
From: Gauntlett, Paul[o=ExchangeLabs/ou=Exchange Administrative Group (FYDIBOHF23SPDLT)/cn=Recipients/cn=20fad69c1be541619bbf58bbefdae198-Gauntlett,]
Sent: Wed 01/12/2021 11:22:36 AM (UTC)
To: Barnes, Gerald [REDACTED] **GRO**
Subject: RE: Summary of Riposte and HNGx

Hi Gerald,

As discussed please review

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.

Audit Gathering Method up to approx 2010

- Originally a system called Riposte was used 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 to other it was switched off because there was already two copies of each transaction (one in each data centre).
- Due to 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 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).
- 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
 - whether or not there are any duplicates or gaps in the sequence of transactions.
- Due to the pre rewrite harvester issue, for queries on data from 2010 or earlier, there may be gaps in the ARQ query results.
- If this occurs then a Peak is raised and the ARQ request is rerun on the other server. This will 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 GRO
Sent: Tuesday, November 30, 2021 5:10 PM
To: Gauntlett, Paul GRO
Subject: Summary of Riposte and HNGx

Hi Paul,

1. Originally a system called Riposte was used to gather audit transactions from all the Post Office counters. It was a big distributed database.
2. There were two data centres. Each evening all transactions were harvested and put in files. Different files were produced in each data centre but they each contained the same transactions all be it in a different order and for different FADs.
3. One audit server gathered one file and the other the other. Although there was an option to robocopy files from one to other it was switched off for these files because there was already two copies of each transaction.
4. Around 2010 the contract for looking after Post Office transactions came up for renewal. Fujitsu's proposal was a rewrite called HNGx which eliminated Riposte (for which there was a big annual licence fee). As a part of this rewrite all transactions were stored in one file only.
5. The auditing of these files was done by one audit server only and what it stored each evening was robocopied and stored on the other audit server.
6. Once all Post Office counters were converted to HNGx as a tidying up exercise the audit configuration files were changed such that only IRE11 actively gathered and all files that were stored were robocopied to IRE19 each evening.
7. When ARQ requests for a FAD code are made all relevant files are got back from the audit server being used (they may have been robocopied, they may be gathered directly). Then they are processed by the Query Manager service on the audit server. Now whether Horizon or HNGx each transaction has a unique number associated with it. The Query Manager checks that the transactions have not been tampered with and also checks whether or not there are any duplicates or gaps in the sequence of transactions.
8. It is very rarely the case that there are gaps in the Horizon transactions on one audit server but not the other.
9. These gaps are caused by the rare malfunction of the Harvester process on one data centre.
10. I went through PEAKs yesterday using a free text search and found one such instance reported. I looked in detail at that. Now Horizon nodes have two sorts – counter transactions which are nodes 1-32(or maybe 31) and some sort of administrative correspondence server nodes which have a bigger number. The gaps were in the correspondence server nodes only. There is good reason for that. In Riposte all transactions have a build in retain time and it is more for real transactions than correspondence server transactions. So if the harvester is down for a period it is more likely that the correspondence server transactions disappear than the real ones. So maybe the harvester was down for some reason or other and by the time it got going again the correspondence server transactions had disappeared.

Regards,
Gerald Barnes

Gerald Barnes

Senior Software Engineer

Post Office Account

Fujitsu

Lovelace Road, Bracknell, Berks, RG12 8SN

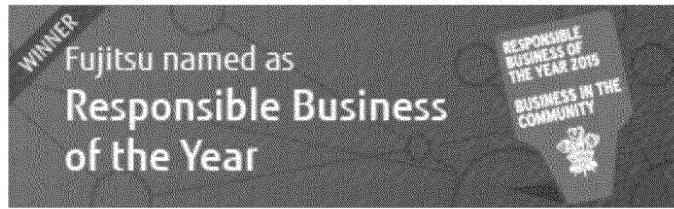
Tel: **GRO**

Mob: **GRO**

Email: Gerald.Barnes@fujitsu.com

Linked in: **GRO**

Web: <http://uk.fujitsu.com>



Fujitsu is proud to partner with [Action for Children](#)

[I-CIO](#): Global Intelligence for the CIO. Fujitsu's online resource for ICT leaders

Sponsors of the 2015 Rugby World Cup



Please consider the environment - do you really need to print this email?