



To: Dave Hollingsworth
cc: Simon Bond
Tony Houghton
Date: 06 August 1999
From: Bob Booth
Subject: AI 372 - System management of LT1 to LT2

As you are aware from the Acceptance Incident Workshop on Wednesday and subsequent brief meeting on Thursday, I have been nominated to own AI 372.

At our meeting I agreed to set out observations on the supplied report ["Report on Upgrade from LT1 to LT2 on 10th and 11th July" version 2 dated 16/7/99] for ICL Pathway to respond to.

The observations are laid out below, using the section names in italics and a paragraph count where required. The points are in places purposefully laborious to ensure that the responses should be able to close the majority down in one hit. Unfortunately the number of observations reflects the concern on the effectiveness of the system management.

We recognised that the report had depth and that good System Management of the estate was mutually beneficial to both ICL Pathway and POCL with shared pain and benefits.

We also touched upon the difficulty of proving that where deficiencies have been noted, that they have been adequately addressed (or are in the process of being addressed). We agreed that the ability to compile and execute a comparable upgrade, proving the improvements, is not feasible in the time available, and therefore some statement of monitoring would be required.

1. Could you please propose a monitoring schema that ICL Pathway could achieve for POCL to consider.

POCL could view the plan for the next major upgrade. They had sight of the plan for this upgrade also so this may not resolve the detail issues.

Observations

2. *1 Introduction* - this makes note of further activities scheduled for 19th July but does not report upon them. An update is requested on these activities.

Nothing contained herein shall be deemed or construed as affecting existing contractual obligations or creating new contractual obligations between ICL Pathway and POCL.

These activities have been completed in the expected timescales. Verbal confirmation was made to Rod Stocker.

3. *2 Management Summary - §1: quote: “a considerable success”. This statement is at odds with the body of the report which indicates that there were considerable difficulties but they were overcome. At times there were judgement calls (e.g. fault not known but carry on) which went well, allowing the late completion of 288 outlets and failure to complete 11 outlets (3.7% or 700 outlets at full rollout). Bearing in mind the body of the report, do ICL Pathway still consider the upgrade a considerable success ?*

Yes – the aim of the weekend was to upgrade the data centre and the counters. The datacentre was completed and the majority of the counters were also completed. The difficulties found have allowed us to learn and improve things for the future.

4. *3 Weekend activities - detailed - §1: a copy of the log of achieved activities and revised activities against the plan is requested to underpin and clarify the observations made in the report.*

Rod Stocker had a copy of the plan that was used for the weekend activities. To produce a revised plan would need considerable effort. Copies of the checkpoint sheets could be made available as they were to John Bruce while he was on site for the Sunday and Monday morning.

5. *3.1 Saturday - §4: it is unclear why the decision to proceed without remaining EODs was taken and whether this was a planned contingency or not. Clarification is requested.*

I saw no reason to suspend committal to the remaining estate due to missing end of day markers. The end of day markers could have been missing for a number of different reasons. However as specified in Section 5.2 in future upgrades the offices which fail to produce an end of day marker will be suspended within Tivoli and their committal will be scheduled during the normal overnight committal window.

6. *3.1 Saturday - §4: it is recognised that whilst some events are not foreseeable, if there is a clear decision point (having all the markers in this case), then a prepared contingency option would be expected at such a point. Comment on this view is requested.*

See response to point 5.

7. *3.1 Saturday - §4: it is unclear from the report if the scheduled activities were confined to those outlets from which markers had been received. Clarification is requested.*

Committal continued on all counters which were expected to be closed on the Saturday afternoon.

8. *3.1 Saturday - §10: it is noted that SMC continued with the committal of other LT2 fixes whilst the issue of the missing ‘D’ data was investigated.*

Confirmation on the level of testing / confidence that this route could have been regressed if required is requested.

I saw no reason to suspend the committal of the remaining two counter fixes while the missing D data was investigated. At the time no accurate timing data was available to estimate the time required to commit all three counter fixes across the live estate so I was loath to suspend the process. Similarly regression timings were not available.

Section 5.1 point 2 also deals with this issue. Tivoli timings have now been taken and a sizing model is under production making it possible in the future to include counter regression and decision points in an upgrade plan.

9. 3.1 Saturday - §11: clarification is requested if the statement 'replicated across the counter estate' refers to replication within an outlet or to an outlet.

'replicated across the counter estate' refers to replication to all outlets.

10. 3.1 Saturday - §16: an update on the cause of the MIS extract failure is requested.

The underlying fault has not been identified.

11. 3.1 Saturday - §16: confirmation that the estate could have rolled back should the failure have been determined to be an LT2 fault is requested.

The MIS fault was experience in the Saturday overnight while the system was still running LT1. However data centre failover was tested the weekend prior to the LT2 upgrade weekend so that failover could have been achieved if necessary.

12. 3.1 Saturday - §21: given the precise timing of the desktop reload (currently) clarification on why a 60 minute window was left for this activity.
Given the move to allow the user to temporarily defer this, unless some form of checking for quiescence is used this 'pause' may need to extend from 03:00 to 05:00. A view on this impact is sought.

I decided to leave a sixty minute window to ensure the counters had time to reload before we attempted counter committals. This is currently normal operating process. If a change is implemented such that the user can defer this reload the normal operating process will have to be reviewed. However committal attempts while the counter is down will just return an error through the distribution mechanism and the application can be re-targeted later.

13. 3.2 Sunday - §3: clarification is sought on why LT1 elements were still running.

As stated in section 5.1 point 4 although the need for a change in the system management script had been recognised it was not included in the correct point in the plan. Hence the system had some LT1 elements running.

14. 3.2 Sunday - §5: given the previous 'need' for live proving, was there any proving performed before the 14th ?

Yes – As per section 3.2 paragraph 5. This testing was specified by John Bruce and the tests were completed as expected.

15. 3.2 Sunday - §13 & 14: given that testing was undertaken at 22:00 on Sunday, and decisions taken around 03:00 Monday this would appear to have been a potentially high risk option. What surety was there that the counters could be regressed to be operational for Monday morning if required ?

See response to point 8. The majority of the counters had the two main fixes applied by 0700 on Sunday morning. I believed we had sufficient time to regress the fixes should it be required. It was a potentially high risk option.

16. 3.2 Sunday - §15: given the rapid improvement in committal rate after 04:00 Monday, why was this method not used earlier ?

This method was not used earlier as it leaves the counter inventory in an inconsistent state which is obviously not an ideal situation. The inventory had to be manually updated later to ensure it reflected the live estate. The decision was left until 4am as committals were halted at 3pm to permit the counter reloads and this seemed a suitable point to review. The progress up to 3am had been continually reviewed and discussions were held with development to suggest the change in process. At 4am the decision was made to change the process, having first clarified that the only effected area would be the inventory.

17. 4.1 Monday - §1: given the target was 07:00 what number of outlets were complete at this time ? (will this be on the log against plan requested above ?)

This figure could be made available if required but will take a while to produce.

18. 4.1 Monday - §2: given that 29 outlets (10%) were still extant at 09:00, at full volume how would this situation be addressed ? (1,900 outlets extant ?)

I can't help here but I think Glen probably can.

19. 5 Issues: the impact in parenthesis alongside each issue appear at odds with the issue. If indeed there was no impact why was there an issue. In some cases there were internal matters that needed immediate resolution and in others the end user was affected. The impacts should be revisited e.g. 5.10.

The impacts were viewed in the overall weekend activity. Although a number of areas had to be addressed during the weekend and as part of a follow up activity most of the issues had no impact on achieving the target.

20. 5.1 Plan deficiencies - : This section infers that there was inadequate walkthrough / dry run / testing of the upgrade. A major improvement would be to perform such verification work in future but does not appear to be being suggested. Comment on this observation is requested.

Numerous reviews of the plan were held with internal Pathway staff and POCL personnel. However some areas were missed which must be included in plans for any future upgrades of this sort.

-
21. *5.2 End of day markers:* it is assumed that 'outlets' rather than 'counters' is meant here. Given the high percentage (8%+ or 1500+ at rollout) confirmation is sought that such suspension will be automatic and the number of missing EOD markers achieved at 18:00 is requested.
(Supplementary why 18:00?).

As specified suspension will be used in future. The definitive approach to be used in future has yet to be decided as a number of options are available. This will be discussed and finalised at the planning stage for future upgrades.

The figures for the number of missing EOD markers at 1800 is not available.

22. *5.3 Downloading fixes to working outlets:* given the problems, are there plans to introduce either checking in the upgrade that the counter (within outlet) is quiescent or to notify an active counter of an imminent upgrade. Given the volumes at full rollout, it is likely that some counters will be active and tearing down the application in mid-operation could cause subsequent difficulties. A response is requested.

This is closely related to point 21. Glen may be able to add more here.

23. *5.3 Downloading fixes to working outlets:* it is also worth noting that some outlets will not produce an EOD because of a failure in the outlet (e.g. node 4 is disconnected from the LAN in the outlet). Clarification on the decisions that would be made in respect of all such outlets nodes is requested.

If the outlet has a LAN failure then not only will the EOD marker fail to arrive but we will also be unable to contact it via Tivoli. This happened with a number of counters over the weekend and they were updated in the early part of the following week.

24. *5.4 Missing 'D' type reference data files:* it is unclear how, in a managed estate the status of key elements was not known and required investigation. A précis of the amended instructions and any automated pre-checking that will be done in future is requested.

Amended instructions are being provided to OSD regarding the application of 'D' type reference data files. A sample is detailed below.

"OSD supply an .NRG steering file for the virtual Post Office unless specific FAD codes are listed below.

OSD run_install on all agents to copy files to 'delivered' directory, then load to messagestore using Bootle Agent1's r_id_coll job. Confirm success via event log."

25. *5.5 Release notes and work packages - bullet 1:* I was not aware that the estate was not fully managed by Tivoli. The statement could be interpreted to be that this is still the case and as such a timetable is requested for when the estate will be fully integrated and managed by Tivoli as has been set out at various meetings with ICL Pathway.

The whole estate is not currently managed by Tivoli. All of the counter outlets are but not the remaining NT boxes. I am not sure when this will be implemented – Glen may know.

-
26. 5.5 Release notes and work packages - bullet 1: clarification is sought on why this will remove the problem and not just move the discovery window (e.g. when Tivoli starts to manage the rig).

Yes this will probably just move the discovery window but Tivoli knows where it is supposed to install software so will find out at time of installation which was not the case here.

27. 5.5 Release notes and work packages - bullet 2: the absence of such packages suggests additional rigour in preparation of upgrades is required. A walkthrough and test upgrade checklist would seem to be a route for teasing out such basic issues. Clarification on how this situation arose is requested.

The packages were available on one of the signing servers however a problem was experienced with this box so they had to be redelivered onto the other signing server.

28. 5.5 Release notes and work packages - bullet 3: see previous comment. It is unclear why the checking was not completed. Clarification is requested.

The checking was not completed due to human error.

29. 5.5 Release notes and work packages - bullet 4: there are two items for consideration here.

- a) Firstly what would be the impact of distributing the fix twice ?
- b) Secondly a more fundamental question over how the configuration management system failed to trap this fairly basic issue of tracking changes along different branches of a development ?

- a) *There would be no impact in applying a fix twice as it just overwrites code*
- b) *The CM system did not fail to track changes. However a number of changes were made to LT1 just before the LT2 upgrade. As the LT2 baseline had been cut this necessitated some additional checking.*

I assume this actually referred to bullet 5 as it doesn't seem to be relevant to bullet 4

30. 5.5 Release notes and work packages - bullet 5: it is unclear from the statement what the additional LT2 fixes were. Clarification is requested. Additionally, reading the problem summary, it appears a potentially error prone method to use documentation to apply packages. The impression has always been given that the configuration management would result in packages being built, tested and then passed to Stevenage for live dissemination. Clarification on the actual process that was used and will be used in future is requested.

I assume this actually referred to bullet 6 as it doesn't seem to be relevant to bullet 5
There were no additional fixes. A package was delivered for application but all code was already on the system so no change was necessary .

31. 5.5 Release notes and work packages - bullet 6: the resolution of this is requested. Clarification is also sought as to the extent of such OSD CP changes allowed, is it to one or a few machines or across the whole estate ?

I assume this actually referred to bullet 7 as it doesn't seem to be relevant to bullet 6

Nothing contained herein shall be deemed or construed as affecting existing contractual obligations or creating new contractual obligations between ICL Pathway and POCL.

As stated OSD CP changes normally relate to hardware changes or repairs being actioned on the system and as such it is normally on one or a number of machines. The complete process is audited but is under review to ensure that actions taken are visible to all the required parties.

32. *5.6 Reference data rig:* it is noted that the specific faults have been remedied, however, assurance is sought that all other rigs have been verified as being correctly built or aligned to correct status if incorrectly built.

All other rigs were checked and verified before the reference data rig.

33. *5.7 Counter fix application - §1:* it is unclear why the committals were slow and how the delays will be removed whilst maintaining integrity; indeed it would have been anticipated that prior testing would have uncovered such timings. Comment on these observations is requested.

The committals were slow due to the application of the Risposte fix being done via a the single shot method. This is much slower than targeting a number of counters at once. The single shot targeting was started due to error 221 being returned by the counters. The error 221 were returned as the ATE had hung. Glen should be able to add more here.

34. *5.7 Counter fix application - bullet 2:* as ICL Pathway are aware, the difficulty in managing a large estate needs to be largely hands free and automated. The Automated Targeting Engine has been described in previous discussions on systems management as a key element. An update on the investigations, and mitigation should there be further problems with the ATE when the estate is larger is requested.
At 299 outlets, they could be hand updated in a weekend if need be (this is not an option at larger - pre-full rollout - volumes).

Glen should answer this.

35. *5.7 Counter fix application - §4:* the assertion that moving from 8 to 40 concurrent machines (5 fold capacity increase at best *if* linear) in the full solution with a conclusion that “this will obviously increase the ability of the system..” is misleading. The estate will have grown from 299 to 19,000 i.e. 63 fold. Given the events it would lead to a blunt scale up of an additional 500 machines (not 32) - assuming a linear gain. Comment on how this will be managed in reality is requested.

Glen should answer this.

36. *5.7 Counter fix application - § last:* the report does not give any evidence to support, or qualify, the statement “...when the system is working correctly the live estate can be updated effectively”. Quantification of this statement is requested.

Glen should answer this.

37. *5.8 Feedback from SMC:* this issue is not referenced in the body of the report. From some of the other comments, e.g. human error notifying NBSC, it would suggest a wider observation of communication and expectation should have

been planned in, thus in the example of NSBC they would have expected feedback and sought it in this instance of an oversight. Comment is sought on this observation.

As specified the requirements will be more closely specified in future to ensure this type of error does not recur.

38. *5.9 Host resources:* the services that were 'thought' to be running on the Hosts. Has investigations been concluded to confirm this or could there be an unknown underlying cause still present ?

This problem has not recurred.

39. *5.10 DLL failures:* an update on the cause of the DLL problem is requested.

This problem has not recurred but the underlying fault was not identified.