WITN09870100

Witness Name: Gerald Barnes

Statement No.: WITN09870100

Dated: 30 August 2023

POST OFFICE HORIZON IT INQUIRY

FIRST WITNESS STATEMENT OF GERALD BARNES

I, *MR GERALD BARNES*, will say as follows:

INTRODUCTION

1. I am currently employed by Fujitsu Services Limited ("**Fujitsu**") as a Software Developer, a position I have held since 1998.

2. This witness statement is made to assist the Post Office Horizon IT Inquiry (the "**Inquiry**") with the matters set out in the Rule 9 Requests provided to Fujitsu on 16 June 2023 and 31 July 2023 (together, the "**Requests**"), to the extent I have or had direct knowledge of such matters. I was assisted in preparing this statement by Morrison Foerster, who represent Fujitsu in the Inquiry.

3. Many of the topics set out in the Requests relate to events that occurred between 5 and 15 years ago. In the limited time available, I have set out my best recollection of these events in this statement, which relate to my work in Fujitsu's audit team and the processes relating to audit queries (also known as Audit Retrieval Queries or "**ARQs**"). I have tried to remember the

WITN09870100

detail of relevant events to the best of my ability, however, there are areas where my recollection is unclear or limited.

4. Where I have referred to documents to assist my preparation of responses to the Requests, the URNs of the relevant documents are set out in this statement.

PROFESSIONAL BACKGROUND

5. I joined Fujitsu's Post Office Account team ("**Post Office Account**") as a developer in June 1998 or thereabouts. My very first job was to look after a database of reports produced by Post Office clerks. After that, I became involved in supporting the Electronic Point of Sale Service ("**EPOSS**") software for transacting at the counter and balancing, as well as looking after related reports.

6. I became interested in balancing and took an evening class in bookkeeping attaining 'OCR Accounting 2 Pass' and 'Pitman Accounting Level 3 Pass' accreditations. I remember writing a component called 'Operation Launch' to facilitate the use of debit and credit cards which was being introduced in the earlier version of the Horizon system (known as "**Legacy Horizon**"). I also remember writing the migration software that enabled a counter to transition from using Escher's Riposte software platform to the new system (known as "**HNGx**" or "**Horizon Online**"). Because of this piece of work, I believe I was the last member of the EPOSS Riposte team, which was a large team during the time of Legacy Horizon.

7. In 2009 or thereabouts, whilst supporting the migration software for the remaining counters to transition to HNGx, I also started looking at the audit

WITN09870100

system in HNGx, which was a completely new area for me. It was around this time that I then joined the audit team. I recall that there was already an audit system in Legacy Horizon for Riposte that I knew little about then. When I joined the team, this audit system was being rewritten as part of the transition to HNGx. For this reason, I have limited experience and knowledge regarding the systems and processes relating to audit and ARQs in relation to Legacy Horizon.

8. In 2012, I wrote much of the migration software to move all the audited data files from Legacy Horizon's audit archive (Centera) to Fujitsu's audit archive for HNGx (Eternus). Centera and Eternus are specialist computers designed for long term storage of files and store each file in multiple locations in case of a fault in one location. Centera was provided by a third-party company, while Eternus was an in-house solution.

9. In 2020, there was a large transfer of jobs and many of the technical functions of the Post Office Account were moved to teams in India, and I was the only one left in the audit team in the UK.

10. In recent years, I have been involved in moving the audit system into the AWS (Amazon Web Services) Cloud and gained "AWS Certified Security – Specialty" accreditation.

THE AUDIT QUERY PROCESS AND THE ARQ SPREADSHEET

The ARQ Spreadsheet

11. In the Requests, the Inquiry has asked a number of questions in relation to an "**ARQ Spreadsheet**" (LCAS0001383, pages 17 and 18), which I understand to be filtered ARQ data for the Marine Drive Post Office that was

WITN09870100

provided to Mr Lee Castleton by way of disclosure in the civil proceedings commenced by Post Office Limited ("**POL**") against Mr Castleton (*POL v Castleton*). I have been made aware that the ARQ Spreadsheet forms part of a larger document. However, for the avoidance of doubt, in preparing this statement, I understand that the Inquiry requires me to consider only pages 17 and 18 of LCAS0001383 and does not require me to consider the remaining pages of the document.

12. The ARQ Spreadsheet appears to have been prepared before I joined the audit team in around 2009, and I was not involved in responding to any ARQs that POL sent to Fujitsu in relation to *POL v Castleton*. Before I joined, the audit team generated these types of spreadsheets using 'RQuery' ("**RQuery**"), which was an Escher tool supplied specifically for the purposes of querying a regenerated message store that contained messages in Riposte Attribute Grammar format. The ARQ Spreadsheet would have been generated by RQuery.

13. As I mention above at paragraph 7, when I joined the audit team, Fujitsu was changing from using Escher's Riposte software to Fujitsu's own bespoke software, HNGx. Following this change, the audit team moved from using RQuery to perform ARQs to 'XQilla' ("**XQilla**"), which is an open-source software program that queries files conforming to Extensible Markup Language ("**XML**") standards. The audit team also wrote a program known as the 'Query Manager service' which converts files using Riposte Attribute Grammar to XML so that the files could be queried. The audit team continues

WITN09870100

to use XQilla (run from the Query Manager service) to run audit queries today.

14. When I joined the audit team in around 2009, I recall that the developed software, (which included the Query Manager service and changes to the audit extractor client) was starting to be prepared for release. For this reason, I did not use RQuery nor am I familiar with the audit query processes that were used to produce the ARQ Spreadsheet. I am, however, familiar with XQilla and the audit query processes that have been used by the audit team since approximately 2009.

15. At the time, my understanding was that Fujitsu wanted to change the software from RQuery to XQilla (i) because HNGx, which was replacing Legacy Horizon, did not use Escher's Riposte software and produced XML output files, and (ii) to stop paying licensing fees to Escher. RQuery was designed to query Riposte message stores, which were no longer going to be used with HNGx. The Query Manager service, which was a key new component of HNGx, parsed the Transaction Management Service ("**TMS**") and Branch Database (BRDB) files and then queried them with XQilla. The audit retrieval processes relating to the period following the change to XQilla are described in the "Audit Extraction Client User Manual" (see, for example, FUJ00158710), the "Audit Data Retrieval High Level Design" (see, for example, FUJ00123758), and the "Audit Data Collection & Storage High Level Design" (see, for example, FUJ00123759).

WITN09870100

The audit query process

16. In relation to HNGx, the process of generating a spreadsheet of transactions (similar to the ARQ Spreadsheet) is as follows:

    a. Files to be audited are placed on many 'shares' across the estate. A share is a folder of a computer that is accessible by another computer.

    b. 'Gatherers' on the audit server bring the files into the audit server, where they are stored on a special long term storage device (known as an audit archive — Centera to begin with, which was later replaced by Eternus) and indexed using a Structured Query Language ("**SQL**") database on the audit server. A checksum of the file is also stored too (a checksum is effectively a unique numerical identifier that is allocated to a file).

    c. A special tool on audit workstations can then be used to display stored files and retrieve them. As these stored files are retrieved, their checksum is checked. Some of the stored files are files of transactions, and extra software is available to generate spreadsheets of transactions.

17. While I am not familiar with RQuery, I have reviewed the "Audit Data Retrieval High Level Design" relating to Legacy Horizon (FUJ00117508), and the process used to generate spreadsheets such as the ARQ Spreadsheet would have included the following steps:

    a. the file of Riposte transactions was imported by 'NWB_TMS_Generate.exe';

    b. this program in turn called 'AgentLoadAuditData.exe' to generate a Riposte message store; and

WITN09870100

c. RQuery was then used to process the message store and produce the spreadsheet.

THE HEADINGS USED IN THE ARQ SPREADSHEET

18. In the Requests, the Inquiry has asked for an explanation of the various headings in the ARQ Spreadsheet. For the reasons I explain above at paragraph 14, I did not personally use RQuery to produce spreadsheets in relation to audit queries. Having reviewed the ARQ Spreadsheet, however, there appear to be similarities in the headings used in the spreadsheets that were generated using RQuery and XQilla. I have also reviewed the following documents to assist my understanding of the headings listed in the ARQ Spreadsheet:

a. "High Level Design Specification for Track and Trace (T&T) Agents" version 4.0 DE/HLD/015 dated 27 August 2008 (FUJ00171843);

b. "High Level Design Specification for Agents for NBX, the NBE Replacement" version 2.0 NB/HLD/017 dated 13 March 2006 (FUJ00171845);

c. "Management of the Prosecution Support Service for Audit Record Queries" NB/PRO/003 dated 20 November 2007 (FUJ00122457);

d. "TPS Agents for BI3: High Level Design" (AD/DES/041) version 7.0 dated 9 August 2005 (FUJ00090934);

e. "EPOSS Attribute Grammar Catalogue" EP/DES/002 version 6.0 dated 12 December 1997 (FUJ00079219);

WITN09870100

f. “Riposte Attribute Gramme Catalogue – Messages” version 0.1 EP/LLD/019 dated 19 July 1999 (FUJ00171846); and

g. “OBCS Counter Component Design Report” OB/DES/0004 version 2.0 dated 15 December 1997 (FUJ00171847).

19. I set out below my understanding of the various headings in the ARQ Spreadsheet. While I have tried to explain these headings based on my experience with XQilla and the documents noted above at paragraph 18, there are certain headings that I am unable to explain in detail. In addition to this, where I can, I have explained my understanding of what the numeric values in the ARQ Spreadsheet mean, which were entered into the message store.

a. The first heading ‘GroupId’ is a six-digit code identifying the collection point, in other words, corresponding to a particular Post Office branch. I understand that the GroupId in the ARQ Spreadsheet relates to the Marine Drive Post Office. The heading ‘Id’ denotes the number of the counter used in a Post Office branch.

b. The headings ‘Date’ and ‘Time’ relate to the date and start time of a given transaction. The heading ‘User’ records the ID of the Horizon system user logged in.

c. The heading ‘SU’ relates to the stock unit reference. My understanding is that there are stock units that the Horizon system user can attach to, and each transaction made by the user records this stock unit. Post Offices are required to check stock in the branch every month for each

WITN09870100

stock unit used before going into a new balance period or cash account period. This process is often called 'rolling over'.

d. The headings 'EPOSSTransaction.T' and 'EPOSSTransaction.Ti' relate to the event description and event result respectively.

e. The heading 'SessionId' is a unique identifier for all transactions within a customer session. I understand a session contains all the transactions that a Post Office branch carries out before settling. It appears that all SessionIds are prefixed with "44-" which identifies the Horizon system, followed by the GroupId and Id as described above. The following number relates to the last message committed to the message store before the start of the customer session. The final number is a uniqueness factor, starting at 1.

f. The heading 'TxnId' is a unique identifier for all messages within a customer transaction using a similar algorithm as the SessionId. I understand that the final number in a TxnId advances one at a time for each transaction in the session.

g. The heading 'Mode' is populated with 'SC' which stands for 'Serve Customer' which is the most common mode of using the Horizon system but there are others, for example, RV for a reversal of a transaction.

h. The heading 'ProductNo' is the product reference number. Each product has a different product number, for example, first class stamps and second class stamps each have different product numbers.

WITN09870100

i. The heading ‘Qty’ records the quantity of products, a positive value meaning stock unit leaving the Post Office branch. The heading ‘SaleValue’ records the monetary value of a transaction, and similarly, a positive value means leaving stock unit.

j. The heading ‘EntryMethod’ records the way in which the data is captured:

    i. 0 = barcode,

    ii. 1 = manual,

    iii. 2 = magnetic card,

    iv. 3 = smart card, and

    v. 4 = smart key.

k. Relatedly, the heading ‘State’ records the method of manually keyed entries (where, for example, 4 means encash and 5 means non-barcode).

l. The heading ‘IOP_Ident’ is the ‘Instrument of Payment’ number.

m. The heading ‘Result’ is the result of the order book transaction where:

    i. 1 = OK,

    ii. 2 = IMPOUND,

    iii. 3 = UNREAD (i.e., unreadable), and

    iv. 4 = INVALID.

WITN09870100

20. The heading 'ForeignIndicator' indicates whether an Order Book Control Service (OBCS) payment was made at a local Post Office or foreign outlet where 0 means local and 1 means foreign. The foreign indicator defaults to 0 for all manually keyed entries. In practice, what I understand by this is if someone normally collects their benefit payment from a particular Post Office, this would be described as local. However, if they collect the payment from another Post Office, that would be foreign.

THE RELIABILITY AND ACCURACY OF THE ARQ SPREADSHEET

21. In the Requests, the Inquiry has asked whether the data that is presented in the ARQ Spreadsheet has been reliably and accurately parsed from the original raw form to the way in which it is presented on the ARQ Spreadsheet. In relation to audit, I would understand 'original raw data' to mean the files of transactions on the audit archive.

22. One way to test whether the data that is presented in the ARQ Spreadsheet, which appears to relate to 2 February 2004, has been reliably and accurately parsed from the original raw data, would be to generate a new spreadsheet with different software that was not used to produce the ARQ Spreadsheet (i.e., XQilla) and compare its results with the ARQ Spreadsheet. Since I have joined the audit team, it is a very common practice to rerun ARQs after any release likely to affect ARQs and check that the spreadsheets generated before and after the release provide the same results. Such checks are performed in all environments — development, live system testing (LST), and live.

WITN09870100

23. Unfortunately, however, I understand that the message store and audit trail data relating to the Marine Drive Post Office branch for 2 February 2004 is no longer available. Without this data, it is difficult to test and confirm whether the data was reliably and accurately parsed onto the ARQ Spreadsheet.

EXPECTED REPORTING IN LOG FILES WHILST HORIZON IS OFFLINE

24. In the Requests, the Inquiry has asked what the expected reporting in the log files should be when transactions have taken place whilst the system was offline. I am taking log files to mean TMS files put on shares for the audit system to gather. I am aware of the processes in place in relation to the audit servers going offline following the implementation of HNGx, however, I am not able to provide information on the processes that are in place where other parts of Horizon (e.g., the harvester agents) go offline. For the reasons I explain above at paragraph 7, I am not aware of the processes that were in place in relation to the audit servers going offline during Legacy Horizon. However, having reviewed the "Audit Data Retrieval High Level Design" relating to Legacy Horizon (FUJ00117508), it would appear the audit system operated in a similar way during Legacy Horizon.

25. The processes relating to the audit servers going offline in relation to HNGx are as follows:

    a. There are two audit servers that gather files from shares and delete them a configurable time later. If either audit server is offline, that audit server is not able to gather files from the shares it monitors. The "Archive Server Configuration Information and Operations Notice" lists all shares that

WITN09870100

have been audited and a description of what files are in those shares (see, for example, FUJ00089176).

b. While the audit server is offline, extra files may be added to the shares. When the audit server comes back online, it will gather the files and save them to Centera (which was later replaced by Eternus, as I explain in paragraphs 8 and 16 of this statement).

c. Similarly, in the evening, files are copied from one server to the other using a process called 'robocopying', (short for robust copying), which is a more resilient process of copying that is more suitable for batch (unattended) operations. No files get missed out; they are just robocopied late. The audit team advise that all audit shares should have sufficient capacity to cope with the audit server being offline for five days.

d. When the audit team performs audit retrievals, it allows extra days in the ARQ date range to allow for files gathered late. The spreadsheets generated as a result of the query (e.g., the ARQ Spreadsheet) would not indicate whether any files were gathered late. The Query Manager service is the current system that checks that there are no missing transactions (each transaction is numbered incrementally) before it generates the spreadsheet.

ISSUES AND INCIDENTS RELATING TO ARQS

26. In the Requests, the Inquiry has asked me to confirm whether I am aware of any incidents where an audit log (whether ARQ log, a log produced by RQuery or XQuilla, detail from the ARQ interface or equivalent) has been provided to POL or Royal Mail for court or disciplinary proceedings or in an

WITN09870100

investigation relating to a postmaster, manager or assistant that was, or may be, unreliable. I understand 'ARQ log' to refer to the data that was provided by Fujitsu to POL in response to ARQ requests which sought data from the audit servers.

27. During my time at Fujitsu, I was not personally involved in responding to ARQs that were submitted by POL to Fujitsu in relation to investigations, court proceedings, or disciplinary proceedings against postmasters, managers or assistants. I was, however, involved in two issues that were identified that could affect the accuracy of ARQ data: (i) an issue relating to the 'CABSProcess' that could cause potentially incorrect data to be presented to the audit system, which was recorded in Peak PC0152376, and (ii) issues relating to 'duplicate transactions' being presented to the audit system.

28. I am aware the CABSProcess issue may have affected the accuracy of the data that was transferred from the message store to the audit archive. I am also aware that the duplicate transactions issue did cause inaccurate reporting of duplicate transactions in audit logs (i.e., the spreadsheets produced as part of the ARQ process). The CABSProcess and duplicate transactions issues could have resulted in unreliable audit logs being provided to POL or Royal Mail for a particular branch. However, I am not aware of any incidents where Fujitsu provided POL or Royal Mail an audit log for an investigation, court proceeding or disciplinary proceeding against a postmaster, manager or assistant, that was or may have been unreliable, including in relation to the CABSProcess and duplicate transactions issues.

WITN09870100

Peak PC0152376

29. Peak PC0152376 (FUJ00154684) concerned an incident that was reported in December 2007 relating to a feature in the CABSProcess (which ran at 7.00pm every evening) that sometimes caused other components to fail silently. Basically, after a particular fix, the CABSProcess wrote its transactions atomically, which meant it created a lock on the message store for a period.

30. The CABSProcess issue was unrelated to audit—while the files of transactions that the audit system stored were not complete, this did not mean that the audit system was not properly recording what happened at the counter.

31. The fact that the failure was silent was really bad error handling. Good programing practices would be to abort (i.e., for the code to stop running) with a clear error message. It is better to produce no results than incorrect results, and good error handling should be coded from the start. However, my understanding is that in Peak PC0152376, an error was written to the audit log and then processing continued, so although the operator at the Post Office branch would not know anything had gone wrong, a detailed analysis of the audit log after the event would have revealed the problem.

32. I cannot recall the specific incident and the fix that was applied. I have reviewed Peaks PC0152376 (FUJ00154684) and PC0164429 (FUJ00155366) (which was the Peak that the fix was attached to). It appears that I proposed a solution to the issue, and it was initially deemed unnecessary. The fix was later applied on or around 25 September 2008,

WITN09870100

and I was involved in applying the fix (FUJ00155266). Following the fix, I was involved in undertaking manual checks of event logs in January 2009 (see FUJ00154836 and FUJ00155402). I can recall performing manual checks of event logs, which I often did as part of my work on the audit team. However, I cannot now recall being aware at the time that any specific checks related to the CABSProcess issue.

33. The CABSProcess issue highlighted a problem that could easily be caused by another system process at any time of day. In retrospect, error handling should have been tightened generally. For example, when I wrote the software to migrate from Legacy Horizon to HNGx, I kept this in mind. The postmaster pressed the migration button which appeared on migration day and if anything went wrong the postmaster got a message displayed saying something to the effect of: “An error has occurred please contact the help desk”. The program then stopped further processing and detailed evidence was recorded that would enable the help desk to identify the issue (possibly after escalating the issue to me). In my opinion, this sort of error handling is the safest. When something goes wrong everyone knows about it immediately and nothing is committed — in this case, the Post Office branch would not be migrated and needed to continue using Legacy Horizon a bit longer. The problem could then be analysed and if necessary, the code could be fixed and a new release produced with the problem eliminated. Then the postmaster could press the migration button on a later chosen day.

WITN09870100

Duplicate transactions

34. The duplicate transactions issue concerned duplicate transaction messages not being reported as part of the ARQ process. As a result of this issue, multiple instances of one transaction could appear on a spreadsheet generated as part of the ARQ process, and it would not be clear that they were the same transaction. An example of a Peak relating to the duplicate transactions issue is Peak PC0205805 (FUJ00171848) opened on 27 October 2010.

35. The issue is not a fault in the audit software as such—the issue was that duplicate transaction messages were being generated. In my experience, this usually occurred because TMS files sometimes had more than one copy of a transaction in different files, and it was a bug in the harvester that caused this.

36. As recorded in Peak PC0205805 (FUJ00171848), the issue affected 'Fast ARQs', which refers to a software tool that replaced 'Slow ARQs'. I recall the differences between Fast and Slow ARQs being as follows:

a. Slow ARQs: The user performing the audit query would supply a Post Office branch code (known as a "**FAD**") to the relevant software and set off a query to retrieve the files. Once the files were retrieved, the user would cause them to be filtered for the FAD, at which point duplicates and missing transactions were shown on the screen. The user would then select a query to be run. The spreadsheets generated as a result of the Slow ARQ process would show duplicate transaction messages and missing transactions. Peaks PC0205353 (FUJ00171892) opened

WITN09870100

on 14 October 2010 and PC0207787 (FUJ00171894) opened on 18 January 2011 indicate that this correctly working functionality for HNGx was not completely in place initially but was implemented soon after.

b. Fast ARQs: Fast ARQs automated the Slow ARQ process in a one-step operation. Fast ARQs did not initially display duplicate transaction messages and missing transactions during the processing of the audit query. As a result of Peaks PC0205353 (FUJ00171892) and PC0207787 (FUJ00171894), the spreadsheet generated was modified to show the correct missing transactions and duplicates. In the final design of Fast ARQs, that has been running for many years, missing transactions (gaps in the message run) cause the Fast ARQ to fail with an error report. Duplicates, however, do not cause the spreadsheet generation to fail, the duplicates are just reported in the resulting spreadsheet. This is because whenever duplicates were checked they were always exactly that — duplicates in every respect, including the message numbers. Hence, duplicates were not considered an error, but it needed to be clear that there were transactions listed twice in the spreadsheets.

37. According to Peak PC0205353 (FUJ00171892), the issue was fixed in or around November 2010. As a result of the fix, duplicate messages were reported in the spreadsheets generated as part of the ARQ process. In addition, duplicate messages were logged in the generated log files.

WITN09870100

Potential further incidents

38. Further to the matters outlined above, I have been pointed to a number of additional Peaks in which I am referenced. Below I have listed the Peaks that I consider may have had an impact on the audit log:

    a. Peaks PC0152828 (FUJ00155211) and PC0153009 (FUJ00155224): The latter of these Peaks is a clone of the former, each recording a Horizon stock unit rollover issue that I investigated in January 2008. This incident appears to have indicated deficiencies in the rollover error handling process and highlighted a risk that this could affect the figures presented by audit.

    b. Peaks PC0225071 (FUJ00173057) and PC0225656 (FUJ00172286): The latter of these Peaks is a clone of the former, each of which were opened on or around 16 April 2013. By way of background, the SQL database on the audit server is shut down each evening so that a backup of the system can be taken. It is necessary for the Query Manager service, which needs the SQL audit server to function, to cope with this backup process. A code issue in relation to the backup process was identified and a fix was tested and subsequently deployed in or around November 2014. All ARQ spreadsheets generated when this issue might have occurred were checked to determine whether they were affected (I note, however, that many of the ARQ spreadsheets produced would have been run outside of the period when SQL server was in shutdown in any event).

WITN09870100

c. Peak PC0272681 (FUJ00173183): This Peak was opened on 1 August 2018 and relates to a Slow ARQ which was run during the SQL audit server evening shut down. The Peak records me noting that a release is required to improve error handing, although those working on the prosecutions had not complained about this issue. My view is that this issue is unlikely to have had financial implications.

d. Peak PC0276698 (FUJ00173185): This Peak was opened on 12 February 2019 and relates to failed ARQ queries which, when rerun, were then successful. It appears that this was an intermittent issue which did not occur consistently. My view is that this issue is unlikely to have had financial implications.

INFORMATION CONTAINED IN ARQ LOGS

39. In the Requests, the Inquiry has asked me to confirm whether in my view (i) the ARQ logs that were provided by ICL/Fujitsu to POL or Royal Mail in respect of court proceedings involving issues relating to transactions from Legacy Horizon and Horizon Online were sufficient to enable a postmaster, manager or assistant to understand whether the system was operating correctly, and (ii) the sources of information from which the ARQ logs obtained were sufficient.

40. The ARQ logs (e.g., the ARQ Spreadsheet) do not provide enough information to enable a postmaster, manager or assistant to understand whether there is an issue with the Horizon system. The ARQ log is a report of transactions as presented to the audit system and it does not contain information based on checks for problems with the Horizon system. The

WITN09870100

primary way of checking for systems issues (that I am aware of) involved checking all the event logs from the counter in Legacy Horizon, and the 'Branch Access Layer' (BAL) in HNGx. While preparing this statement, I have come to understand that the Software Support Centre ("**SSC**") can comprehensively check all system logs if asked to.

41. When I joined the audit team, events from the event log were checked manually for suspicious looking events. Soon after I joined, this process started to become automated in stages so that all aspects of the process were eventually automated. In this regard, all events relevant to the time period requested in the relevant ARQ request were extracted automatically as part of the ARQ process. Each event was then checked against the signatures of those known to be benign from previous investigations. These benign events were then eliminated. If the relevant toggle button on the audit software was 'switched on', those events that were not eliminated were reported in a separate spreadsheet of events, which would be provided to the SSC to conduct additional checks. If the relevant toggle button was 'switched off', the audit software would not automatically generate a separate spreadsheet of events, but the SSC could still be asked to conduct additional checks in support of the ARQ log (I have raised concerns about the 'switching off' of this facility as routine, and the matter is being reconsidered). I do not recall what the SSC's additional checks involved, as the audit team was not involved in this part of the process. While preparing this statement, I have remembered, however, that the checking of the separate spreadsheets of events by the SSC stopped in about 2020.

WITN09870100

42. By way of example, a spreadsheet of events (FUJ00176752) was produced by this process in relation to a Peak (PC0263160) (FUJ00176751). In this case, it appears the SSC reviewed the events and concluded all were benign.

43. I recall that, in some instances, the SSC would provide feedback on the event checking process by raising Peaks and sending them through to the audit team. We would then look to enhance the filters applied when excluding certain patterns of events from the spreadsheets that the SSC would consider. This feedback would generally be to exclude event patterns that the SSC had previously identified as being benign.

44. I have been pointed to a number of Peaks in which I am referenced in this regard. Outlined below are those Peaks that I consider to be relevant to the event log checking process described above:

    a. Peak PC0260335 (FUJ00173091): This Peak was first raised on 7 July 2017, and relates to issues with the format of the event files processed by the audit team to produce the spreadsheets of events to be checked by the SSC. These issues were fixed on or around 19 July 2017.

    b. Peaks PC0280466 (FUJ00173193) and PC0280793 (FUJ00173184): The latter of these Peaks is a clone of the former, both of which relate to the event files used by the audit team to check for suspicious events being gathered late in respect of an ARQ data extraction. This incident occurred in or around October 2019. The audit team noticed the issue because it checks that events are present at the beginning and end of the query run for each ARQ. However, there is an unlikely possibility that

WITN09870100

event files the middle of the requested time period would be gathered late, as opposed to those at the beginning or end of the period, and that would escape notice. A more sophisticated checking system was developed because of this whereby it was checked that events were present for every day of the query. This Peak was closed on 10 March 2020.

c. Peak PC0261282 (FUJ00173153): This Peak was opened in August 2017, reporting an issue regarding event logs being filled with security events. A fix was tested and then rolled out in or around October 2017.

45. I am not fully familiar with the processes that produced the sources of information from which the ARQ logs for Legacy Horizon and HNGx were obtained and whether they were sufficient. However, in relation to XQilla, I do know that each message has a message number and the current XQilla-based audit software checks that there are no gaps or duplicates in the message number sequence.

Statement of Truth

I believe the content of this statement to be true.

Signed: GRO

Dated: 30 August 2023

WITN09870100

INDEX TO THE FIRST WITNESS STATEMENT OF GERALD BARNES

Exhibit No.	Document Description	Control No.	URN
1.	Bundle of documents relating to Mr Lee Castleton		LCAS0001383
2.	Audit Extraction Client User Manual (DEV/GEN/MAN/0015) v9.0 dated 21 May 2018	POINQ0164887F	FUJ00158710
3.	Audit Data Retrieval High Level Design (Audit Data Retrieval High Level Design) v2.0 dated 11 June 2010	POINQ0129972F	FUJ00123758
4.	Audit Data Collection & Storage High Level Design (DES/APP/HLD/0030) v1.0 dated 5 October 2009	POINQ0129973F	FUJ00123759
5.	Audit Data Retrieval High Level Design (SD/HLD/002) v1.0 dated 26 November 2004	POINQ0123679F	FUJ00117508
6.	High Level Design Specification for Track and Trace (T&T) Agents (DE/HLD/015) v4.0 dated 27 August 2008	POINQ0178024F	FUJ00171843
7.	High Level Design Specification for Agents for NBX, the NBE Replacement (NB/HLD/017) v2.0 dated 13 March 2006	POINQ0178026F	FUJ00171845
8.	Management of the Prosecution Support Service for Audit Record Queries (NB/PRO/003) v2.1 dated 20 November 2007	POINQ0128671F	FUJ00122457
9.	TPS Agents for BI3: High Level Design (AD/DES/041) v7.0 dated 9 August 2005	POINQ0097105F	FUJ00090934
10.	EPOSS Attribute Grammar Catalogue (EP/DES/002) v6.0 dated 12 December 1997	POINQ0085390F	FUJ00079219
11.	Riposte Attribute Grammar Catalogue – Messages (EP/LLD/019) v0.1 dated 19 July 1999	POINQ0178027F	FUJ00171846

WITN09870100

Exhibit No.	Document Description	Control No.	URN
12.	OBCS Counter Component Design Report (OB/DES/0004) v2.0 dated 15 December 1997	POINQ0178028F	FUJ00171847
13.	Peak PC0164429	POINQ0161560F	FUJ00155366
14.	Archive Server Configuration (DEV/INF/ION/0001) v15.0 dated 27 July 2015	POINQ0095347F	FUJ00089176
15.	Peak PC0152376	POINQ0160879F	FUJ00154684
16.	Email from Phil Budd with subject title “CounterDev - WP29300 'T86i1' now RFB - PC164429” dated 26 September 2008	POINQ0161460F	FUJ00155266
17.	Email chain involving G. Barnes, P. Thomas and A. Chambers with subject title “ARQ 499-509 – 475329 – LPD 19 Jan 09” dated 5 January 2009	POINQ0161031F	FUJ00154836
18.	Email chain with subject title “Audit Issue” dated 8 January 2009	POINQ0161596F	FUJ00155402
19.	Peak PC0205805	POINQ0178029F	FUJ00171848
20.	Peak PC0205353	POINQ0178073F	FUJ00171892
21.	Peak PC0207787	POINQ0178075F	FUJ00171894
22.	Peak PC0152828	POINQ0161405F	FUJ00155211
23.	Peak PC0153009	POINQ0161418F	FUJ00155224

WITN09870100

Exhibit No.	Document Description	Control No.	URN
24.	Peak PC0225071	POINQ0179238F	FUJ00173057
25.	Peak PC0225656	POINQ0178467F	FUJ00172286
26.	Peak PC0272681	POINQ0179364F	FUJ00173183
27.	Peak PC0276698	POINQ0179366F	FUJ00173185
28.	A spreadsheet of event logs in relation to Peak PC0263160	POINQ0230987F	FUJ00176752
29.	Peak PC0263160	POINQ0230986F	FUJ00176751
30.	Peak PC0260335	POINQ0179272F	FUJ00173091
31.	Peak PC0280466	POINQ0179374F	FUJ00173193
32.	Peak PC0280793	POINQ0179365F	FUJ00173184
33.	Peak PC0261282	POINQ0179334F	FUJ00173153