
From: Gauntlett, Paul [/o=ExchangeLabs/ou=Exchange Administrative Group (FYDIBOHF23SPDLT)/cn=Recipients/cn=20fad69c1be541619bbf58bbefdae198-Gauntlett,]
Sent: Thur 02/12/2021 2:54:04 PM (UTC)
To: Boardman, Phil [GRO]; Browell, Steven [GRO]
Barnes, Gerald [GRO]
Cc: Mistry, Manisha [GRO]
Subject: RE: Historical Issue with Audit Data

Hi Steve – I have no further comments.

From: Boardman, Phil [GRO]
Sent: Thursday, December 2, 2021 1:59 PM
To: Browell, Steven [GRO]; Gauntlett, Paul [GRO]; Barnes, Gerald [GRO]
Cc: Mistry, Manisha [GRO]
Subject: RE: Historical Issue with Audit Data

Looks good to me.

Regards, PhilB

From: Browell, Steven [GRO]
Sent: Thursday, December 2, 2021 1:34 PM
To: Gauntlett, Paul [GRO]; Barnes, Gerald [GRO]; Boardman, Phil [GRO]
Cc: Mistry, Manisha [GRO]
Subject: RE: Historical Issue with Audit Data

I attach my amended notes for your critique please before I pass to various parties for review and permission to release.

Steve Browell

Mob: [GRO]

From: Gauntlett, Paul [GRO]
Sent: Wednesday, December 1, 2021 5:01 PM
To: Barnes, Gerald [GRO]; Boardman, Phil [GRO]; Browell, Steven [GRO]
Cc: Mistry, Manisha [GRO]
Subject: RE: Historical Issue with Audit Data

Hi All,

Please see updated statement below.

- Accepted the changes that Phil added directly
- Orange updates based on comments below from Phil/Gerald.

Steve/Gerald/Phil – please confirm if happy with the below. I'm guessing there may be some further tweaks needed so will setup a short call for tomorrow am so we can finalise - I can cancel if not needed!

Problem Summary:

Post Office counter audit transaction files gathered prior to the HNGx software rewrite in 2010 are not consistent across both IRE11 & IRE19 servers. Holistically no data is missing but in rare instances transaction data exists only on one server or the other. The problem was caused by the occasional malfunction of the Harvester process of the previous Horizon system - Riposte.

Background:Audit Gathering Method in Horizon up to approx 2010

- Originally a system called Riposte was used, as part of Horizon, to gather audit transactions from all the Post Office counters.
- Riposte was a big distributed database.
- At that time Riposte was deployed into Bootle & Wigan data centres.
- Each evening all Post Office counter transactions were harvested and stored in files.
- By design transactions were either harvested to Wigan or Bootle or both data centres.
- Different files were produced in each data centre. The files were different because although the same transactions were harvested they were processed in a different order and for different FADs. Filenames were also different.
- The "CopyData" option (robocopy) was not set for transaction files which were gathered to both data centres.
- It is now known that due to the occasional malfunction of the Harvester process, for those transactions harvested to both datacentres, data may be missing from the files gathered in one or other data centre.

- There is no recorded instance of both harvesters failing at the same time so although data may be missing from one server it will be present in the other.
- NOTE: The data gathered using this method, in accordance with the contract, should have been deleted by now. It is only being retained currently, beyond Fujitsu Services' contracted obligations, at Post Office's request (below text from CWO0395b)
Post Office have requested under RTQSR0003106 (previously RTQSR0002349 and RTQSR0002456) that Fujitsu Services continue to preserve data (including Post Office Personal Data) generated on the HNG-X System as per the previous work orders (CT2616a & CW0251a) until 30 April 2022.

Audit Gathering Method for HNG-X from 2010 onwards

- Around 2010 the contract came up for renewal.
- Fujitsu's proposal was a rewrite called HNGx which eliminated Riposte (for which there was a big annual licence fee) and to migrate the datacentres from Wigan and Bootle to IRE11 and IRE19.
- As a part of this rewrite each Post Office counter transaction was only processed into a one file on a single server.
- This drove the decision to change the Audit gathering approach.
- From this point forward all transaction and non-transaction files were gathered by the IRE11 Audit Server only and robocopied to the IRE19 Audit Server.
- Therefore from 2010 all audit files (transaction and non-transaction) are consistent across both audit servers.

Impact on ARQ requests

- When ARQ requests for a FAD code are made all relevant files are retrieved from the target audit server .
- Then the files are processed by the Query Manager service on the audit server.
- Each Horizon or HNGx transaction has a unique number associated with it. The Query Manager checks that the transactions
 - have not been tampered with
 - that there are no duplicates or gaps in the sequence of transactions.
- Due to the Riposte harvester issue, for ARQ queries from 2010 or earlier, there may be gaps in the results from one or other of the servers.
- If this occurs then a Peak is raised and the ARQ request is rerun on the other server. This resolves any

issues with data gaps.

- There have been no reported instances of unsuccessful ARQ request once a query has been executed on the second server.

From: Barnes, Gerald <[REDACTED]>
Sent: Wednesday, December 1, 2021 1:29 PM
To: Boardman, Phil <[REDACTED]>; Gauntlett, Paul <[REDACTED]>; Browell, Steven <[REDACTED]>
Cc: Mistry, Manisha <[REDACTED]>
Subject: RE: Historical Issue with Audit Data

Hi Phil,

Answered inline.

Regards,
 Gerald Barnes

From: Boardman, Phil <[REDACTED]>
Sent: Wednesday, December 1, 2021 12:45 PM
To: Gauntlett, Paul <[REDACTED]>; Browell, Steven <[REDACTED]>
Cc: Barnes, Gerald <[REDACTED]>; Mistry, Manisha <[REDACTED]>
Subject: RE: Historical Issue with Audit Data

Hi Paul

I think these two bullet-points (in the pre-2010 section) are liable to cause confusion/consternation ...

- The Bootle audit server gathered Bootle transaction files and the Wigan audit server gathered Wigan transactions files.
- Although there was an option to robocopy files from one server it was switched off because there were already two copies of each transaction (one in each data centre).

... as I understand it (from Gerald's explanation (and I may have mis-understood)), by design transactions were either harvested to Wigan or Bootle or BOTH ... and the robocopy was ONLY turned off for those transactions configured to be copied to both datacentres ... I think it's important that we show that this was not designed to leave single copies of data anywhere. Please correct me if I'm wrong there Gerald.

My understanding (from reading historic copies of DEV/INF/ION/0001 and from practical experience) is that the main case of the "CopyData" option (the robocopy) not being set was these transaction files. The logic behind that I guess (I was not the designer) is that two copies of each transaction were made anyway – one on Bootle and one Wigan. If we need more detail on this I will need to get out from Dimensions historic copies of the configuration file to examine.

If I'm not wrong, then this statement "due to the occasional malfunction of the Harvester process in one or other of the data centres transactions may be missing from the files gathered in that data centre" needs to be clarified that that's only true for those transactions configured to be harvested to BOTH datacentres.

That is right. Pretty certain that was all Horizon (pre 2010) transaction data.

I've also proposed some changes to the text inline below (in this colour), trying to make it more clear that this was a change instigated by the Horizon to HNG-X change.

I think we also need to consider how we should include the details that the data in question is only being retained currently, beyond Fujitsu Services' contracted obligations, at Post Office's request (below text from CWO0395b) ...

Post Office have requested under RTQSR0003106 (previously RTQSR0002349 and RTQSR0002456) that Fujitsu Services continue to preserve data (including Post Office Personal Data) generated on the HNG-X System as per the previous work orders (CT2616a & CW0251a) until 30 April 2022.

... and (in accordance to our contract) should have been deleted, by now.

Yes that is right – normally the files would have been deleted long ago!

Regards, PhilB

From: Gauntlett, Paul <GRO>
Sent: Wednesday, December 1, 2021 11:51 AM
To: Browell, Steven <GRO>
Cc: Barnes, Gerald <GRO>; Mistry, Manisha <GRO>; Boardman, Phil <GRO>
Subject: Historical Issue with Audit Data

Hi Steve

I am Migration Lead for the Audit Migration to AWS.

Myself and Gerald Barnes (Audit SME) are currently working with POL to define requirements for the migration of historical Audit data

An historical issue has been identified and is detailed below.

This issue was discussed yesterday with John Nelis the POL PM & also Dean Bessell who I understand is your POL opposite.

POL requested that we write up what was communicated in the meeting so they can take it to their legal team ahead of making a decision regarding the scope of the data to be migrated.

On that basis I don't wish to send anything over without it being reviewed and agreed by relevant parties. Happy to have a call to discuss further.

Problem Summary:

Post Office counter audit transaction files gathered prior to the HNGx software rewrite in 2010 are not consistent across both IRE11 & IRE19 servers. Holistically no data is missing but in rare instances transaction data exists only on one server or the other. The problem was caused by the occasional malfunction of the Harvester process of the previous Horizon system - Riposte.

Background:

Audit Gathering Method in Horizon up to approx 2010

- Originally a system called Riposte was used, as part of Horizon, to gather audit transactions from all the Post Office counters.
- Riposte was a big distributed database.
- At that time Riposte was deployed into Bootle & Wigan data centres.
- Each evening all Post Office counter transactions were harvested and stored in files.
- The Bootle audit server gathered Bootle transaction files and the Wigan audit server gathered Wigan transactions files.
- Different files were produced in each data centre. The files were different because although the same transactions were harvested they were processed in a different order and for different FADs. Filenames were also different.
- Although there was an option to robocopy files from one server it was switched off because there were already two copies of each transaction (one in each data centre).

- It is now known that due to the occasional malfunction of the Harvester process in one or other of the data centres transactions may be missing from the files gathered in that data centre.
- There is no recorded instance of both harvesters failing at the same time so although data may be missing from one server it will be present in the other.
- TBC - In a DR situation data could have been lost was this captured as an operational risk and communicated to POL?

Audit Gathering Method for HNG-X from 2010 onwards

- Around 2010 the contract came up for renewal.
- Fujitsu's proposal was a rewrite called HNGx which eliminated Riposte (for which there was a big annual licence fee) and to migrate the datacentres from Wigan and Bootle to IRE11 and IRE19.
- As a part of this rewrite each Post Office counter transaction was only processed into a one file on a single server.
- This drove the decision to change the Audit gathering approach.
- From this point forward all transaction and non-transaction files were gathered by the IRE11 Audit Server only and robocopied to the IRE19 Audit Server.
- Therefore from 2010 all audit files (transaction and non-transaction) are consistent across both audit servers.

Impact on ARQ requests

- When ARQ requests for a FAD code are made all relevant files are retrieved from the target audit server .
- Then the files are processed by the Query Manager service on the audit server.
- Each Horizon or HNGx transaction has a unique number associated with it. The Query Manager checks that the transactions
 - have not been tampered with
 - that there are no duplicates or gaps in the sequence of transactions.
- Due to the Riposte harvester issue, for ARQ queries from 2010 or earlier, there may be gaps in the results from one or other of the servers.
- If this occurs then a Peak is raised and the ARQ request is rerun on the other server. This resolves any issues with data gaps.
- There have been no reported instances of unsuccessful ARQ request once a query has been executed on the second server.

Regards,

Paul Gauntlett
Customer Solution Architect
Cloud Transformation & Development - AMCS

Fujitsu
Central Park, Northampton Road, Manchester, M40 5BP
United Kingdom

     
paul.gauntlett@gro
www.fujitsu.com/uk