

RESTRICTED - COMMERCIAL

PWY/JFO/002**POCL Infrastructure Demo - Meeting Report****Supplier:** Pathway**Date:** 8 November 1995**Attendees:****BA/POCL**Jeremy Folkes  
Bob Booth**Supplier**Martin Johnston  
Martin McAdam (An Post)  
Dave Cooke  
Liam Church (An Post)  
Fionnuala Higgins (An Post)  
Joe Cosgrave (An Post)  
Mike Murphy (Escher - from 16.05)**Purpose:**

Second "POCL Infrastructure" demonstrator meeting - specialising on technical aspects of Riposte.

**Items of Note:**1. Informally reviewed outstanding requirements queries:

- Mobile offices - steer given not to let consideration of these distort their solution. Mobiles would be likely to be on a per situation basis. Pathway stated GSM would be their preferred solution.
- Scales - steer is to assume outwith this procurement, just connectivity to be demonstrated. May be potential for synergy when installing kit, but outwith this tender. Discussed again when type approval may be required - timetable being sought.
- Support for smart cards / keys - steer given that keys are unlikely to become nationally justified and more likely to be regional along lines of the Work load document provided. Pathway requested number of offices, transactions per office type figures. Confirmed that current view is that SP are unlikely to get access to technology partners until after selection.
- 2D bar-code - steer given that there is no current requirement to either read, or produce 2D bar-codes. However, a cost implication would assist the requirements team in decision process.
- OCR of data on bills - steer that this is with sponsors, but not currently seen as likely major area; clients use mag cards and move is towards bar codes; noted An Post demonstrably use OCR and again a costed option would assist consideration. Steer given that issue was commercial as much bill business done for Girobank rather than directly for end client.
- Thermal printer - question still with programme, may depend on ability to do parts as impact. Pathway need to supply data on fade, storage requirements etc. This ties in with apportioning of risk, as SP may want access to evidence. It was recognised that the cost of printers was significant, and the cost to maintain and for consumables also factored.
- Compatibility with TCDs - no requirement at present. ECCO kit should be available at Terminal House this week, and scales may be attached. TCDs are not currently attached to ECCO. Bob to determine number TCDs deployed. Bob to brief JM prior to

RESTRICTED - COMMERCIAL

PWY/JFO/002

next week on requirement implications such as cash holding, replenishment and ownership within the office.

- Sample APPU was requested and should be available by the end of the month. Will Russell is organising, has in stock and about to send out.
- Electronic purse. Steer given that we would just want evidence that it can be added (e.g. a serial port).
- Consumables - steer given that as they would be for use in the SP's kit, these were currently considered as their responsibility - though distribution via POCL channels could be investigated with commercial strand. SP needs to propose what as POCL's responsibilities.
- General nudge - up to service provider to propose solutions and to clearly state assumptions made about BA/POCL responsibilities to make the proposals viable. Prefer before ITT so no surprises and avenues can be smoothed.
- Colour printing - steer not required
- C(TMS)1 schedule reference for message switching between PO's. No current requirement apart from any need predicated by the SP's solution. Would not expect it to be excluded but any use would be the subject of change control. Example use might be use for transmitting remittance information between offices.
- Requirement for other on-line products. Outline of possible stock ordering given, e.g. messages upwards of 500 bytes, up to 20,000 messages per day across the system in semi-real time (a few minutes to deliver). Workload Brief *may* hint at others, but depends on re-engineering and look to the service provider to innovate and provide capacity.

## 2. Discussion on user administration:

- Pathway saying that terminals *may* be initialised by a magnetic card, and then used to allow a till to follow the clerk if they move (without forcing a close-down - just swipe and the terminal "follows" the swipe) - this being a facility introduced in Rip23 and under consideration for An Post. Storage of these cards was unclear, as was ownership / production / distribution.
- Pathway saying terminals *may* have their personality stored on a Smart Card that the engineer uses on swap out etc.. Storage of these cards was unclear, engineer or office.

## 3. Riposte - Mike Murphy (Escher) joined late on and altered understanding of some areas. The "distilled" version is below:

- Riposte is fundamental to the counter, back through TMS and for the provision of PAS
- Single file JOURNAL.TXT contains variable length records, it is only appended to (see Archive), and there is one file per machine. Other machines journals have their entries replicated across the system so all entries should exist at all machines - though the order will differ depending on when the entries arrived, and there are potentially "additional" (but redundant) entries after a failure scenario - all the transaction data is retained. If not all nodes present - unsure how this is determined at broadcast time - then the transaction is marked as such to assist in collision reconciliation.
- The journal file is moving away from being text based in Riposte 2 to binary (though still with human readable content in most cases) under Riposte 32.

RESTRICTED - COMMERCIAL

PWY/JFO/002

- All records have 64-bit CRC, some have digital signatures RSA / DES unclear which and where and how chosen - guess in the grammar describing the entry.
- Replication is achieved (in Riposte 32 by a single IP multi-cast, on Riposte 2 by individual NetBEUI calls) by telling all other nodes that appear on a local list and a global list held locally on the distributing machine.
- All machines can and will replicate to other machines. The position 1 machine is the only one that can talk to the correspondence server - which is effectively another office node (from the office viewpoint) that happens to be remote - and this is controlled by its GLOBAL file.
- An attribute in the journal server shows whether a transaction should force a connection to be made over ISDN. Whilst the connection is open (for duration of whole units) other data is replicated, but the connection will be closed (even if the other data is flowing) when the unit expires and there is no data of high enough priority to open the link. How is this configured, given changes in telecomms tariffs?
- When a machine comes up, it broadcasts to all its known nodes, both its own latest state (i.e. the CRC of its last journal entry) and the latest state of all its nodes *as known to it*. If there is a discrepancy, then nodes report back, for both themselves and other nodes whose data has been replicated back to them, to bring the new machine up to date, or conversely to get data they have missed off of the new node. If there are collisions over the sequence numbers (e.g. Pos 3 seq 32 appears twice with different CRC's) the owning journal server is prevented from serving until the situation is sorted out.
- When a new entry is made, it is broadcast. If another node receives it, the entry is rejected if the sequence number is not the next one expected (still true ?) and the receiver then re-synchs.
- Concept of entries having mortality - where they can be deleted (see Archive later) and this will vary dependant on the attribute (type) of entry. Some entries are Persistent Journal Objects - typically describing bill payments or reports and are self contained. A PJO is added via an API (also termed Retail Broker) that knows about CRC's, security etc. A PJO can only be added if it is new ('plu') or points back to the PJO it replaces.
- Correspondence servers replicate in the same manner - according to a local list - amongst themselves. Correspondence servers will be beefed up (4 processor ...). It is envisaged that there will be two physical sites (using n x 2M between sites), each split into two campuses (on site 100Mbps, with separate machine room, power etc.), and replication will be to partner on other campus and then onto other site - possibly to both campuses. One node may be a write only node (i.e. not replicating out) for audit purposes.  
Management of new offices to be taken up in system management next week.
- Rough sizing was for each correspondence server to handle 43 tps, based on 20M over 8 hours (695 tps over the system).
- Riposte has indices on the journal. Either permanent, updated each time an entry is made (though this can be disabled!), and temporary, where the index is built on demand and then discarded.
- Indexing was rather confused. There appears to be different type of index depending of the potential size, and they plan to uses a hierarchical one on the correspondence server where the volumes are largest, and a quicker to write, flatter index at the counter where volumes are smaller. The permanent index will be updated when an entry is written, adding to a B-tree, hash table. The hash table is CRC'd, and in the hierarchical case this is done in lumps, each lump indexed and

RESTRICTED - COMMERCIAL

PWY/JFO/002

CRC'd so corruption may only need partial rebuild. Is all this effort is it because they keep getting corrupt indices so need to rebuild often?

- Inference was that they can create parts of index table on different machines and then merge them, proposed for taken on from CAPS when "10 million entries may hit the system". How does this tie into journals having same stuff but in different orders ?
- New indexing is done when corruption is detected at the counter.
- Archive will be an on the fly operation. Riposte was original Windows for Workgroups - not true multi-tasking - so archives were done with the journal server off-line. Now intend to apply weighting to entries (base on mortality and type), combine with disk space and then determine what, if anything to delete. When delete updates the hash index, and can return blocks from the middle of files back to NT, and uses NT transaction system to ensure integrity "standard NT stuff".  
Contacted PO/iT/TAS to see if they can clarify/educate on this.  
Envisage the workstations archive would discard, the centre would archive to CD-ROM. Unclear on performance impact whilst archiving.
- For EFTPOS (true on-line) an agent would be used, from Riposte to another piece of software within / at the boundary of its trusted network within the correspondence layer. Whilst doing this the counter can continue and get the response later. Software designed for 16 events extant but throttled back to 3 as a usability issue.
- Riposte API to be supplied by Pathway. Documentation !! Also requested last week.
- Binary data can be included in the journal (but UUENCODE or Base 64 to preserve binary sentinels) so Schlumberger key images etc. can be stored. But data will be expanded in the process!
- Software is intended to be the same between journal and correspondence layers, but with different features enabled.
- Riposte written in C++; front ends in Visual Basic 4.
- Audit log longevity was questioned. 18 months is best steer, but probably not in the office.
- Escher scorned Pathway concern at one position office and threw a PCMCIA hard disk to Pathway member saying "*this will be in the back of the one position office machines and will be swapped on machine failure*". Apparently news to Pathway. Actually a good idea (if costs in). The MiniScribe disk being thrown around was 260M giving "*3.5M journal entries*". Riposte 2 would update by having 2 copies of journal server, one acting as normal and the other replicating to the spare disk. Riposte 32 can either do that, or write to two disks at once. No thought to what other than the journal may be stored. Take up along with software distribution at System Management meeting next week.
- The network was unclear, with mentions of ISDN CUG, and PRI's (though not all were sure what a PRI was) feeding into the correspondence servers, or into FEPs or .... Still modelling (about to ?) ISDN links required, we will need surety that the servers can handle the traffic levels and the replication, and archiving, and recovering offices, and CAPS feed, and system management ...
- Confusion in Pathway between Escher and rest with regards to terminal having "strong identity" - did Pathway/An Post understand ? - and what dangle that was thought to have been discounted re-surfaced, though may now - decided in meeting - use a smart card to allocate blocks of sequence numbers for each terminal to underpin Riposte's design basis of CRC and sequence number.

RESTRICTED - COMMERCIAL

PWY/JFO/002

- Shock from Mike Murphy at thought of a 9" screen on the counter - view was "*big screen small keyboard*" at odds with Pathway direction. Obvious cost and size implications here, need to wait for EPOS to pan out. Who is defining the hardware?
- When queried about sizing model Pathway had agreed to produce paper, when this was mentioned later in day with Escher present, the Mike Murphy response to Pathway co-members was "***how can you do that, Pathway haven't got a clue how it works***". Lack of demonstrable thought processes.
- Watch this space as the Pathway paper **should contain** information on network sizing (in the office, office to correspondence layer, within the correspondence layer), disk sizing, move to NT, move to IP (so IP multi-cast can be used to reduce network traffic instead of individually addressing NetBEUI), and Topology Supervisor as a starter.
- Lack of cohesion between the people at the meeting, must be doubt over ability to manage project if this interface to their customer is so weak.
- General problem that there is no documentation about the system, and late arrival of Mike Murphy and his contradiction / clarification re earlier explanations call into severe doubt the knowledge of the consortia about what they are proposing and then how they may develop it / support it in any time scale.
- Large difference between Riposte 2 and Riposte 32 - OS, indexing, ISDN, move from DDE, data stores to name but a few that came out. We may get incremental change notes.
- Large difference between An Post system and what is being proposed.
- Mention of Bill Gates presenting to Andrew Stott on 2 December. Suggestion that JF might benefit from visit to Cambridge Mass. to meet Escher designers etc.

#### 4. Reviewed Risk Register:

- PWY7 - previously transferred to John Meagher
- PWY8 - Card reader mechanism to be taken up with peripherals - 15/11.
- PWY11 - keyboard size relating to transaction speed and suitability of touch screen for EPOS to be covered with peripherals - 15/11. Note Pathway subsequently have queried this with the SLM - their risk team seem out of touch with this meeting.
- PWY19 - to be taken up with peripherals - 15/11.
- PWY29 - CLEAR
- PWY48 - software distribution, response paper in progress.

#### Papers Received:

None.

RESTRICTED - COMMERCIAL

PWY/JFO/002

**Requirements Queries:**

Query	Steer/Recommendation	Strand	Date Req'd
Would we be interested in using the installation logistics for the rollout of further scales..... do we have any guidance for the number of scales.....	Assume baseline is we wouldn't, but worth considering (and with other kit). SP should raise with Implementation Stread.	Infra / Implem	24/11
Smart key figures in Workload Brief do not give indication of spread of offices or likely duty cycles - design of receptacles will depend. Further info needed.	Can Workload Brief be extended in this area?	Infra	24/11
BA receipt - processing, storage etc. - is this in the requirements? Will SP be allowed access to receipts for investigative purposes, if they are expected to take the risk.	Suggested SP talk to BPS	BPS	24/11
Requirements for Teller Cash Dispenser connectivity? How many are actually in use, likely to be in use? What would we actually want connectivity to mean?	Connectivity at infrastructure level probably shown by serial port  RB to brief JM on applic issues	Infra / Applies	24/11
Possibility of using the current PO distribution channels for consumables etc.	PO would be open to suggestions.	Infra	24/11
Archive period length, along with who would enquire, using whose facilities and how quickly is response given?	Suggested 18 months as first stab at retention.	Infra	24/11

Jeremy Folkes &amp; Bob Booth 10/11/95