

---

**From:** Ashford, Edward[/O=EXCHANGELABS/OU=EXCHANGE ADMINISTRATIVE GROUP (FYDIBOHF23SPDLT)/CN=RECIPIENTS/CN=D359CE8E14A94C9BA3EB1066E6C67C41-ASHFORD, ED]  
**Sent:** Fri 14/01/2022 10:58:50 AM (UTC)  
**To:** Gauntlett, Paul [REDACTED] GRO  
**Subject:** FW: Mirroring the audit archives in IRE11 and IRE19

for info

---

**From:** Browell, Steven [REDACTED] GRO [REDACTED]  
**Sent:** Friday, January 14, 2022 10:45 AM  
**To:** Barnes, Gerald [REDACTED] GRO [REDACTED]; Ashford, Edward [REDACTED] GRO [REDACTED]  
**Cc:** Hemingway, Andrew R [REDACTED] GRO [REDACTED]; Gibson, Andrew R [REDACTED] GRO [REDACTED]  
**Subject:** RE: Mirroring the audit archives in IRE11 and IRE19

Thanks Ed & Gerald.

I do not fully understand all of what was shared as I lack the experience and skill.

The concern is that we do not have exactly the same Audit Archive data in both IRE11 and IRE19. For the gap we are aware of, I can see that Gerald has wisely addressed that and taken a copy. However, there MIGHT be other gaps and if we lost a DC and were running an ARQ on the remaining DC and it saw a gap then we would not be able to check the other store to hopefully find the absent content and hence have a complete result set.

I believe Fujitsu has a responsibility to ensure that both DCs have identical data – for the Audit Archive specifically for now – and we should take steps to ensure that is the case.

Given that HNG-X writes to the IRE11 Audit Archive and then that is replicated to IRE19, we have mirrored data for HNG-X.

For pre HNG-X we do not. Hence, I think we need to copy the pre HNG-X Audit Archive that is in IRE19 that came from Wigan across to IRE11 – and vice versa.

I realise that the AWS migration will likely copy both pre HNG-X Audit Archive stores (one from IRE11 and one from IRE19) plus the HNG-X Audit Archive from either DC but that does not address our current issue.

I see the following options:

1. Do nothing – knowing that the known gap has been addressed by Gerald's actions
2. Treat AWS S3 as the solution – which may indeed be the best option depending on lead times to deploy option 3
3. Add storage to IRE11 to receive the IRE19 data, and storage to IRE19 to receive the IRE11 data

Hence, we need to quantify option 3 and invite Dan/Wendy to support the activity.

- **How much storage would we need to purchase for each DC and what type should it be (I would like the shopping list so we can get a price and lead times)?**
- **How long MIGHT it take to install and make the new storage available?**
- **How long MIGHT it take to copy the data IRE11>IRE19 and IRE19>IRE11?**

---

Steve Browell

Mob: GRO

---

**From:** Barnes, Gerald <GRO>  
**Sent:** Wednesday, January 12, 2022 1:27 PM  
**To:** Ashford, Edward <GRO>; Browell, Steven <GRO>  
**Cc:** Hemingway, Andrew R <GRO>; Gibson, Andrew R <GRO>  
**Subject:** RE: Mirroring the audit archives in IRE11 and IRE19

Hi,

Another reason for slight discrepancies would be orphans.

Sometimes the API for storing a file fails. An automatic retry is done. I have looked at some of these failures and in some cases it is a network issue such that although failure was reported the file was actually saved. So then you have two copies of the file on the Eternus – one pointed at and one an orphan. Whether or not this happens is quite unpredictable.

Regards,  
Gerald Barnes

---

**From:** Barnes, Gerald  
**Sent:** Wednesday, January 12, 2022 1:00 PM  
**To:** Ashford, Edward <GRO>; Browell, Steven <GRO>  
**Cc:** Hemingway, Andrew R <GRO>; Gibson, Andrew R <GRO>  
**Subject:** RE: Mirroring the audit archives in IRE11 and IRE19

Hi Ed,

There is only one known gap in the transaction files that is known of and that I have covered.

In terms of the other files I have not done an in depth study but I do note that the audit configuration files can be configured to either take two copies or one copy on a sub-point basis so that may account for it. When I looked back at early versions of DEV/INF/ION/0001 it said that the main use of not having “Xfer Data” set on was for the transaction files (where normally you get two copies of each transaction in different files by other means) but it did not say it was the only use. Maybe some other log files were considered not important enough to take two copies. After all when these were backed up to tape designers may well have put more importance to whether there should be one or two copies. Now “Xfer Data” is on for all sub points and I do remember a long time ago having to click it on for a few sub points (it was at the time of the replacement of Horizon by HNGx) after we had decided to be consistent in this.

Regards,  
Gerald Barnes

---

**From:** Ashford, Edward <GRO>  
**Sent:** Wednesday, January 12, 2022 12:35 PM  
**To:** Barnes, Gerald <GRO>; Browell, Steven <GRO>  
**Cc:** Hemingway, Andrew R <GRO>; Gibson, Andrew R <GRO>  
**Subject:** RE: Mirroring the audit archives in IRE11 and IRE19

A summary for convenience.

I notice that even for the 3, 4, 5, 6 shares which I would have expected to be HNG-x only and to match for number of files and space used (look at the Premigrated columns) there is a slight discrepancy between IRE11 and IRE19.

For the shares that would contain Horizon data the discrepancies are in hundreds or thousands of files, so I think the safest thing to assume is that we just need a copy of the whole lot.

There is probably space to "cross-mirror" the W and B Correspondence server files, but this is a different exercise than the one we did previously.

I think Audit Tactical is probably doomed to copy both sites no matter what, just in case.

Regards

Ed

---

**From:** Barnes, Gerald <[REDACTED]**GRO**>  
**Sent:** Wednesday, January 12, 2022 12:04 PM  
**To:** Ashford, Edward <[REDACTED]**GRO**>; Browell, Steven <[REDACTED]**GRO**>  
**Cc:** Hemingway, Andrew R <[REDACTED]**GRO**>; Gibson, Andrew R <[REDACTED]**GRO**>  
**Subject:** RE: Mirroring the audit archives in IRE11 and IRE19

Hi,

Short term there is only one gap that I am aware of so I have run an ARQ to save all the missing files on IRE19. So therefore right now we have two copies – one in the Eternus and one on the F: drive of IRE19. These files could be just copied to the special MisArcs folder of IRE11 where they will be regathered and stored on the IRE11 Eternus and then at least we have two copies of all files.

Longer term we intend to make copies of the data in S3 storage anyway as Ed mentioned. The software to do this is already written. It needs to go through test first though. We might be able to bring this work forward a bit I suppose.

Regards,  
Gerald Barnes

---

**From:** Ashford, Edward <[REDACTED]**GRO**>  
**Sent:** Wednesday, January 12, 2022 11:52 AM  
**To:** Browell, Steven <[REDACTED]**GRO**>; Barnes, Gerald <[REDACTED]**GRO**>  
**Cc:** Hemingway, Andrew R <[REDACTED]**GRO**>; Gibson, Andrew R <[REDACTED]**GRO**>  
**Subject:** RE: Mirroring the audit archives in IRE11 and IRE19

Hi Steve,

The basic process is similar to that used for EMC to ETERNUS migration, which is also being proposed for ETERNUS to AWS S3 migration.

<https://pol-jira.atlassian.net/wiki/spaces/CO/pages/762380521/ARC+PCI+Migration+Design#Migration-Methodology>

We would just about have that kind of storage available in Belfast if we were to resurrect the POLMI disks into the DX8700, but all of the arrays are EOSL so this is not an ideal solution. This could present either one enormous drive or 10 small (2TB) ones (to mimic the pattern used on ETERNUS shares) directly to the Audit server at each site.

It strikes me that we want to make the copy of IRE19 reside at IRE11 and vice versa. This is possible by presenting

IRE11 storage to ARC202 and IRE19 storage to ARC201.

Purchasing storage for this would need to purchase a small array, preferably capable of NAS presentation, and install it first. This probably has a lead time of at least 8 - 12 weeks, and prior to that we would need Legal approval which can potentially take a lot longer.

Once the storage is presented its out of my hands and into Gerald's.

The pragmatic approach would seem to be to put a bit of urgency around initiating the copy to AWS S3.

Regards

Ed

---

**From:** Browell, Steven <[REDACTED]**GRO**[REDACTED]>  
**Sent:** Wednesday, January 12, 2022 11:15 AM  
**To:** Ashford, Edward <[REDACTED]**GRO**[REDACTED]>; Barnes, Gerald <[REDACTED]**GRO**[REDACTED]>  
**Subject:** Mirroring the audit archives in IRE11 and IRE19

Ed,

As part of the work you both helped me with recently I came away believing that the Audit Archive in IRE11 (that came from Bootle) is not exactly the same as the Audit Archive in IRE19 (that came from Wigan) for the pre-HNG-X data. Therefore, in the unlikely event of a catastrophic loss of any single data centre, we would lose one of the Audit Archives and hence could have gaps in data when performing ARQs.

I believe we need to create a copy of the pre-HNG-X Audit Archive datastore for both Bootle and Wigan in both IRE11 and IRE19.

What do you think are the bullet point list of steps we should take to address this, and what approximate amount of additional storage would be needed in each Data Centre please?

From this we will look to create a CP and run a proper project to address it.

---

**Steve Browell**  
Post Office Account  
Management Consultant & CISO

**Fujitsu Enterprise & Cyber Security**  
Fujitsu Services, Trafalgar House, Temple Court, Risley, Warrington, Cheshire, WA3 6GD, United Kingdom

Mob: [REDACTED]**GRO**[REDACTED]  
E-mail: [Steven.Browell@fujitsu.com](mailto:Steven.Browell@fujitsu.com) [REDACTED]**GRO**[REDACTED]



Planned leave: None