I don't remember that I was involved in this decision, it appears that the RMF may have made the working assumption that the significant update to Riposte (as part of Release BI2 and related to network banking) would encompass a fix of these issues.

b. The Inquiry has referred me to an email I sent on 22 March 2001 (D43) [FUJ00083592] and has asked what "RER" meant. I made this reference to RER in the context of a PinICL assigned to me: I said that this PinICL *"is on the RER as a 'nice to have.'"* I cannot be completely certain, but I believe that RER meant "Riposte Enhancement Register". I believe that the PinICL to which I was referring in this email was PC0061665. The Inquiry has also asked why this PinICL was closed. The answer to this appears to be contained in Chris Wannell's response to my email dated 25 March 2001 (D43) [FUJ00083592], namely that *"being on the RER suggests that it is not a PinICL but a design change to Escher code".* In other words, it was not a bug, but a request for new

WITN00460200

functionality. If we wanted Riposte to be changed in that way, then we would need to commission Escher to make a change.

c. The Inquiry has referred me to an email exchange dated 17 April 2001 between Brian Orzel and me (D44) [FUJ00083596]. This email exchange contains a list of a number of PinICLs, some of which relate to locking errors, to be sent to Escher for fixing. The Inquiry has asked why PC0056922, PC0057478 and PC0057957 were not included in this list. It has been drawn to my attention that one of those on the list is PC0058994 and is described as a copy of PC005795. Fujitsu only needed to send a copy of one PinICL to Escher to be fixed so this suggests that one was selected for being sent on.

Second stage: 2010

71. At some point in 2003, my involvement in the lock errors seems to have ceased. I cannot be completely sure why that was, but I believe it was linked to the fact that I moved from the Agents team and onto Project Impact full-time. It could also be that Brian Orzel left and the interface to Escher moved on to Iain Janssens. It may also have been because the SSC stopping sending PinICLs about lock errors to fourth line support.

72. Prior to this Inquiry, I did not recall having any conversations or learning anything new about the lock errors between 2003 and 2010. However, amongst the materials I was provided with on 27 April 2023 is a lengthy email chain from October 2008 (E19) [FUJ00083712], although the Inquiry has not provided me with the attachments. I was copied into some of the emails in the chain but not others. The chain is a discussion of ARQ data relating to seven different branches: the identifying FAD codes for those branches are listed in the first email in the chain. The Inquiry has said that this email chain relates to the Callendar Square bug, but I am struggling to see the connection. None of the seven branches is the Callendar Square branch. It appears that some of the seven branches suffered from lock errors, though without further details, I cannot say whether this is the same lock error experienced at Callendar Square. By the start of 2010, I think I was aware in very general terms that a small number of branches were still occasionally experiencing single or limited event lock errors, which were very different from the event storms seen at Callendar Square, but I did not know the details.

WITN00460200

73. Putting that to one side, my memory of events is that I knew very little about the Callendar Square issue (in terms of the specific lock error experienced at the Callendar Square branch in Falkirk in 2005) until 2010. I think I had heard about it or heard it mentioned before then but I had no detailed grasp of it until 2010.

74. The reason the Callendar Square branch was brought to my attention in early 2010 was because I was asked to comment about it in the criminal proceedings brought by POL against Mrs Seema Misra. I am aware that this goes to phase 4 of the Inquiry, but I cannot answer this question without some reference to Mrs Misra's trial. In summary, I was asked to provide some evidence about the Callendar Square bug in the context of this trial. I had no first-hand knowledge of it and so I conducted some research about what had happened. Having ascertained what I understood to have occurred at Callendar Square, I set out my understanding of the position to Anne Chambers and asked that she confirm whether my understanding was correct. The email correspondence at (E20) [FUJ00083721] was part of this communication. I did this for the purpose of assessing whether the same problem had affected the branch where Mrs Misra had worked (West Byfleet).

75. The understanding I developed about Callendar Square was that a fault in the Riposte software had affected transfers of stock between stock units. It manifested itself by the receiving stock unit being unable to "see" the transfer made by the sending stock unit, resulting in the transfer message not being written and the counter carrying on regardless with an incomplete view of the accounts. Because the transfer in was not visible, the SPM (unnecessarily but entirely understandably) repeated the transaction, and this caused a receipts and payments mismatch. In other words, what the SPM did – and I do not criticise the SPM in the slightest for doing so – contributed to the problem because the repeated transfer contributed to the missed balance. This would have been visible to SMC since the underlying Riposte issue created a flood or storm of NT events (about 1 every 10 seconds until the counter was restarted) as Riposte repeatedly attempted to write the messages.

76. I don't recall that, in 2010, I drew a connection between what had happened at Callendar Square in 2005 and the lock errors I had seen between 2000 and 2003. I have asked myself why, if my memory is right, I did not draw that connection. I believe the answer is that the circumstances in which the problem at Callendar Square arose

WITN00460200

(the failed transfer between multiple stock units) were very different from the circumstances in which the earlier lock errors had occurred between 2000 and 2003. I also think that the exercise I was conducting in 2010 was the focused one of analysing whether the Callendar Square issue had occurred at Mrs Misra's branch and that I did not research the issue more broadly. It is also possible that I had forgotten the detail of the earlier lock errors by 2010. It was only in 2018-2019 that I recall drawing a connection between them (see further below).

77. Because I was not involved in investigating the Callendar Square bug when it arose in 2005, I cannot be certain who else (apart from Anne and others at the SSC) were aware of it at the time. I became aware in 2010 that Fujitsu had informed POL about the error when it arose but I do not know what steps Fujitsu took to inform POL, the SPM at Callendar Square or any other SPMs.

78. I also understood in 2010 that Escher had developed a fix for the underlying software problem causing the specific event storms seen at Callendar Square which was rolled out by Fujitsu in the S90 release, which went live in March 2006. As can be seen in the email at [FUJ00083721] (E20), Anne told me that:

> *"Anyway it stopped happening once S90 was installed (around 4th March 2006, according to info below). This particular problem would only affect branches with more than one stock unit. It happened several times at Callendar Square, though we never found why they were so badly affected."*

79. As mentioned above, from memory, I may have had some awareness of other occasional time out waiting for lock incidents at this point in early 2010 but had not linked these to Callendar Square.

Third stage: 2018-2019

80. As part of the civil proceedings brought against POL, I was asked by POL's lawyers to look in more detail at the lock errors experienced in Riposte up to 2010. As a result, I reviewed a large number of PinICLs, PEAKs and KELs that I do not believe I had previously reviewed. The Inquiry has recently provided me with a note dated 9 February 2019 (E90) [POL00028911]. I think that the table in section 2 of this note may have been prepared by someone else, but section 3 records my reflections

WITN00460200

having carried out this work. As I observed in section 3, it was only by doing this work in 2018-2019 that I realised the *"full scope"* of the issue. This work developed my understanding of the locking errors in two main ways:

a. I became aware that the locking errors had affected more branches than I had previously realised. I recall suggesting to POL's lawyers that they should refer to the spreadsheet of affected branches (prepared by Anne Chambers just before she retired in 2015) to illustrate how many branches it had affected. I embedded the spreadsheet in the note as I thought that it might be a useful document for the POL lawyers to see.

b. Whilst many incidences of the locking errors were not associated with any discrepancies in branch accounts (e.g. they resulted in operational problems such as screen freezes), I became aware that it was not just the Callendar Square branch that had suffered from discrepancies but other branches where there were counter occurrences of the time-out waiting for lock problem.

81. I have considered why I did not have a fuller picture at the time. I think this reflects the nature of fourth line support whereby not every PinICL on a same or similar issue was referred back to the person who may have dealt with the same or similar issues previously. As noted above, I also believe that my move into full-time work on Project Impact in 2003 may explain why I had no role in investigating Callendar Square and had no grasp of the detail of the Callendar Square issue until 2010 (and was not generally involved in lock issues after 2003).

82. As a result of my work in 2018-2019, and having read the evidence that Anne Chambers gave to the Inquiry about the lock error, I agree with her assessment that it reflects poorly on Fujitsu that it took until 2006 for a fix to be implemented.

Receipts and Payments Mismatch bug

[Rule 9 request: questions 36 to 42]

83. The Inquiry has asked me a number of questions about the Receipts and Payments Mismatch bug. This occurred only in Horizon Online (not Legacy Horizon).

WITN00460200

First awareness of bug

84. The Inquiry has asked me about an email exchange with Mrs Chambers on 6 May 2010 with the subject line "Receipts payments mismatches" (D10) [FUJ00081602]. A Receipts and Payment mismatch is a symptom of something having gone wrong as opposed to the actual problem. Despite the subject line of my email, I do not believe that the events described in this email were caused by the bug discovered several months later (which Mr Justice Fraser called the Receipt and Payments Mismatch bug). I am fortified in that view because when Fujitsu responded to the Receipts and Payments bug several months later, the SSC checked all occurrences of the events which demonstrated the mismatch to have occurred. This exercise highlighted a few instances of other events but which were not the same issue. I have checked the spreadsheet of branches ((F12) [FUJ00081220]) which were affected by the Receipts and Payments bug and the two branches mentioned in the May email are not on it. I agree with Anne Chambers (having regard to the evidence that she gave to the Inquiry) that this was not an early incident of the Receipts and Payments Mismatch bug.

85. I had picked up on the issue referred to in the May email because I had been keeping a watch on event logs and I spotted these error messages. I decided to double check that SSC had been informed of these and that there was a process in place to monitor these events. It isn't clear from the response if that was the case, however Anne did re-assure me that there was a KEL such that SMC should have been picking up such events and raising calls.

86. I have looked through the documents given to me by the Inquiry and tried to work out when I first become aware of the Receipts and Payments Mismatch bug. I believe that this was on or around 27 September 2010. I recall that I had been away on holiday for two weeks and returned to the office on 27 September 2010 (which was a Monday). I recall that the SSC told me that Horizon Online had been generating unexpected error events arising from a receipts and payments mismatch when balancing the accounts. I was immediately tasked with trying to understand this bug. I rapidly produced a note about the bug (D25) [FUJ00083353], spoke to development about it and circulated it to the SSC at 1.50pm on 28 September 2010 (D24) [FUJ00082443]. It is probable that others in Fujitsu (particularly in the SSC) were aware of the bug prior to 27 September 2010 but I do not know who they were or what they knew about it.

WITN00460200

87. In summary, whilst I believe that this bug would have been present in Horizon Online from the beginning of its roll out in early 2010, I believe that I personally became aware of it on or around 27 September 2010.

My note entitled 'Correcting Accounts for "lost" Discrepancies'

88. The Inquiry has asked me a number of questions about my note dated 28 September 2010 entitled 'Correcting Accounts for "lost" Discrepancies' (D25) [FUJ00083353].

89. My understanding of the Receipts and Payments Mismatch bug, as reflected in my note dated 28 September 2010, was that it arose during the balancing process at the end of a balancing or trading period. It arose if a branch cancelled the completion of a trading period, and then re-attempted the completion, which caused discrepancies from that trading period to disappear from the Horizon Online counter, whereas those same discrepancies remained visible on the back-end branch account.

90. It was clear to me when I first became aware of the bug and wrote my note that it could cause discrepancies and thereby affect the reliability of Horizon Online. I think that this was reflected in the speed with which it was responded to by Fujitsu.

91. The Inquiry has asked me about the fix for the bug. Two separate issues arose. In summary, the fix was to correct the bug in the code (to prevent it arising again). Separate to that, there was the need to correct the accounts for those branches which had been affected. In sections 6 and 7 of my 28 September 2010 note, I suggested that this could be done by injecting new data into the main branch database (called the BRDB). I noted as follows:

> *"Fixing the Data for each Affected Branch*
>
> *The data can be corrected by adjusting the appropriate Opening Figures and BTS Data that relates to the current TP. This will result in the Discrepancy needing to be processed when rolling over into the next TP.*
>
> *I propose that if we are to do this then we take a copy of the data for one branch and check out the proposed changes on a test system and then rollover the*

WITN00460200

branch on the test system to ensure that the discrepancy is handled correctly before we attempt to correct Live data. Having done one example in this way, we then need to agree a timetable with Post Office Ltd to correct the other branches and ensure that this is communicated with the Branches to ensure that everyone involved is happy."

92. My note proceeded on the basis that, if there was to be this sort of adjustment, the branches affected should be told of the correction. Ultimately, Fujitsu did not use this method, and instead POL adjusted the accounts using the back-end POL SAP accounting system. I agreed that this was a better way to correct the discrepancies caused by the bug.

93. The Inquiry has asked me why, at this stage (i.e. 28 September 2010), Fujitsu had not informed POL about the bug. My email dated 28 September 2010 made clear my view that POL should be informed of the issue:

> "*We probably need to formally raise this as a problem with POL. I'm not sure how this is done, but presumably you can initiate that."*

94. In section 6 of my note, I recommended that Fujitsu should carry out the steps at section 4 so to identify the *"full scope"* of the problem before communicating it to POL through the usual problem management mechanisms. These were steps that could and should be carried out swiftly. It made sense to me that Fujitsu should approach POL as soon as it was able to explain the full extent of the problem properly and provide recommendations about how to deal with it. I understand that Fujitsu did inform POL and sent POL my note one day later on 29 September 2010 (see below). This was the sort of timeframe I had in mind.

95. The Inquiry has asked me to expand on the proposal I made in my note to monitor 116, 117, 902 and 903 events. Whenever the bug arose, a 902 or 903 event was written into the NT event logs, which were monitored by operations staff (in the SMC), and a 116 or 117 business event was written into the BRDB, which was monitored, albeit not routinely, by the SSC. After discovery of the bug, I recommended that all four events (i.e. 116, 117, 902 and 903 events) should now be monitored more closely both by operations staff and the SSC during the investigation phase to ensure that all occurrences of the bug were detected. I should add that the SMC should have monitored 902 and 903 events in any event and raised a call for SSC to investigate

WITN00460200

(as indicated in the KEL Ballentj1759Q referred to in the May email (D10)). However, I suggested that SSC actively check for all four events to make doubly sure that we picked up every occurrence of this issue. As explained above, I recall that the SSC also looked back in time to check if there were earlier occurrences that they had missed.

96. The Inquiry has asked me why I considered that the full extent of the bug could be identified by Windows NT events (i.e. the 902 and 903 events) and the Counter Business Events (i.e. the 116 and 117 events). The simple answer is that I believed at the time – and I still believe now – that this was a failsafe method of scoping for any occurrences of the specific bug. In other words, I believed that actively monitoring for these four events, both in the past and in the present, would reveal the full extent of the bug.

Emails on 29 and 30 September 2010

97. The Inquiry has asked me a number of questions about emails I sent and received on 29 and 30 September 2010.

98. The Inquiry has asked me to explain the reference to "*Penny in prosecution support*" in my email dated 29 September 2010 (D11) [FUJ00081135]. This is Penny Thomas, who worked in the Fujitsu prosecution support team. Whilst investigating the Receipts and Payments Mismatch bug, I had come across a tool that the SSC had developed, which allowed them to produce a report in a readable format rather than as raw data. I thought that this tool could potentially be very useful in analysing evidence passed to POL for prosecutions and so I flagged it for discussion with Penny. This had nothing to do with the main issue discussed in the email chain (hence I used the words *"as an aside"* in my email).

99. The Inquiry has asked me to expand on my reference to *"(by inserting data into BRDB)"*. If Fujitsu was required to re-insert the lost discrepancies, then new records would have to have been inserted into the BRDB. This would probably have been done using a specially developed script by the BRDB Host development team. However, it was soon agreed (I can't remember exactly when – but this was the outcome of the teleconference with POL the following week) that this was not the best

WITN00460200

approach and that it was better to make corrections to the POL SAP system by POL using normal business processes.

100. The Inquiry has referred me to bullet point 4 in my email of 30 September 2010 and has asked whether I took any steps to investigate whether the events referred to in this bullet point were a separate issue and/or had been fixed. I don't recall whether I was involved in any investigation, but looking at the email, I would agree that these were events unrelated to the Receipts and Payments Mismatch bug which were fixed.

101. The Inquiry has referred me to bullet point 7 in the same email and has asked me whether I can describe the nature and results of the investigation into these two branches (122946 and 374632). As I explained in my email dated 6 October 2010 (D13) [FUJ00081211], branch 122946 lost its opening figures when it migrated (i.e. from Legacy Horizon to Horizon Online), whilst branch 374632 lost its opening figures after 60 days.

102. As can be seen from this email chain, this was in reply to John Simpkins asking me to identify which branches required further investigation by the SSC (see his reference to raising PEAKs for SSC to investigate further). Having passed the information on for SSC investigation, I don't believe I had any involvement in their investigation. However, I can see that the problem at branch 113459 became the subject of further emails in 2012 (see below).

103. The Inquiry has referred me to bullet point 10 in the same email and has asked me to expand on my comment *"lets [sic] wait until POL ask us to fix something before worrying about adding them into the list".* I cannot now recall why I made this comment but looking back on it, I think I was referring to those branches which could get the mismatch in the interim (pending the problem being fixed). This is what Mark Wright seemed to be suggesting in his earlier email, when he said: *"I would suggest that we wait until we know what POL want us to do before we collect any more info for the new offices".* The responsibility for managing this ongoing issue lay with the SSC and Problem Management (and Mark Wright was the main person doing that). I think I was agreeing with him that we could hold off on analysis of any new branches until we knew what POL wanted to do by way of a fix.

WITN00460200

Interaction with POL and SPMs about the bug

104. The Inquiry has provided me with a number of documents relating to Fujitsu's interaction with POL and SPMs about the bug and has asked me questions about these documents. In particular, the Inquiry has asked me to set out details of all conversations and/or meetings that I attended with POL concerning the bug.

105. I am afraid I cannot remember the details of every conversation that I had or meeting I attended 12 years ago about this event, but what I can recall or deduce from the documents are the following significant events:

    a. I first became aware of the bug on or around Monday 27 September 2010 (see above).

    b. I put together the note that described the bug and sent it to the SSC on Tuesday 28 September 2010 (see above).

    c. I believe that Fujitsu first told POL (Emma Langfield) about the bug by email on Wednesday 29 September 2010 at 11.04am: (D12) [FUJ00081137]. This email attached my note about the bug. In the same email, Fujitsu proposed a conference call between POL and Fujitsu on Thursday 30 September or Friday 1 October 2010.

    d. There was a delay of around 48 hours before POL responded to this email at 12.15pm on Friday 1 October 2010: (D12) [FUJ00081137].

    e. I believe that the proposed conference call, which I attended, took place on Monday 4 October or Tuesday 5 October 2010: (E16) [FUJ00081584] (see below).

    f. I do not recall having much further input until early 2011, when I was asked by POL (although I cannot remember whether this approach was made directly or more likely via someone at Fujitsu) to put together a detailed "storyboard" of exactly how the bug had arisen and what the SPMs would have seen on their screens.

WITN00460200

g. I produced a "scoping" document on 11 February 2011: (D17) [FUJ00081527]. I did this to make sure that POL agreed what they wanted me to do before I started work. I sent this scoping document to Rod Ismay at POL.

h. On 18 February 2011, Mr Ismay asked me a number of questions / clarifications about the scoping document I had produced: (D19) [FUJ00081544].

i. On 22 February 2011, I responded to Rod Ismay's questions: (D19) [FUJ00081544].

j. I believe that I probably had one or more conference calls with POL in the next few days about the "storyboard" they had asked me to produce: (D20) [see FUJ00081545].

k. I produced a more developed version of the scoping document, which by now had become the storyboard, on 25 February 2011: (D21) [FUJ00081550]. Looking at this now, however, I cannot be certain whether this was the final version.

l. After I produced the final version of this document, I do not recall having any further interaction with POL about the bug until April 2012, when Andrew Winn at POL asked me questions about the list of affected branches (see below).

m. After I dealt with those questions, I do not recall having any further interaction with POL about the bug until May or June 2013, when I remember having a number of calls with POL senior management and POL's lawyers, about informing Second Sight of the bug.

106. I am aware that the Inquiry is interested in the typed note made of the conference call that took place between POL and Fujitsu on 4 or 5 October 2010: (E16) [FUJ00081584]. I have not been asked questions about this note in the Rule 9 request. I did not write this note. I have a vague memory of this call but I cannot recall the exact details. What I can recall is that the basic approach to investigating the bug was agreed, as well as how to correct the accounts of the affected branches. I then left it to the SSC, problem management and POL to implement what had been agreed. In terms of how POL and Fujitsu agreed that the accounts should be corrected, my

WITN00460200

recollection is that there was very little debate about which of the three "Solutions" set out in the note should be used because everyone agreed on "Solution Two", i.e. an amendment to the accounts by POL using POL SAP. "Solution One" in this note did not reflect what was set out in my note dated 28 September 2010, which proceeded on the basis that SPMs should be informed of the change to the data.

107. In relation to what Fujitsu or POL told SPMs about the bug and when this happened, I do not believe that I was involved. However, I believe that POL did communicate with all of the affected branches in order to explain the bug and to reassure the SPMs. This was envisaged in "Solution Two". However, I cannot recall exactly what these communications involved.

108. I understood that Fujitsu rolled out the fix to all branches in late 2010. My recollection is that Fujitsu identified a precise scenario that had resulted in the bug, and so it was relatively easy to run that same scenario on the fixed code to ensure that the fix was successful. As with any other fix, I believe that the fix for this bug also went through a set of standard regression tests to ensure that there were no adverse side effects.

109. I do not know the detail of all of the investigations, but I recall that the SSC continued to monitor the events raised by the bug until late 2010 and circulated a weekly email with an updated list of affected branches. Fujitsu applied the fix to all branches and then continued to monitor for any further occurrences of the tell-tale events for at least a further month to ensure that all instances of the bug were identified. As far as I recall, this monitoring did not identify any further branches as being affected by the bug after the fix was rolled out.

110. To my knowledge, Fujitsu did not exercise its remote access powers at any point in investigating or fixing the bug. No doubt the SSC will have viewed live data from within the BRSS (a read-only copy of the live BRDB system) in order to investigate and diagnose the problem. However, viewing live data on the BRSS would have had no impact on live data in the BRDB. I also assume that the SSC may have retrieved diagnostic logs from the counters, although I do not know for sure whether this happened. If it did happen, this was a standard process that simply retrieves – but does not amend – live data. As noted above, I understood that POL adjusted their version of the accounts of affected branches using the back-end POL SAP accounting system. This had no impact on the accounts as seen in the branches.

WITN00460200

111. I cannot identify with confidence every team or person within Fujitsu who knew about the bug. All I can say is that I understood that various individuals within development, the SSC and problem management were aware of it, but there may have been others.

Identifying affected branches

112. I understood that Fujitsu set up a detailed monitoring exercise to detect any other instances of the bug and continued this monitoring until it was confident that the fix had been rolled out effectively to all branches. Initially the number of affected branches appeared to be in the region of about 40. After the investigations described above, I believe that the SSC calculated that the bug had occurred on 64 occasions in 62 branches in total and produced a final list of those branches. The bug had caused some of those branches to suffer losses and some of them to experience surpluses. I understood that where the branch suffered a loss, the SPM was reimbursed, and where it had experienced a surplus, the SPM was allowed to keep the surplus. To put that figure of 62 branches into context, there were around 12,000 branches in the whole estate at that time.

113. I note that in the email correspondence at (D14) [FUJ00081214], Antonio Jamasb of POL sought an update on the number of affected branches because he was speaking to senior stakeholders within POL. Mark Wright stated that he had been sending a report every week to the POL Duty Manager. This reflects my understanding that POL received weekly reports about the extent of the impact of the bug.

114. As I have noted above, this specific Receipts and Payments Mismatch bug arose if a branch cancelled the completion of a trading period, and then re-attempted the completion, which caused discrepancies from that trading period to disappear from the Horizon Online counter, whereas those same discrepancies remained visible on the back-end branch account. The issue would always occur in the event that the final rollover was partially cancelled, then re-attempted without exiting the process. That this was a rare sequence of activities was demonstrated by the issue having only occurred around 64 times (in the content of 12,000 branches rolling over successfully every month without hitting this sequence of activities). This would help to explain why it had not been discovered during testing. That it was rare did not matter – it was responded to swiftly and decisively.

WITN00460200

115. A small number of branches that were initially included in this list were subsequently removed from it. These included branches 122946, 113459 and 374632, and the Inquiry has asked me why they were not considered to be affected the bug. Looking back at the emails, I believe that the SSC reached this decision because these branches did not generate both: (a) a 116 or 117 event in the BRDB, and (b) a 902 or 903 event in the event logs. Both (a) and (b) needed to be present in order for the bug to be present.

116. I believed at the time that the SSC had identified every branch affected by the bug. From time to time, I saw the reports that the SSC produced that identified the affected branches (because I was copied in) for the purposes of updating POL. I believed that the investigations that the SSC conducted were thorough. If there is evidence to suggest that the SSC did not identify all affected branches, I would be very concerned by that and I would welcome the opportunity to comment on it.

The other branches

117. On 2 April 2012, Andrew Winn at POL emailed me to ask about two branches that had been removed from the list of affected branches early on in Fujitsu's investigation into the bug: (E95) [POL00029718]. The two branches were 122946 and 113459. Having investigated the problem, I replied on 10 May 2012, explaining that these two branches had been affected by separate problems not related to the Receipts and Payments Mismatch bug. In relation to branch 122946, I noted that Steve Parker had initiated an investigation into the problem on 5 January 2011 using PEAK PC0207483. I said that I had been unable to work out what had happened to the investigation into branch 113459. I suggested to Mr Winn that POL and Fujitsu's problem management teams should liaise with each other to get to the bottom of the matter.

118. My review of the documents recently provided to me by the Inquiry include:

    a. An email dated 16 April 2013 from Andrew Winn to Steve Bansal: (F124) [POL00098016]. I was one of a number of people copied into the email. In the email, Mr Winn again asked about the same two branches and sought confirmation that there was no unresolved problems in relation to them. I note

WITN00460200

that Mr Winn stated that POL asked for Fujitsu's investigation into branch 122946 not to proceed. I do not know why this happened.

b. An email dated 2 July 2013 from me to Mark Wright and Pete Newsome noted that Andrew Winn was chasing Steve Bansal about the two branches: (E96) [POL00029719]. Later that same day, as part of the same email chain, I can see that Mark Wright unearthed an email from Steve Parker to Andrew Winn dated 6 January 2011 which gave the results of the investigation into branch 122946. It seems that Andrew Winn was unable to find this email in 2012/2013 when he asked me about branch 122946. I can see from the email chain that Pete Newsome forwarded this 6 January 2011 email to Rod Ismay on 3 July 2013.

119. Ultimately it appears that these two branches were affected by problems which were different from the Receipts and Payments Mismatch bug. That is why they were removed from the list of affected branches in late 2010. The investigation into branch 122946 concluded in January 2011 and Mr Winn appears to have been notified of this. I am unclear what happened to the investigation into branch 113459, although I note that, in his email dated 3 July 2013 to Rod Ismay, Pete Newsome said that he had an answer in relation to that branch and would forward it shortly.

Email of 12 November 2010

120. In the context of the Receipts and Payments Mismatch bug, the Inquiry has asked me about an email I sent on 12 November 2010: (D14) [FUJ00081214]. This email was a reply to a request from Antonio Jamasb at POL, who (amongst other points) requested a *"summary from Fujitsu stating why we have no other integrity issues with Horizon and why we couldn't see this issue".*

121. Mr Jamasb's email was forwarded to me and John Simpkins and I then forwarded it to Mark Wright (because it was he and not John Simpkins who was dealing with the Receipts and Payments mismatch bug). I said in my email to John Simpkins, Mike Woolgar and Mark Wright that I didn't think we could make such a statement and that what we could do was to check through what *"known integrity issues"* we have and

WITN00460200

also make the more general statement that when integrity issues arrive, then they leave a trail enabling them to be identified and their scope to be ascertained.

122. By *"known integrity issues",* I meant any issues that had come to light in the very early stages of Horizon Online which had the potential to cause discrepancies to branch accounts. I asked Jon Simpkins and Mark Wright whether they were aware of other integrity issues that we hadn't fixed because I could not think of any. Mark Wright responded, in turn, that he did not think there were any more integrity issues but that we still had recovery issues and duplicate JSNs in the audit trail. To put these in context, recovery issues relate to problems with communications failures (and which I deal with below) which will occur occasionally in all systems. They were spotted by the reconciliation mechanisms and SPMs were prompted to follow certain processes where they arose. As regards duplicate JSNs, I recall that there had been some issues in the pilot phase of Horizon Online about duplicate JSNs but that these had been fixed some time before this email correspondence. I don't understand these to have impacted the branch accounts (and recall that there were checks for them as part of the ARQ retrieval process). I left it to Mike Woolgar to respond back to Antonio Jamasb as he considered appropriate.

123. What I was suggesting in this email was that Fujitsu inform POL of any integrity issues we had and also point towards the audit trail as a means of identifying those integrity issues. A number of problems had become apparent during the pilot of Horizon Online and the Receipts and Payments Mismatch bug had been discovered thereafter. This made me wary of giving a blanket assurance, at what was still an early point in the life of Horizon Online. The audit trail could however be relied upon to demonstrate if there was a bug going to the integrity of the system.

Suspense Account bug

[Rule 9 request: questions 44 and 45]

124. The Inquiry has asked me a number of questions about this bug, which only affected Horizon Online (not Legacy Horizon).

125. I believe that I first became aware of the bug at some point in March 2013, although I am afraid that I cannot be more precise than that.

WITN00460200

126. In order to refresh my memory of the bug, I have read my report dated 15 May 2013 (D26) [FUJ00083375], which was largely based upon work previously conducted by Anne Chambers (and others in the SSC). This contains a description of the nature of the bug and its consequences.

127. To summarise, the bug was caused by a failure to fix an earlier problem in 2011. My report described this earlier 2011 problem as follows: "*in April 2011 a problem was found with the archiving strategy related to Stock Units that have been deleted in Branch".* What I meant by this was that if a branch deleted a Stock Unit, the removal of the data associated with that Stock Unit by the archiving process was faulty. The changes implemented in 2011 were designed to cure this problem by amending the metadata in BRDB_ARCHIVED_TABLES. Whilst these changes fixed this problem, it created another problem, which was that, if a branch deleted a stock unit between the dates on which PEAK PC0203522 and PEAK PC0208783 went live, temporary data used in calculating the local suspense account was not archived when it should have been, and so was erroneously re-used a year later, leading to new reconciliation problems.

128. My note said that this problem would occur *"under some specific, rare circumstances"*. What I meant by this was that, for the bug to arise, a stock unit would need to have been deleted within a relatively narrow window of two or three months (i.e. between the fixes of the two PEAKs mentioned above). It would also have been unusual for a branch to delete a stock unit in any event. Hence the circumstances in which the bug could arise were rare, as reflected by the fact it was only found to have occurred on 14 occasions in 12 branches.

129. Dealing with the other questions the Inquiry has asked me about this bug:

    a. I agree that the bug resulted in discrepancies in these 12 branches: I noted in my report that five branches had losses and seven had gains (and two had both). This was the totality of the affected Branches. However the problem would re-occur annually until the offending data was removed.

WITN00460200

b. I do not know every team or individual within Fujitsu who may have known about the bug, but I believe that individuals within development, the SSC and problem management were all aware of it.

c. I understood that POL was already aware of the bug before it was first reported to Fujitsu in early 2013. I do not know when POL first became aware of it but I understood that an SPM had raised the issue with them in 2012, but for reasons I do not understand, POL did not pass the problem to Fujitsu. I understood that another SPM subsequently raised the issue in 2013 and this led to Fujitsu becoming aware of it.

d. As far as I can recall, the bug was resolved by SOT, who applied a script developed and tested by the host development team to modify the database to remove the old records that should have been archived. This is what I meant when I said in my note that *"these records have been manually deleted".* I did not think of this as a form of remote access involving the deletion of live financial data; it was a change to temporary data held in Horizon Online's back-end systems that should have been deleted many months earlier. This was done to prevent SPMs experiencing further balancing problems.

e. I understood that POL informed the SPMs in the 12 affected branches about the problem, including that it had been fixed, but I do not know exactly what they were told. My paper contained draft terms for a letter to the SPMs which explained what had happened. I think this was drafted by Andy Winn and was incorporated into my note. The note refers to the proposed letter being approved by POL's legal department prior to sending (this does not surprise me as the letter contained information which could only have been provided by POL).

f. I don't think I was involved in specifying or applying the fix. Anne Chambers had identified the offending records and arranged to have them removed. I may have been consulted, but Anne did all the hard work. We also introduced a new test into the balancing process that would detect if there were any potentially spurious old records around, so that if in future we encountered something similar, the SPM would have been made aware of it. This can be seen at paragraph 2.3 of the note.

WITN00460200

g. I believe that the bug was not picked up in earlier testing because the issue did not become readily visible until Horizon Online had been running for at least a year.

130. In summary, whilst the Suspense Account bug was a serious issue, I believed that Fujitsu reacted quickly to the bug (as soon as they became aware of it) and that there were only 12 branches where it caused actual discrepancies, all of which were rapidly rectified. I am not aware of any recurrences of the bug thereafter.

Dalmellington Bug / Branch Outreach Issue

[Rule 9 request: questions 46 and 47]

131. The Inquiry has asked me a number of questions about this bug. This affected Horizon Online (not Legacy Horizon).

132. I first became aware of the bug after I retired in February 2015. After I retired, Fujitsu kept me on a retainer for ad hoc consultancy work and I was asked to spend a couple of days assisting them with this issue. Partly because I had retired, I am unsure who within Fujitsu knew about the bug before I did. I prepared a note about the bug on 15 November 2015 (D56) [FUJ00085882]. I believe this sets out a full description of the bug and its consequences.

133. To summarise, the bug occurred when an SPM owned both an outreach branch and a core branch. Outreach branches are typically only found in rural or remote locations: they are opened by an SPM for a few hours each week, usually in a village hall or pub. The SPM needs to remit cash and stock from their core branch (which is open all week during normal office hours) to their outreach branch, and then, if there are any surpluses, remit those surpluses back again from the outreach branch to the core branch.

134. Outreach branches are typically not very busy, which may lead to the counter sitting idle for 75 or more minutes. If this happened during the Log On process, or if a user forced another user's profile to log out of the counter, the Log On script would remain on the stack whilst incomplete. This was the precondition to a lengthy and complicated series of further steps that might sometimes (but not necessarily) arise. Assuming

WITN00460200

these steps did arise in a particular order, when the user hit the enter button on the screen (ostensibly to complete the remittance process of cash or stock from the outreach branch to the core branch or vice versa), the function completed but did not tell the user that it had completed. This may have led to a false belief in the user that the cash or stock had not been remitted. If, as a result of this false belief, the user hit the enter button again, the cash or stock was remitted again, leading to the function being repeated, causing a shortfall of cash or stock on the branch accounts due to the duplicate remittances. This would have been fairly obvious to SPM as a new receipt was printed for each spurious remittance. As a result, the bug, exacerbated by the (entirely understandable) acts of the SPM, could result in multiple remittances of cash or stock between outreach branches and core branches that would have shown as cash or stock losses if not rectified, as the cash or stock logged on the system would not match up with the actual cash or stock in the branch.

135. When I first became aware of the bug, it was immediately obvious to me that it had the potential to cause discrepancies in branches. However, as explained below, it appears (save for the branches in respect of which I do not know the position) that it did not result in actual discrepancies.

136. Dealing with the other questions that the Inquiry has asked me about the bug:

    a. I do not know when Fujitsu first informed POL about the bug but I understood that the reason I prepared my note was so that it could be provided to POL to assist with their understanding of the problem.

    b. I do not know exactly when Fujitsu or POL contacted the affected SPMs or what they said, but I understood that Fujitsu support staff were in contact with the affected SPMs because the call had been raised by those SPMs in the first place. Indeed, I recall that Fujitsu contacted the affected SPMs directly with a suggested avoidance (namely to reverse the duplicate remittances in) whilst the bug was being fixed.

    c. I understood that the Fujitsu development team subsequently identified the actual bug in the code and developed a fix. Because I had retired, I was not involved in developing this fix and so I do not recall any details about how it worked, or when it was rolled out to branches.

WITN00460200

d. I do not know what steps Fujitsu took to identify the extent of the bug. However, I recall that, on my recommendation, Fujitsu reviewed its archives looking for the symptoms of the bug since Horizon Online went live in 2010. I believe that the results of this review were fed into the Power point presentation at (D27) [FUJ00083379], which documents the findings provided to POL on 10 December 2015. I believe I first saw this presentation in 2018 (I do not recall that I had any part in preparing it). The presentation stated that there had been 112 occurrences of the bug in the previous five years. Based on the analysis carried out by Mr Justice Fraser in the Technical Appendix, I believe that this figure is wrong. My understanding based on the work I did in 2018 is that the first 65 of these occurrences, which happened between February 2010 and January 2011, were caused by two different problems with similar symptoms, all of which had been fixed by January 2011. I believe that the number of occurrences of this bug was 47, and of those 47, my understanding (which seems to be supported by Mr Justice Fraser's analysis) is that POL corrected nearly all of them at the time by means of transaction corrections (I think that some may also have been corrected by the SPM just reversing the duplicate transaction). The presentation states that there were four uncorrected instances of the bug. I do not know what was done in relation to those four cases.

e. I believe that the bug had not been identified in testing because the sequence of events required to create the bug was very obscure. I outlined in my note dated 15 November 2015 the twelve steps that had to occur in the correct order for the bug to arise. Moreover, this bug only arose when dealing with remitting cash or stock to or from outreach branches, which were a very small percentage of POL branches nationwide.

137. In summary, whilst the Dalmellington bug was a serious issue, I believed that Fujitsu reacted quickly to the bug after they first became aware of it, and that there were only four branches where it may have caused actual discrepancies. From reading the second Fujitsu corporate statement given to the Inquiry [FUJ00126035], I understand that the fix was rolled out in January 2016. I am not aware of any recurrences of the bug thereafter.

WITN00460200

Withdrawn stock discrepancies

[Rule 9 request: question 43]

138. The Inquiry has drawn my attention to PEAK PC0208335 (D76) [FUJ00086720] and asked me to explain the *"known problems with declarations containing withdrawn products".* I do not recall the specific problem described in this PEAK, but reading it now, I believe that the *"known problems"* were that, if POL withdrew a stock product, after a period of time Horizon Online removed the reference data for that product. This could cause problems when SPMs conducted their stock declarations. If an SPM attempted to update an old declaration a long time after the withdrawal of a product, there would be no reference data to identify the product. I don't think that the failure in the declaration process could lead to discrepancies in the branch accounts.

139. The Inquiry has asked me how the particular problem described in this PEAK was fixed. I proposed how to fix the problem in my entry on the PEAK dated 16 February 2011 at 11.45am, but I cannot tell from the PEAK whether this actually fixed this problem, or whether Fujitsu used an alternative fix. Either way, the PEAK records on 10 September 2012 that the problem was fixed.

A hypothetical problem

[Rule 9 request: question 48]

140. The Inquiry has asked me questions about PEAK PC0187425: (D22) [FUJ00081770]. I do not recall this PEAK, but reading it now, I believe that, in early 2010, I was reviewing some code on Horizon Online and came across a hypothetical problem whereby, if a record was not locked before updating, another user could potentially amend it from elsewhere. This problem seemed to me to be hypothetical because it could only arise if there were two counters accessing and updating the same database entries simultaneously (within the same fraction of a second), which in normal operating conditions should not happen. Nonetheless, I raised a PEAK to look into whether my hypothesis was correct. My concern was that, had this problem occurred in live data, only one of the two operations would succeed, and the other would be ignored without that user being informed. In other words, the other user would mistakenly think that he or she had updated the data.

WITN00460200

141. This PEAK was hypothetical in the sense that the factual scenario that might give rise to it would not occur. As far as I am aware, the problem never occurred on any live data. Even if the problem had occurred on any live data, I do not believe that it would have caused discrepancies in branch accounts because the updates were not to transactional tables.

142. I do not know whether POL or SPMs were informed about this PEAK but given its hypothetical nature I would be surprised if they were.

Giro payments

[Rule 9 request: question 50]

143. The Inquiry has asked me to consider two documents which relate to the investigation by Second Sight. My understanding is that the Second Sight investigation will be explored in phase 5 of the Inquiry, but I will do my best to answer the Inquiry's questions on this specific issue.

144. I understood that at some point in 2012 POL commissioned Second Sight to investigate the Horizon system. I recall having a few conversations with Ian Henderson, who was an investigator with Second Sight, and I recall assisting with the technical aspects of some reports that I understood POL sent to Second Sight. I believe that (D9) [FUJ00080537] was one of these reports, since the identifier SR011 was used for "spot reports", which was the name given to reports produced by POL and sent to Second Sight to assist them in their investigation. POL's legal team sometimes asked me to carry out technical analyses of Horizon data, which I understood POL would use in these spot reports. I believe that this particular report contains some of my analysis, including the section that describes the Horizon Online counter processes, but without seeing the accompanying emails, I cannot identify the specific wording I suggested.

145. The Inquiry has asked me about my *"feedback"* on this report (the word used in the email dated 25 March 2013 at (D8) [FUJ00080536]). I do not recall what feedback I gave, and from what I can see, the report does not contain amendments in track

WITN00460200

changes or any other information that would assist me in working out what feedback I gave.

146. The Inquiry has asked me why the email and attached report were marked as *"subject to legal professional privilege"*. I did not send the email or write the report, so I do not think I can answer that question. What I can recall is that POL's lawyers asked Fujitsu to mark such spot reports as legally privileged. I assumed that there were good legal reasons for doing that.

Remote access

[Rule 9 request: questions 16 to 20]

147. The Inquiry has asked me a number of questions about remote access. My understanding of remote access differed between Legacy Horizon and Horizon Online, so I will address the two separately. Before doing so, however, I would like to make some general comments, all of which apply to both Legacy Horizon and Horizon Online:

    a. I have never used remote access. I did not have the permissions or privileges to be able to do so. None of my jobs ever required me to have this functionality. As a result, I have only ever had a general (and indirect) understanding of the nature of the remote access rights POL or Fujitsu could and did exercise. I have read the evidence to the Inquiry given by Richard Roll, Mik Peach, Steve Parker and Anne Chambers and I would defer to their direct knowledge of remote access. For the same reasons, I do not know who had remote access rights, the periods during which they had remote access rights or the details of their security vetting/qualifications.

    b. There has always been a basic distinction in my mind between remote <u>business</u> access exercised by POL and remote <u>support</u> access exercised by Fujitsu. An example of the former was POL's use of transaction corrections in Legacy Horizon, which was a type of external input into Horizon which the user of the system had to authorise. The latter encompassed different kinds of scenarios – from rolling out software updates including fixes, improving the operational performance of counters and updating the back-end databases. It also encompassed injecting data to correct a problem, and to my mind, this involved

WITN00460200

either restoring data that had been lost, or alternatively adding new data to reverse or amend the effect of existing transactions.

c. I never understood or perceived that there was anything secretive or sinister about the fact that Fujitsu had used remote support access. I do not know how it would have been feasible to operate Horizon without it. When I met Second Sight in September 2012, I told them this. As far as I am aware, every computer system, large or small, has a form of remote access built into it.

d. My understanding when I was employed by Fujitsu – and my understanding now – is that Fujitsu never had the capacity to delete data from branch accounts. I only understood that it had the capacity to inject and thereby add data to branch accounts (which would not cause any deletion of existing data). (Fujitsu was also able to inject data that would effectively amend Persistent Objects in Legacy Horizon, for example to change the current TP of a Stock Unit).

e. My understanding when I was employed by Fujitsu – and my understanding now – is that whilst Fujitsu had the capacity to inject data to branch accounts, this happened very rarely.

f. When Fujitsu injected data, I believed at the time – and I still believe now – that all such injections would be visible in the audit data.

g. When Fujitsu injected data, I believed at the time that POL approved such additions. I now understand that there were occasions when POL was not approached and did not approve it.

h. I don't recall giving thought at the time to what SPMs were told about Fujitsu's exercise of its remote access powers. Looking at it now, however, if Fujitsu was simply making background changes to improve the operational performance of the counter, I don't see why Fujitsu would necessarily inform the SPM (just as many computer systems run software updates without the user being asked explicitly to approve them). However, I can see that the injection of data to change branch accounts raises entirely different issues.

WITN00460200

i. If the relevant teams within Fujitsu injected data to enable action to be taken to correct a discrepancy caused by a BED or other problem, I cannot see how this could have affected the reliability of the branch accounts. The only purpose of injecting data was to present a true picture in the branch accounts. My assumption was that Fujitsu only ever injected new data for this reason. I don't think that I contemplated or considered the possibility of malicious access until many years later.

Remote access in Legacy Horizon

148. My understanding of Legacy Horizon was that there were two teams in Fujitsu, the SSC and Systems Operations Team (**SOT**), that had and exercised remote access rights. I understood that both teams could access the correspondence servers, which received live data from the branch counters in the form of messages. If, for example, the SSC discovered a system failure or BED, I understood that it may have been necessary for the SSC to inject messages into the system to correct data which had been affected by the failure or BED. To do this, my understanding was that the SSC or SOT could write a new (corrected) message that either replaced lost messages or that reversed the effect of the original (incorrect) message without deleting it. I understood that this could be done in two ways: either by injecting the new message at the correspondence servers or by injecting the new message directly at the counters.

149. When I was employed by Fujitsu, my belief was that, in practice, on the very rare occasions it arose, the SSC and SOT nearly always injected new messages at the correspondence servers, not at the counters. In my experience, messages injected at the correspondence servers would contain information that identified who had injected the message and why (usually in the form of an incident reference). However, in the civil litigation in 2018, I learned that the SSC and SOT had injected new messages at the counter more frequently than I had previously realised. I cannot now recall the detail but I remember that there were technical reasons why the messages had to be injected at the counter and not at the correspondence server. These technical reasons made sense to me at the time of the civil litigation. If messages were injected at the counter, in my experience, those messages would have attributes in the raw audit data which meant that they were identifiable as such. However, messages injected at

WITN00460200

the counter were generally more difficult to spot in the audit data than messages injected at the correspondence server.

150. My understanding was that adding a new message to the correspondence servers or to a counter did not delete the original message (the effect of which it was reversing). The original message would remain visible on the correspondence servers. I also understood that Fujitsu's archive server received a copy of all data transmitted from the branch counters to the correspondence servers, including any new messages injected at the counter. This understanding was the consequence of how Riposte was structured: it simply did not allow any data to be deleted. All that could be done was to write further messages that undid the effect of the original messages. I understood that this would leave a complete and visible audit trail on both the correspondence servers and the audit server, clearly showing all additions and modifications.

151. My understanding at the time was that there were procedures and processes that governed when and how the SSC or SOT could exercise remote access in Legacy Horizon. I never knew the detail of these procedures and processes. I only knew that they existed and I have a memory that they were tightened up over time. I do not know who was responsible for ensuring compliance with these procedures, but I believe that it was the Security team.

Remote access in Horizon Online

152. Horizon Online operated a main branch database (**BRDB**) that stored all live data. I understood that certain members of the SOT could access the BRDB, but only for the purposes of applying software updates. As far as I was aware, this type of access did not have the effect of adding, modifying or deleting any of the live data on the BRDB.

153. I understood that, with one proviso, the SSC could view but not amend live data on the BRDB. The proviso was that the SSC had access to a tool that enabled them to inject new transaction data onto the BRDB in an emergency to fix urgent system problems or BEDs. Rather confusingly, this tool was known as 'Transaction Corrections', but it was completely different to the POL function of issuing Transaction Corrections to a branch in the event of a discrepancy. During my employment by Fujitsu, using the Transaction Corrections tool was the only way that I understood that the SSC could remotely inject live data onto the BRDB. I recall learning, at the time it

WITN00460200

happened, that the SSC had used this tool once in March 2010, during the pilot of Horizon Online. I do not recall why this was necessary other than to fix a BED. I recall being told that the SSC's use of the tool on this occasion was in accordance with an audited process, but I knew nothing about what this process involved.

154. Any transactions injected by SSC using the Transaction Corrections tool would be included in the audit trail and, by virtue of the Node number (99, which is shown in the audit trail as having added the message), would be immediately identifiable.

155. As a result of my involvement in the civil litigation in 2018-2019, and now this Inquiry, I have become aware that the SSC had other tools that they could (and did) use to remotely inject live data onto the BRDB. I believe that use of the AppSup Role when accessing BRDB was one of these tools. I may have heard the name AppSup when I was employed by Fujitsu, but I did not know at that time what it was or how it worked, and I still don't understand it fully now. This probably reflects the fact that I was not an expert in the Oracle-based software that is relevant to database access in Horizon Online. In particular, I have never had detailed knowledge of precisely how this software was configured at any one time to permit or prevent remote access to different applications and to different groups of people.

156. I do not know what procedures or checks were in place at different times to control remote access by the SSC or SOT in Horizon Online. This was not my responsibility and I did not use remote access. I was only aware that there were procedures and processes that governed when and how the SSC or SOT could exercise remote access, and that when the SSC used the Transaction Corrections tool to access live data in the BRDB, this process was audited. I never knew the detail of these procedures and processes. I only knew that they existed. I also recall that, after my retirement in February 2015, I learned that improvements were made to the auditing of changes to the BRDB to make it clearer what changes had actually been made to data. I do not know who was responsible for ensuring compliance with these procedures, but I believe that it was the Security team.

157. For completeness, I should add that I was also aware that the SSC (and others) could access copies of the live data stored on a replica support database (**BRSS**). As far as I was aware, if the SSC accessed the BRSS, this did not add, modify or delete any of the live data on the BRDB. In particular, I had read-only access to BRSS but I do not consider this to be a form of remote access.

WITN00460200

Reports on Horizon Data Integrity

[Rule 9 request: questions 21 and 22]

158. The Inquiry has asked me questions about two reports I prepared concerning data integrity in Horizon.

My report on Legacy Horizon

159. (D6) [FUJ00080526] is my report dated 2 October 2009 concerning Legacy Horizon. When the Inquiry first provided me with this report, I could not remember how it had come about. I have now listened to the evidence given to the Inquiry by Dave Smith of POL, in which he describes a call in which he briefed me to prepare the report. I do not remember this call but what he told the Inquiry makes sense to me and I am content to take his word for it. Mr Smith's evidence as to why he commissioned this report appears to reflect its contents: he recalled that he wanted information on whether Horizon could be affected by power interruptions and about the security of the audit file.

160. Under the heading "Purpose", the document explained that it was a technical description of the measures built into Horizon to ensure data integrity including a description of failure scenarios. As the paper makes clear, ensuring data integrity in this context meant ensuring the integrity of data committed to the audit trail. As such at section 2, I described the process by which data was stored on hard discs within the post office counter, then at the data centre and then sealed within the audit trail. The report contemplated circumstances in which data could be lost. Section 3 of the report dealt with a series of equipment failures. I pointed out that the effects on data integrity were contingent upon whether the counter could be restarted or not and then went on to explain the different scenarios that could arise (and the checks that could be made to see whether, for example, a transaction could be completed).

161. As is clear, the paper was not a report or survey on all of the issues or BEDs which had affected Legacy Horizon. It was not a report on the integrity of Legacy Horizon in terms of it being affected by BEDs which might affect branch accounts. It was sought for a more limited purpose and its contents reflect that. I note that this report was reviewed and approved by a number of senior and technical people within POL and Fujitsu (and bears the emblems of both organisations).

WITN00460200

162. The Inquiry has asked me why the words *"without prejudice"* appear on the report. I am not a lawyer and I am not completely sure what without prejudice means in this sort of context. I don't think I would have used these words unless someone told me to write them. If someone did tell me to write them, I can't recall who that was.

My report on Horizon Online

163. (D7) [FUJ00080534] is my report dated 25 November 2011 concerning Horizon Online. I believe that there are earlier and later versions of this same report. This is reflected at section 1.1, which refers to Horizon Online having been operational for 12 months (which would have been correct at around the time of the first draft of the report in January 2011).

164. I recall that I drafted this report because Fujitsu senior management decided to obtain an independent audit from KPMG as to the integrity of the audit data produced by Horizon Online. I believe this was probably linked to press reports (in particular a BBC documentary) raising concerns about the reliability of Horizon. My report was intended to be included in the briefing materials for KPMG, so that they had a high-level technical overview as to how data was recorded in the audit trail. I recall meeting KPMG on three or four occasions to discuss the audit trail. Ultimately, however, this KPMG audit did not happen. I cannot now recall the reasons but I believe that Steven Long (who was the Post Office Account Director at the time) took the decision not to proceed with it.

165. This report was intended to address the same core issue which was addressed in my earlier sister report on Legacy Horizon. This was spelt out at section 2 ("Purpose") where I stated that *"The scope of this paper is restricted to showing the integrity of the audit trail and that it accurately reflects the transactions entered at the counter."* The scope of this report was to demonstrate that Horizon Online was configured to create an audit trail, and that this audit trail had integrity because it captured and secured all data arising from all transactions carried out by SPMs at the counter.

166. Section 3 of this report was a technical description of how data was committed to the audit trail. Section 4 was the specific part of the report dedicated to describing the storage of data within the audit system. This part of the report referred to missing or

WITN00460200

duplicate JSNs and I referred to these as potentially occurring as a result of a bug or someone tampering with the data in the BDRB. This was mentioned as potentially relevant to the audit trail.

167. Again, as is clear from the contents of this report, it was not and was not intended to be a report on the integrity of Horizon Online in terms of it withstanding or being impacted by BEDs that might affect branch accounts. It served a narrower purpose. Again, I note that it had a number of reviewers, including senior members of the Fujitsu team.

168. The Inquiry has asked me why I added the words *"legally privileged"* to the report. I cannot recall but had probably been advised to do so.

169. The Inquiry has asked me why it did not occur to me that there might be a problem with the Oracle software. As noted above, I claim no technical expertise in Oracle, but I was aware at the time that many other large companies and applications used it and I do not recall being aware of any major technical problems with it.

The Horizon audit trail

170. In this statement I have referred on a number of occasions to the audit trail or audit data created by Horizon. In his Horizon Issues judgment, Mr Justice Fraser agreed with Dr Robert Worden's description of *"audit data"* as a *"gold standard"* upon which to identify and or investigate BEDs. I also considered that the audit data was a gold standard, but Mr Justice Fraser is incorrect when he subsequently says that audit data captured every keystroke carried out by the SPM at the counter. To be clear, this is not a criticism of Mr Justice Fraser (and I note that some witnesses at the Inquiry also thought this). I hope it is helpful to expand on what I mean when I talk about an audit trail and audit data, as a basis for my belief that it had integrity. I will deal with Legacy Horizon and Horizon Online in turn.

Legacy Horizon audit trail

171. There were many hundreds of audit points in Legacy Horizon and many different ways of filtering and extracting audit data. In Legacy Horizon, I designed the mechanism on

WITN00460200

the correspondence servers which would listen to and pick up new Riposte messages (including all transactions and business events), which would then be written onto a serial file. There were tens of thousands of messages per serial file. All of those messages were then sucked into an audit server, and this became the audit trail. The audit trail contained every single message written to Riposte, i.e. a whole universe of raw data. Fujitsu developed tools that filtered this audit trail and extracted parts of it. These extracts are the ARQ data. To summarise, there were five main data groups:

a. The trace logs (also referred to as diagnostic or counter logs). These comprised specific messages output by the code to indicate the paths that were taken. The trace logs were retained for a period of days and assisted the SSC in looking at immediate problems raised by SPMs in relation to counter failures. They were not audited and were unavailable after a short period of time (probably about 30 days, but I can't remember exactly).

b. The raw audit trail (also referred to as transaction logs). This did not pick up every single keystroke at the counter, but it did comprise all "transactions" (such as cash receipts) and "events" (such as log-ins) at the branch, as well as a myriad of other things such as working data generated during balancing and cash declarations. However, it excluded "events" captured by the NT event logs (see below). The raw audit trail was retained for seven years.

c. Transactional ARQ data. This comprised all "transactions" (but not "events") extracted from the raw audit trail (see above) when required. This was retained for seven years.

d. Business Event ARQ data. This comprised all "application business events" (such as the printing of receipts) extracted from the raw audit trail (see above) when required. This was retained for seven years.

e. NT event logs. These comprised the vast majority of system errors or warnings, sometimes also called "error events", such as lock errors in Riposte. They were retained for seven years. It was part of the system design as to what the NT event logs captured.

WITN00460200

Horizon Online audit trail

172. Horizon Online sought to replicate the data extraction system for Legacy Horizon: as soon as the data centre received data, it would be written to an audit table, which would contain all transactions and all events. Horizon Online enabled the transactions to be presented in a near identical format to the ARQ data tables used in Legacy Horizon. As with Legacy Horizon, Horizon Online also produced NT event logs.

Further bugs, errors and defects

[Rule 9 request: question 49]

173. The Inquiry has asked me to consider the list of 29 BEDs identified by Mr Justice Fraser in the Technical Appendix to his Horizon Issues judgment and has asked me detailed questions about each one. Of those 29 BEDs, I have already addressed in this statement: (a) Data Tree Build Discrepancies (BED 10), (b) the Callendar Square/Falkirk bug (BED 2), (c) the Receipts and Payments Mismatch bug (BED 1), (d) the Suspense Account bug (BED 3), (e) the Dalmellington bug (BED 4) and (f) withdrawn stock discrepancies (BED 13). I will not say anything further about those six BEDs here.

174. Of the remaining 23 BEDs, it is difficult for me to answer detailed questions with confidence about what I knew or might have known. As I have set out above, the four BEDs that have always stood out in my mind as being particularly significant were the Callendar Square/Falkirk bug, the Receipts and Payments Mismatch bug, the Suspense Account bug and the Dalmellington bug.

175. That said, I wish to assist the Inquiry as best I can. My lawyers and I have reviewed the broader span of Rule 10 documents which I was given notice of (on 20 and 27 April 2023), to check whether they demonstrate that I had contemporaneous involvement in responding to other BEDs mentioned by Mr Justice Fraser in the Technical Appendix. Whilst I cannot guarantee that I have picked up every reference that may have been made to me on PEAKs and other technical documents, I have noted the following.

WITN00460200

Remming Out (BED 6) issue 2

176. The materials disclosed include a paper that I wrote entitled "Rem Misbalance" (F104) [FUJ00121073], which I had forgotten until I saw it again. I wrote this on 13 February 2007 and it was timed as written at 4.45pm. As set out in the paper, the issue that arose was a consequence of the introduction of LFS_COUNTER 35_6. This went to a limited number of branches for a pilot from 4 February 2007 to 11 February 2007 and then to the whole estate on 12 February 2007. The problem first occurred in two pilot branches over the weekend of 10 February 2007 but was not investigated until after the weekend (this was noted in the paper as being normal practice). I recorded that it occurred in 47 other branches on 12 February 2007 and that therefore 49 branches needed to have their accounts corrected. The paper was focused on the question of correction and set out that the simplest option was to send a Transaction Correction to correct the amount in Suspense and to use the Cash In Transit account to correct the files being sent to POL FS. The solutions suggested demonstrated that they had been worked out in conjunction with POL. The paper noted that LFS_Counter 35_6 was removed on 12 February but that there were 229 counters in 120 branches where the regression was unsuccessful and that they could have problems until the software was fully regressed. The process for managing such later occurrences would be the same as for the 49 branches known at the time of my paper.

177. In an email of the same day at 4.49pm (F103) [FUJ00121072], I set out that where branches had not tried Pouch Reversals, that everything should be fixed as a result of a phone call from Fujitsu support staff asking them to process a Dummy Rem (and that no Transaction Correction would be required).

178. On 14 February at 8.31am, I emailed Ian Trundell of POL (E60) [FUJ00121074] to inform him of my understanding that a fix was being prepared and that once it was prepared it would be tested and distributed in the normal way. I informed him that *"I'm not normally in the loop for such things, but presumably details of this will be communicated to POL in the normal way (whatever that is)."* This reflected my general lack of involvement in that side of the response to issues arising in Horizon.

179. In Mr Justice Fraser's Technical Appendix, he noted that KEL acha508S provided advice as to how this issue could be manually fixed; that Fujitsu ran automated reports to spot any further occurrences and fixed the issue.

WITN00460200

180. From my perspective, I was involved in producing, at speed, an immediate solution to the remming out issue. I don't think my involvement could have been for more than a few days. As far as I was aware a fix was identified and the issue resolved. In short, the problem was swiftly identified, the branches regressed to the previous version of the software and, to my knowledge, this problem did not recur.

Bureau de change problems (BED 23) issue 1

181. The PEAK for this issue is PC0129767 (E77) [POL00001264] and it appears to have arisen on 6 December 2005. It therefore relates to Legacy Horizon. Anne Chambers recorded on the PEAK (see entry dated 6 December 2005 at 2.39pm) that:

> *"Really this is user error. The transaction log search on the session id displayed 3 entries: for currency, margin, and cash. For the existing reversal, he entered the transaction id for the cash settlement only. If he had entered the transaction id for the currency/margin (they are the same), both parts would have been reversed as he intended. Reversal of a transaction settlement is always allowed but is ineffective - it just reverses cash and settles to cash, as is stated on the screen and on the reversal receipt. Net impact on anything is nil. A balance snapshot taken after this will still show both the currency having been sold, and the margin. This should have indicated that the reversal had not been done as intended. Because the PM then adjusted his stock to remove the currency which he had failed to reverse, the margin for the transaction remained (because margin is not stock). Hence he had a loss to the value of the margin. Can this be corrected via a Transaction Correction?"*

182. I agreed with this, recording (on 7 December 2005 at 11.22am) that:

> *"I'm not really clear as to why this has been raised as a PEAK. As explained above, the root cause is a user error, though it is also clear that the user documentation (which is Post Office Ltd's responsibility) could also be clearer. There are many things that we or Post Office could do (some are simple Ref Data changes as indicated earlier in this PEAK). However we cannot make any changes without guidance from Post Office. All I can suggest is that we make*

WITN00460200

Post Office aware of the analysis carried out above and ask them what (if anything) they want to do about it."

183. Looking at the PEAK now, I don't think this was an error in the system but I can understand why the user was confused and I was certainly not attaching any blame to the user. For Horizon Online, we got POL to change the reversal mechanism, meaning the mistake could no longer occur.

Remming in (BED 4) issue 1

184. I note (and it is referred to in Mr Justice Fraser's Technical Appendix), that I am mentioned in PEAK PC0203085 (F21) [FUJ00081870] in one entry of 18 August 2010 at 6.35pm where Anne Chambers recorded:

"I checked whether there were any exceptions in the BAL OSR logs for any of the messages, there was nothing. Gareth Jenkins thinks that it should not be possible to complete the rem in on both counters. Please investigate."

185. This suggests that Anne and I discussed this issue but I have no idea now what we discussed or whether it involved any of the underlying detail. It would seem to reflect the sort of question I could have been asked without needing background detail.

Recovery issues (BED 8) issue 1

186. I understand this issue to relate to the bug found in April 2010 which Fraser J referred to in his Technical Appendix. It relates to a normal (that is not failed) recovery which is recorded in an incorrect accounting period. It was fixed in early June 2010.

187. The documents recently given to me by the Inquiry include an email chain dated 26 May 2010 (F105) [FUJ00121083]. This email chain refers to KEL acha5650L, which in turn refers to PEAK PC0197769, which was raised during the pilot phase of Horizon Online. My attention has been drawn to an entry on this PEAK on 26 April 2010 in which Anne Chambers states that she has discussed the problem with me and that we were attempting to find out whether any other recovered transactions had been

WITN00460200

similarly affected. I note the observation made by Anne in her witness statement given to the Inquiry (at paragraph 135). I agree with her that this was not an instance of a failed recovery but that the recovered transaction was assigned the wrong Trading Period. I also note that there was monitoring of this issue.

188. Upon review, it appears that this was a specific bug where the accounting period had changed between the original failure and the time of recovery resulting in the recovery transactions being put into the wrong accounting period. This was a relatively straightforward bug in the pilot phase. The email reflects an agreement between Anne and me that all such cases should be identified so that POL could communicate with the SPMs.

Counter Replacement Issues (BED 12)

189. I note by reference to Mr Justice Fraser's Technical Appendix that I was mentioned in PEAK PC0052823 about a counter replacement issue [FUJ00075614]. The PEAK records that Mike Berrisford and I looked at this issue on a test rig. Again, I cannot recall much about this save that there was a fix produced in September 2002, when we made changes to the "install a replacement counter" software.

Local Suspense Account (BED 7)

190. I have been through a number of documents given to me on 20 and 27 April 2023 and I don't think that I am named in them in respect of the Local Suspense Account issue which arose in 2010. However, my attention has been drawn to an email of 29 April 2010 [FUJ00121077] in which I explain to the Post Office what the issue amounted to and the solutions for fixing it. This may well be an example of a case in which I was asked to draft a communication to POL for the purposes of explaining an issue. My email attached a spreadsheet of the 33 branches affected. My email stated that the spreadsheet showed the branches, the amounts and the settlement product used for clearance.

191. I noted that if the branch cleared the Local Suspense at the end of the period which was current, using the same method as they used in the previous period, there would be no lasting effect. The problem would not cause them a discrepancy (though it might

WITN00460200

look as though it had done so) and the Assign to Nominee and other accounts in POLFS would be correct. For branches where there was a current cash discrepancy, where the branch cleared Local Suspense to cash (or made good), this could be resolved immediately by using the Housekeeping function to clear losses / gains from Local Suspense.

192. Where branches chose more than one settlement option, it was noted that they should use the Housekeeping function before the next TP rollover, to clear the loss/gain from Local Suspense that was settled to cash. Then the remainder should be settled to Assign to Nominee (or whatever was appropriate during the TP rollover. If the branch had a new loss/gain and this was not the appropriate settlement product, they may have needed additional help.

193. I asked the MSU to pass this information to POL attaching the spreadsheet. I asked the MSU to tell POL that the NBSC had already been involved with some of the branches (those which were highlighted in yellow).

Other matters

194. As the Inquiry might anticipate, given the passage of time, this statement cannot be definitive as to my state of knowledge. I would not have recalled my involvement in, or communications about, a number of these issues and am only able to comment on them having seen the underlying material provided to me by the Inquiry. If the Inquiry wishes to me to consider any specific document or other documents, I am willing to do so.

195. There are some issues which have been considered during the examination of other witnesses and which the documents provided to me as part of the Rule 10 process touch on. I thought it would assist if I attempted to deal with these in writing.

Secrecy

196. This relates to the suggestion that there was a hesitancy or a culture within Fujitsu of not providing information to POL. I understood that there was a proper concern that there should not be informal or ad hoc communications with POL when issues arose

WITN00460200

about Horizon but that these communications should be raised through the established or formal channels (or upon myself or others being asked to provide an explanation through those channels). In April 2010, an issue was raised as to how POL had come by some information about a recovery. I noted in my reply that it was probably because I had spoken to Ian Trundell at POL and that "*I probably said more than I should, but I'm used to working openly with POL and not keeping them in the dark*" (E42) [FUJ00095095]. As noted elsewhere in this statement, I was clear that POL needed to be told about the duplication of the ARQ data and the Receipts and Payments Mismatch bug.

Duplication of ARQ Data in 2010

197. As was set out in Penny Thomas' report on duplicate data (E51) [FUJ00097058], this issue came to light on 21 June 2010 when duplicate transaction records were identified in an ARQ return. This report was written by the following day with a list of actions. In essence, the issue lay with the ARQ extraction tool. Under Legacy Horizon the audit process could result in transactions being duplicated (identically) in the audit. The Legacy Horizon retrieval process spotted this and filtered them out but the Horizon Online version didn't. None of this affected the underlying audit data. It was the extraction of the data into the ARQ spreadsheets and the possible use of those spreadsheets in legal proceedings that was of consequence.

198. By way of immediate response, Penny's report noted that as the unique identifier 'NUM' was not included in the current ARQ returns, it was agreed that this would be incorporated in the queries used to filter the records until the problem was solved. This would allow the service to continue and duplicated transaction records would be identifiable.

199. The report recorded that I agreed to draft a statement for management review detailing the issue for onward transmission to POL. It also noted that a separate issue was also identified whereby a seemingly duplicated transaction had a different NUM and that I had agreed to look at the detail of this. I believe that this is a reference to Postal Services transactions whereby multiple, identical mails items were accepted, but Postage Labels were printed for each individual item.

WITN00460200

200. In email correspondence of 24 June (E48) [FUJ00097046] at 12.39pm, Penny set out my suggested note to POL:

> *"With Horizon counters, the mechanism by which Data is audited has always worked on the principle that it is acceptable to audit the same data more than once — in particular if in doubt as to whether or not it has been previously audited successfully. The Mechanism used on Horizon to retrieve the data took this into account and only presented one instance of such duplicate data in the ARQ extracts.*
>
> *However it has recently been noticed that the HNG-X retrieval mechanism does not remove such duplicates and a quick scan of the ARQs provided to Post Office Ltd since the change to the new system indicates that about 35% of the ARQs might contain some duplicate data. A Peak has been raised to remove such duplicate data in the future. However until the fix is developed, tested and deployed, there is a possibility that data is duplicated.*
>
> *The reliable way to identify a duplicate transaction is to use the <Num> attribute that is used to generate the unique sequence numbers. Unfortunately, this attribute is not currently included in the Excel version of ARQ data that has been passed to Post Office Ltd in the past. This will be included in all future ARQs until the problem is fixed.*
>
> *Meanwhile all that can be done on existing ARQs is look for transactions that appear to be duplicates. Note that we have identified a scenario with Postal Services transactions where multiple, identical mails items are accepted (ie the Quantity button is set to greater than 1), but Postage Labels are printed for each individual item. This results in separate transactions being generated for each item, which are identical in the ARQ extracts (there is another minor difference in the raw data apart from the <Num> attribute, but this different attribute is not currently included in the ARQ extract)."*

201. In the same email Penny also provided an update on the situation (noting that a fix was expected by 29 June 219). She expressed the wish to communicate with her counterpart in Post Office about this and asked for comments by return. At 5.12pm, Guy Wilkerson emailed addressing Penny and me to tell us that Alan D'Alvarez and Geoff Butts should look at it because of the forthcoming Acceptance Board. He informed Penny that she should hold off until he had spoken to the HNG-X Team.

WITN00460200

202. On the following day, there was some further email correspondence about the viability of the fix and about Post Office not having raised the issue on a call that morning (E52) [FUJ00097070]. I emailed making clear my view that, regardless of the workaround, POL needed to be told and without delay:

> *"Geoff,*
> *The reason POL have not raised this on the call this morning is that we haven't told them there is a problem yet. On Guy's advice, Penny has not said anything to POL until you give us the say so. This was really to allow you to avoid it becoming an Acceptance Incident.*
> *The work around would require us to tell POL about the issue. Whether they feel it is acceptable is a different matter. We need to be careful, since this does relate to evidence used for prosecutions, so I feel that now we know there is an issue we do need to tell POL about it asap. (The problem was found by Penny and Alan Holmes earlier this week.)"*

203. In terms of the use to which such ARQ data may have been put, Penny's report (E51) [FUJ00097058] and emails recognised that we needed to identify which cases provided with ARQ returns since the HNG-X application had been live had progressed to prosecution and identify whether duplicate records were included. It was noted that POL's involvement was required to ensure all instances were picked up. Whilst this goes to Phase 4, I did make a witness statement in the case of Mrs Misra which identified the duplicated data in the transaction logs provided to the defence in her case.

204. In summary, this issue was acted on as soon as it came to light and the potential impact of it on spreadsheets that were provided for legal proceedings was immediately recognised. There was a response to deal with the immediate issues that this presented and the duplication was fixed. The duplication issues did not affect the integrity of Horizon or its audit data but rather the extraction of this data into the ARQ spreadsheets.

Lepton Branch

205. I dealt with Helen Rose in 2012-2013 in relation to a discrete issue which had arisen at the Lepton branch and which she was investigating. The SPM had complained about an in-branch reversal which had occurred during system recovery and which he

WITN00460200

denied making. I investigated this issue and confirmed the SPM's position – the reversal had been system generated. I am conscious that Mr Justice Fraser was critical when in email correspondence with Ms Rose on 30 January 2013 [POL00097444] I said that the *"system was behaving as it should".* When I made this comment, I had entirely accepted that the SPM was not at fault. But the reason why the system was behaving as it should was because when the system failed unexpectedly, it was not viable to wait until it restarted before deciding what (if anything) was owed by/to the customer. For that reason, the design of the system made certain assumptions as to what transactions were deemed to have completed and what were deemed to not have completed and this was documented. But it was clearly a complex and rare scenario and I understood that SPMs could get this wrong. My words were not intended in any way to minimise the situation or attribute blame. They were only intended to convey that from the technical perspective, what happened accorded with the design intention.

Reflections

206. I wish to make clear that I have learned a great deal through the documents provided to me by the Inquiry and by the evidence that I have heard from the SPMs and others. I am conscious of the gap which existed between my work on Horizon and my belief that overall, it worked for the great majority of branches, and then the reality experienced by those who have given evidence to the Inquiry. Reflecting on this and the broader evidence in the Inquiry, I am struck by the lack of support that was afforded to SPMs when they got into difficulties. I agree that the time out waiting for lock issues ought not to have persisted for as long as they did. I can also see, at this distance, that there would have been a real benefit in having a team whose function it was to have sight of all issues across the different levels of support and who could have drawn together PEAKs, had oversight of what front-line support were fielding, understood the practical ramifications of any issues upon those who worked in the branches and who ensured that monitoring was working. I have other observations which are pertinent to phase 4 but will not develop these here.

207. I am aware that the Chair has said that he will keep under review the question of whether to approach the Attorney General for an undertaking. I maintain my application to him in relation to Phase 4 of the Inquiry. I understand that this will be raised in correspondence on my behalf.

WITN00460200

Statement of Truth

I believe the content of this statement to be true.

Signed: GRO

Dated: 01/06/2023

WITN00460200

INDEX TO SECOND WITNESS STATEMENT OF GARETH IDRIS JENKINS

No	URN	Document Description	Control Number
1	FUJ00079301	Email from Gareth Jenkins to John Pope and Roger Donato (cc others) re: EPOSS Reconciliation	POINQ0085472F
2	FUJ00079303	Email from Gareth Jenkins re EPOSS Reconcilliation Design Version 0.7 21/9/99	POINQ0085474F
3	FUJ00079304	Email - subject 'Urgent re EPOSS Reconciliation Design Version 0.7	POINQ0085475F
4	FUJ00079488	Pathway Change Proposal re Cash Account Derivation	POINQ0085659F
5	FUJ00079333	Email, subject Current Issues on CI4 EPOSS	POINQ0085504F
6	FUJ00126038	Email from Chris Allen to Ann Clarke, Ben Gildersleve, Clive Read and others regarding an invite to the branch trading - treatment of suspense (18 Feb 13:00 in room F34)	POINQ0132251F
7	POL00038916	Draft version of Impact Release 3 Design Proposal®Version 1	POL-0035398
8	FUJ00090393	IMPACT Release 3 Design Proposal	POINQ0096564F
9	FUJ00090060	IMPACT Release 3 Design Proposal	POINQ0096231F

WITN00460200

10	FUJ00126035	Second Corporate Statement of Fujitsu Services Limited	FUJ00126035
11	POL00028867	Peak Incident Management System Report	POL-0025349
12	POL00030523	Fujitsu's PEAK Incident Management System log of customer call raising issue of accounting discrepancies and data tree build issues	POL-0027005
13	FUJ00086360	Fujitsu Test Report for Functional and Non-Functional Integration Testing of BI3 S80 v0.1	POINQ0092531F
14	FUJ00086363	Fujitsu SV&I High Level Test Plan for S81 v0.1 with comments from Gareth Jenkins	POINQ0092534F
15	FUJ00086368	Minutes of Fujitsu "Prayers" meeting on 16/08/05	POINQ0092539F
16	FUJ00086490	PEAK PC0146170	POINQ0092661F
17	FUJ00083564	Email from Gareth Jenkins to John Balletyne re: PC0057478	POINQ0089735F
18	FUJ00083548	Email from Gareth Jenkins to Mark Jarosz and Brian Orzel re: PC0057478	POINQ0089719F
19	FUJ00083564	Email from Gareth Jenkins to John Balletyne re: PC0057478	POINQ0089735F
20	FUJ00083544	Email from Mark Jarosz to Gareth Jenkins re: PinICL PC0056922	POINQ0089715F

WITN00460200

21	FUJ00083582	Email from Mark Jarosz to Gareth Jenkins re: PC0057957	POINQ0089753F
22	FUJ00083568	Email from Gareth Jenkins to Mark Jarosz re: PC0057957	POINQ0089739F
23	FUJ00083582	Email from Mark Jarosz to Gareth Jenkins re: PC0057957	POINQ0089753F
24	FUJ00075544	Peak Incident Management System - Copy PC0057957 FAD260801 - Timeout occurred waiting	POINQ0085136F
25	FUJ00083600	Email from Gareth Jenkins to Mark Jarosz re: PC0065665	POINQ0089771F
26	FUJ00083621	Email from Gareth Jenkins to Mark Jarosz re: PC0075892	POINQ0089792F
27	FUJ00083634	Email from Gareth Jenkins to Mark Jarosz re: PC0087709	POINQ0089805F
28	FUJ00083640	Email from Gareth Jenkins to Brian Orzel, Simon Fawkes and Barbara Longley re: PC0087709	POINQ0089811F
29	FUJ00083592	Email from Chris Wannell to Gareth Jenkins and Mark Jarosz re: PinICL Client Call Summary list 21 March 2001	POINQ0089763F
30	FUJ00083596	Email from Brian Orzel to Gareth Jenkins re: priority of Escher-Dev PinICLs	POINQ0089767F
31	FUJ00083712	Email chain between Anne Chambers, Gareth Jenkins and Penny Thomas on Callender Square bug	POINQ0089883F

WITN00460200

32	FUJ00083721	Email from Anne Chambers to Gareth Jenkins re: Callendar Square bug	POINQ0089892F
33	POL00028911	Management report summary re: Callendar Square / Falkirk	POL-0025393
34	FUJ00081602	Fujitsu Portable Appliance Testing Branch Roll Out v1.1	POINQ0087773F
35	FUJ00081220	Spreadsheet with data of receipts and payments mismatch with affected branches	POINQ0087391F
36	FUJ00083353	Report by Gareth Jenkins on Correcting Accounts for "Lost" Discrepancies	POINQ0089524F
37	FUJ00082443	Email from Mark Wright to Gareth Jenkins re: the receipts and payments mismatch issue	POINQ0088614F
38	FUJ00081135	Email from Mark Wright to Mike Stewart and Steve Bansal re: receipts and payments	POINQ0087306F
39	FUJ00081211	Email from Gareth Jenkins to John Simpkins and Mark Wright re: branches affected by the receipts payments and discrepancies issue	POINQ0087382F
40	FUJ00081137	Email from Mike Steward to Gareth Jenkins, Mark Wright and Steve Bansal re: receipts and payments mismatch issue	POINQ0087308F
41	FUJ00081584	Receipts/Payments Mismatch issue notes	POINQ0087755F

WITN00460200

42	FUJ00081527	Report by Gareth Jenkins on the Receipts and Payments Mismatch	POINQ0087698F
43	FUJ00081544	Email from Gareth Jenkins to Rod Ismay, Will Russell and Antonio Jamasb re receipts and payments call- questions	POINQ0087715F
44	FUJ00081545	Email from Will Russell re receipt and payments issue	POINQ0087716F
45	FUJ00081550	HNG-X System: Receipts and Payments Mismatch (25 Feb 2011)	POINQ0087721F
46	FUJ00081584	Receipts/Payments Mismatch issue notes	POINQ0087755F
47	FUJ00081214	Email from Mike Woolgar to Mark Wright, Gareth Jenkins and John Simpkins re: receipts and payments issue	POINQ0087385F
48	POL00029718	Email chain between Steve Parker, Mark Wright, Andrew Winn, Emma Langfield and Gareth Jenkins Re: ISSUE - Receipts & Payments mismatch	POL-0026200
49	POL00098016	Email from Andrew Winn to Steve Bansal dated 16/04/13 re: investigation into receipts and payments problem in 2010	POL-0097599
50	POL00029719	Email chain between Gareth Jenkins, Emma Langfield, Rod Ismay, Paul Dann, Simpkin John, Mark Wright, Pete Newsome, Andrew Winn®Re: Branches affected by Receipts Payments and Discrepancies Issue®	POL-0026201

WITN00460200

51	FUJ00081214	Email from Mike Woolgar to Mark Wright, Gareth Jenkins and John Simpkins re: receipts and payments issue	POINQ0087385F
52	FUJ00083375	Note authored by Gareth Jenkins titled 'Local Suspense Problem' v0.5	POINQ0089546F
53	FUJ00085882	Report by Gareth Jenkins on Duplicate Remittances	POINQ0092053F
54	FUJ00083379	Fujitsu presentation on Branch Outreach Issue (Initial Findings)	POINQ0089550F
55	FUJ00126035	Second Corporate Statement of Fujitsu Services Limited	FUJ00126035
56	FUJ00086720	PEAK PC0208335	POINQ0092891F
57	FUJ00081770	Peak Incident Management System - PEAK PC0187425	POINQ0087941F
58	FUJ00080537	Report: report regarding the apparent loss of audit trial of Giro Payments	POINQ0086708F
59	FUJ00080536	Email from Andy Winn to Gareth Jenkins and Craig Tuthill re Giro Payment Report	POINQ0086707F
60	FUJ00080526	Fujitsu Report: Horizon Data Integrity v1.0	POINQ0086697F
61	FUJ00080534	Fujitsu Report: Horizon Online Data Integrity	POINQ0086705F
62	FUJ00121073	Rem Misbalance Report	POINQ0127265F

WITN00460200

63	FUJ00121072	Email from Gareth Jenkins GI to Gary Blackburn. cc'd Mike Stewart and Anne Chambers re: Rem Misbalance	POINQ0127264F
64	FUJ00121074	Email from Gareth Jenkins to Ian Trundell, from a forwarded email from Gary Blackburn re: RE: Rem Misbalance	POINQ0127266F
65	POL00001264	PEAK - PC0129767 - reversals of foreign currency transactions	VIS00002278
66	FUJ00081870	Peak Report RE: FAD126109 Pouch remmed in on two counters at same time	POINQ0088041F
67	FUJ00121083	Email from Claire Drake to Anne Chambers, cc'd Gareth Jenkins and Steve Parker re: Recovery into wrong TP/BP	POINQ0127275F
68	FUJ00075614	Peak Incident Management System - Lost transaction following SCO replacement®®	POINQ0085212F
69	FUJ00121077	Email from Gareth Jenkins to Barry Evans, cc'd Will Russell, Mark Andrews and others re: BTS issues 14/04/10	POINQ0127269F
70	FUJ00095095	Email from Gareth Jenkins to Graham Allen re: BIMs process	POINQ0101266F
71	FUJ00097058	Report called "Duplication of transaction Records contained in ARQ returns - 22 June 2010 by Penny Thomas	POINQ0103229F
72	FUJ00097046	Emails between Guy Wilkerson, Geoff Butts & Alan	POINQ0103217F

WITN00460200

		Alvarez re duplication of transaction records	
73	FUJ00097070	Email chain between Gareth Jenkins, Geoff Butts, David Cooke and others, RE: ARQs	POINQ0103241F
74	FUJ00097058	Report called "Duplication of transaction Records contained in ARQ returns - 22 June 2010 by Penny Thomas	POINQ0103229F
75	POL00097444	Email chain from Helen Rose to Gareth Jenkins re: transaction log	POL-0097027