FUJ00078768

Document Title: PERFORMANCE SUMMARY REPORT FOR NEW RELEASE 2

Document Type: Report

Abstract: This document summarises the performance position of New Release 2 for the Acceptance Review process, and in support of the Release Authorisation Board. It brings together the findings of all the contributing activities, and refers to the detailed documents holding the supporting information.

Status: Approved

Distribution: Senior Pathway Management
(and further at their discretion)

Author: J Hunt, Performance Manager

Contributors: Allan Hodgkinson, James Stinchcombe

Approval Authority: Mike Coombs
Deputy Managing Director

Signature: _____

Date: _____

FUJ00078768

0 DOCUMENT HISTORY

0.1 VERSION CONTROL

Version	Date	Reason
0.1	17/03/99	Initial draft - general framework – not issued
0.2	26/03/99	Internal review only
0.3	27/04/99	Formal internal review.
1.0	29/04/99	Issued
1.1	21/05/99	Revised to include input from POCL
2.0	07/06/99	Approved for Performance Acceptance Review

0.2 CHANGES FROM LAST ISSUE

Ref.	Change
2.0	Updated with comments from Bob Booth, POCL. The changes applied to issue 1.1 and issue 2.0 are marked with sidebars in the right-hand margin.

0.3 CHANGES FORECAST

Change	Target Issue
Incorporate comments from the Acceptance Review	3.0

FUJ00078768

0.4 ASSOCIATED DOCUMENTS

Ref.	Library Ref.	Description	Iss.
[20 Counter]	VI/REP/042	20 Counter Test Evaluation Report	1.0
[BUS VOLS]	CS/PER/011	Business Volumetrics	1.0
[BRIEF]	POCL Doc.	Workload Brief 5.3	5.3
[DW014]	DW/TRP/014	Data Warehouse Technical Test Report	1.0
[MOD017]	PA/PER/017	APS Bulk Harvesting Model	0.1
[MOD018]	PA/PER/018	OBCS Bulk Harvesting Model	0.1
[MOD019]	PA/PER/019	TPS Bulk Harvesting Model	0.1
[MOD020]	PA/PER/020	BPS Payment Authorisation Bulk Loading Model	0.1
[MOD021]	PA/PER/021	OBCS Stops Bulk Loader Model	0.1
[MOD022]	PA/PER/022	Change of Nominated Post Office Model	0.1
[MOD023]	PA/PER/023	Reference Data Bulk Loading Model	0.1
[MOD024]	PA/PER/024	Host Processing Model	0.1
[MOD025]	PA/PER/025	Counter to Campus Workload Model	0.1
[MOD026]	PA/PER/026	Workload Volumes Model	0.1
[REP001]	AD/TRP/001	Time Taken to Bring a Counter On-line Test 2.1 Performance Report	1.1
[REP002]	AD/TRP/002	Performance Test 3.1 Report	1.0
[REP003]	AD/TRP/003	Harvesting BES Transactions from Correspondence Server Performance Test 3.2 Report	1.0
[REP004]	AD/TRP/004	Correspondence Server Performance Test 3.4 Report – Riposte Archiving with Message Replication	1.0
[REP005]	AD/TRP/005	Performance Report 3.5/4.5	1.0
[REP007]	AD/TRP/007	Correspondence Server Workload Evaluation Report (Test 3.7)	1.0
[REP009]	AD/TRP/009	TPS Harvesting (Single Agent) Test 4.1 Performance Report	1.0
[REP010]	AD/TRP/010	TPS Harvesting Performance Test Report (3.3 & 4.8)	1.0
[REP011]	AD/TRP/011	Harvesting BES Transactions from Agent Server Performance Test 4.3 Report	1.0
[REP012]	AD/TRP/012	OBCS Agent Harvesting Performance Test 4.4 Completion Report	1.0
[REP014]	AD/TRP/014	Reference Data Bulk Loader Agents Performance Report 4.6	0.1
[REP015]	AD/TRP/015	APS Agent Harvesting Performance Report 4.7	1.0
[REP016]	AD/TRP/016	Performance Test OBCS Stops Bulk Loader Performance 4.9 Report	2.0
[REP017]	AD/TRP/017	Change Nominated Post Office Performance Report 4.10	1.0
[REP031]	AD/TRP/031	OBCS Encashment Harvesting Test 5.1 Performance Report	1.0
[REP032]	AD/TRP/032	BES Harvesting Performance Report	1.0
[REP033]	AD/TRP/033	BPS Payment Authorisation Loading Test 5.3 Performance Report	1.0
[REP034]	AD/TRP/034	TPS Harvesting and End of Day Processing (SE70 Host) Test 5.4 Performance Report	1.0
[REP035]	AD/TRP/035	PAS/CMS Host Performance Test Report	1.0
[REP036]	AD/TRP/036	PAS/CMS NUMA-Q Host Performance Test Report	1.0
[REP037]	AD/TRP/037	CAPS/CAS Batch Test Results	0.2
[REP039]	AD/TRP/039	TIP PC Interface Tests 6.2 and 6.3 Report	1.0

FUJ00078768

Ref.	Library Ref.	Description	Iss.
[REP040]	AD/TRP/040	MIS Interface Performance Test 6.4 Report	1.0
[REP041]	AD/TRP/041	OBCS Foreign Transactions Performance Test Report 7.1	1.0
[REP042]	AD/TRP/042	Performance Test NR2 CAPS On-line Workload Testing Report	1.0
[REP043]	AD/TRP/043	Correspondence Server Audit Harvesting – Performance Test 8.2 Report	1.0
[REP044]	AD/TRP/044	Pathway Performance Report – TPS R2 Design Feedback	1.0
[REP045]	AD/TRP/045	Pathway Performance Report –Design Feedback for TPS Agent	1.0
[REP048]	AD/TRP/048	Interim Report for Performance Test 8.1	1.0
[REP049]	AD/TRP/049	Interim Report for Performance Test 8.3	1.0
[REP050]	AD/TRP/050	Riposte Archiving Investigation – Technical Report	1.0
[Requirements]	PA/PER/012	Performance Assurance Requirements	2.0
[Scaleability]	PA/PER/015	Performance & Scaleability Strategy	1.0
[TED]	TD/ARC/001	Technical Environment Description	4.0
[VOLS]	CS/PER/035	Pathway Performance – Business Volumes	0.4
[VOLSTRAT3]	VolStrat3	POCL TIP Interface Testing	1.0
[VID011]	CS/PRP/011	OBCS Counter Transaction Times for NR2	0.2
[VID012]	CS/PRP/012	BES Counter Transaction Times for NR2	0.2
[VID013]	CS/PRP/013	APS Counter Transaction Times for NR2	0.2
[VID014]	CS/PRP/014	EPOSS Counter Transaction Times for NR2	0.2
[CAPS]	tbs	CAPS to Pathway Interface Testing	1.0
[SRDF]	tbs	EMC Symmetric Performance Tests	0.1

0.5 ABBREVIATIONS

A comprehensive list of terms and abbreviations used by Pathway can be found in the [TED]. Any that are discovered during the review of this document not so covered will as an interim measure be added to the table below.

AP	Automated Payments
APS	Automated Payment Service
AS	Agent Server
BA	(DSS) Benefit Agency
BES	Benefit Encashment Service
BPS	Benefit Payment Service
BCV	EMC^2 Business Continuity Volume
CAP	Cash Account Period
CAPS	Customer Accounting and Payments Strategy
CAS	CAPS Access Service
CMS	Card Management System
CPU	Central Processing Unit
DSS	Department of Social Security
E3	34 Mbps ATM link
EMC	Supplier of disks for Sequent Servers
EPOSS	Electronic Point Of Sale Service

FUJ00078768

FTF	File Transfer Facility
FTMS	File Transfer Managed Service
Gb	Gigabyte
ISDN	Integrated Services Digital Network
Kbps	Kilo*bits* per Second
LAN	Local Area Network
Mb	Megabyte
Mb/sec	Mega*bytes* per Second
Mbps	Mega*bits* per second
MIS	Management Information System
NFS	Networked File System
NINO	National Insurance Number
NR2	New Release 2
NR2+	Next release after New Release 2
NTFS	Windows NT File System
NUMA	Non Uniform Memory Addressing
NUMA-Q	Sequent implementation of NUMA architecture
OBCS	Order Book Control Service
PAS	Payment Authorisation Service
PMS	Payment Management System
POCL	Post Office Counters Ltd
POLO	Post Office Log On
RDBMS	Relational Data Base Management System
RDDS	Reference Data Distribution Service
RDMC	Reference Data Management Centre
RDMS	Reference Data Management System
SLA	Service Level Agreement
SLAM	Service Level Agreement Monitor
SMP	Symmetric Multiple Processor
TED	Technical Environment Description (this document)
SRDF	Symmetrix Remote Data Facility
TIP	Transaction Information Processing
UNIX	Widely used operating system available in a number of variants
VME	Virtual Machine Environment

FUJ00078768

0.6 TERMINOLOGY

This section defines the terminology used in the report.

The workload pattern across the year has been defined using the following terms:

- **Peak day in an average week** – The busiest day for business in a typical week during the year. The busiest day is normally a Monday.

- **Peak on peak** – The busiest day in a busy week. This will normally occur just before Christmas when there is a heavy EPOSS load and double benefits are payable.

- **Peak day in a peak week** - As 'peak on peak'

EMC² – EMC²is a supplier to Pathway of leading edge disc technology. The EMC² discs are used on both the Correspondence Servers and the Sequent Host systems to provide highly available & highly resilient data storage.

BCV – The facility in the EMC² disc system for the fast archiving of data using a third mirror for the archive copy. To copy the data to tape the mirror is broken and the third copy is used for archiving whilst the other two mirrors are used in the live service.

FUJ00078768

0.7 CONFIDENTIALITY AND DISCLOSURE

The information contained in this document is confidential. No copies may be made of it or any part of it, nor may the contents of this document be disclosed in whole or in part to any other party without the prior written consent of the author. It may not be copied to, given to, copied by, or discussed with any non-ICL employee, without first obtaining ICL's express permission in writing. Requests to disclose this document, or any part of it, to any non-ICL employee should be addressed in writing to the Performance Manager, ICL Pathway Ltd, stating the purpose and circumstances.

An exception to this general provision is that this document is intended for review by the Horizon Programme as part of the formal acceptance of the Pathway solution, and in support of the Release Authorisation Board. Prior permission is hereby granted that this document can be made available to those persons directly involved in those reviews, and to the Horizon Technical Test Manager and certain others specified by Horizon who will need to contribute toward the process. All such individuals must then be bound by this general provision, and if they have not already done so must first sign an undertaking of non-disclosure.

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

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

FUJ00078768

0.8 CONTENTS

FUJ00078768

1 MANAGEMENT SUMMARY

The Pathway system architecture is not based on conventional client-server models. Nor does it conform to traditional central-system models. It adopts an entirely original and highly innovative four-tier model that effectively merges the qualities of central systems and client server systems. It is tailored specifically to meet the demands of the Horizon system. This four-tier model comprises:

- Counter Clients
- Correspondence Servers
- Agents
- Host Servers.

Each of these architectural layers has been considered both separately and in combination in assessing the projected loads they will encounter and in modelling and measuring their potential performance capabilities against these projected loads.

For New Release 2 the objective of performance modelling and measurement was to provide assurance that the Pathway system could support, at least, the peak workload generated during the rollout of the first 8,000 outlets. For NR2 a 'safe water marker' of 40% of full rollout volumes (equivalent to 8,000 outlets) was selected as the level at which the system would exceed the demands of the NR2 release after taking into consideration all appropriate allowances inc.:

- Single site operation
- Heaviest day/hour business volumes
- Operational headroom to allow for e.g. re-runs of failed jobs
- Pessimistic workload profile inc. full multi-benefit (see §3.2)
- Management safety net.

Beyond this 'safe water mark' the NR2^{+} release contains a number of developments that will allow the system to support the full rollout volumes on a peak on peak day e.g. a double benefit day before Xmas.

FUJ00078768

The performance modelling and measurement for NR2 confirms that the system can support the peak load generated during the rollout of the first 8,000 outlets including:

- The peak hour and peak day at a 20 counter outlet
- The peak hour at the datacentre during the peak working day.
- The peak daily overnight processing inc. bulk harvesting and loading of messages into TMS on the Sequent SE70
- The peak day's interactions with the client systems (file transfer and on-line)
- The peak day at the data warehouse
- Operational backup at peak volumes

The performance modelling of the system was based on the contracted volumetrics that include a rollout rate of 300 outlets per week and the full implementation of multi-benefit by card.

The OBCS service has been measured and modelled against the peak daily load. The peak load on the OBCS service occurs close to 8,000 outlets as cards replace order books (see diagram).

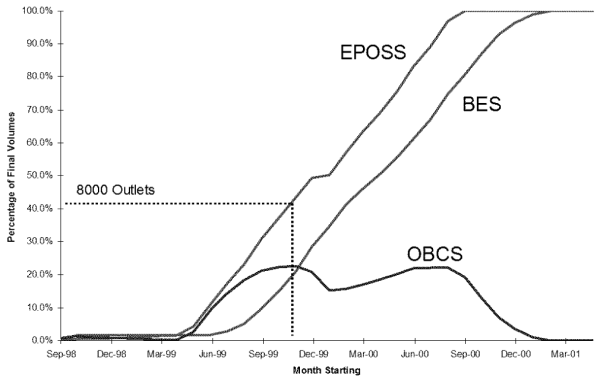

The performance measurement of the system also included many tests that evaluated the behaviour of the system under the peak workload generated by 20,000 outlets at the end of rollout including:

- Peak daily load on a 20-counter outlet with 100% payment by card
- Correspondence server peak hour
- Peak day BES harvesting/loading of messages
- Peak on peak day overnight processing of PAS/CMS on NUMA-Q
- Peak day file transfer to client systems inc. Automated Payments

FUJ00078768

where the current volumes are 2.5 times the volumes in the Workload Brief [BRIEF].

- Peak hour on-line access from client systems.

The results from the measurement and modelling projects have provided the Pathway design and implementation teams with the business and message data volumes and mixes that are required in the planning of the rollout programme to 20,000 outlets. Included in the plans for releases beyond NR2 are activities to replace or enhance components of the system e.g. the Host systems, the Correspondence Servers and the WAN routers to meet the service requirements beyond 8,000 outlets. These developments are compliant with the capacity management strategy defined in the Pathway Performance and Scaleability Strategy [Scaleability].

Pathway has also reviewed the capacity requirements for applications beyond NR2 and is implementing plans to deliver enhanced versions of a number of components including:

- the Reference Data loaders,
- Change Nominated Post Office Agents and
- Riposte message server.

The measurement and modelling of the system used the Pathway workload profile [VOLS] derived from the Business Volumes [BUS VOLS]. The Pathway workload profile includes models of the peak load on each of the services. These models can only be verified against the production system as very little workload data is available to Pathway from the current manual or ECCO systems. The workload models will be refined using data collected from the live system during Live Trial and the early stages of rollout. The data in the models will be refined to more accurately reflect the production environment so that the workload can be modelled with increasing accuracy as it grows towards 20,000 outlets. In some areas e.g. AP transactions, this process is already underway and the volumes used in testing and modelling have been increased in-line with current operational volumes.

FUJ00078768

2 PURPOSE & SCOPE

This document brings together the findings of all the various performance activities conducted for NR2 as a summary of the performance status for the release.

It is intended to serve as the principal report for performance in support of the Acceptance Review process and for the Release Authorisation Board.

It directly references the separate detailed reports produced for all the various performance activities that serve to support this document.

FUJ00078768

3 PERFORMANCE STATUS

3.1 INTRODUCTION

This section walks through all the major components of the solution for which performance related Acceptance Specifications apply. It further walks through all other major components of the system. For each it details (where applicable) the following:

- A brief summary of the current performance status
- The performance assurance requirements specified and agreed, together with the source reference.
- Status of the performance tests conducted
- Their principal findings
- Report reference
- Any performance modelling carried out
- Its findings
- Source Reference

In most cases the timings are given for single site operation (where appropriate) as the objective for Pathway is to ensure that even when one site is out of operation the remaining site can support all of the workload through to full rollout and can maintain the SLAs.

FUJ00078768

3.2 BUSINESS MODEL

For NR2, Pathway has undertaken a considerable amount of:

- Testing/measurement and
- Modelling

of the system. To measure or model a system requires a definition of the workload that is being measured or modelled and in particular a definition of the peak loads that will be applied to the system during each period of the rollout. The definition of the business model given in [BRIEF] does not define the peaks in the workload.

Before Pathway could undertake the modelling and measurement of the solution, the workload was analysed, modelled and documented in [BUS VOLS]. This process involved identifying:

- The workload volumes at pre-determined points during rollout e.g. 8,000 outlets
- The peak hour or peak day load at pre-determined points during rollout
- The workload volumes that would be generated with a modified rollout schedule e.g. if multi-benefit by card is phased in, the delay in the introduction of some benefits by card will reduce the BES workload but increase the OBCS workload during rollout.

Pathway has modelled the workload based on all of the available data. During the process of modelling the workload, many assumptions had to be made about the nature of the business. Pathway has been unable to confirm with the sponsors that the assumptions are valid. A very safe (low risk) approach has therefore been adopted to modelling the workload (see below) that will ensure that the NR2 system has the maximum headroom and can cope with changes to the workload profile. The peak workload volumes for NR2 are documented in [MOD026].

The workload model used in the modelling and measurement of the system has been fixed (see definition below) as any deviation from the workload model would affect both testing and modelling and make the evaluation and presentation of results difficult. In particular, the use of a workload profile that is not changing is essential for developing the data, scripts and tools used in the testing programme.

FUJ00078768

With a constantly changing profile it would be very difficult to assess the affect of changes to applications, systems or the infrastructure. In almost every case (with the exception of OBCS) the workload profile which has been adopted will generate a significantly heavier load on the system than is now expected. All of the measurements and modelling work is based on the following rollout model:

- National Rollout starts on 5th April 1999 [1].
- The rollout rate in the first 6 weeks of rollout is 100, 100, 150, 150, 200 and 250 Post Offices respectively.
- The rollout is then 300 Post Offices per week until 22nd November 1999.
- Over the 7 week Christmas period the rollout is 150, 150, 150, 0, 0, 0, and 150 Post Offices respectively.
- The rollout is then 300 Post Offices per week until all Post Offices are rolled out.

The rollout model assumes that for benefit payments:

- Multi-benefit by card is in operation across all 20,000 outlets.
- Child Benefit by card is available at start of National Rollout.
- All other benefits are card enabled within 3 months of the start of National Rollout.
- The number of people collecting each type of benefit is as given in [VOLS].
- During the first 4 weeks after a Post Office is automated no beneficiary is transferred from payment by book to payment by card.
- All beneficiaries in a Post Office are then transferred from payment by book to payment by card over the next 20 weeks.

The Automated Payments model assumes that 200M payments are made per year not the 68M payments defined in the Workload Brief [BRIEF]

1 Since this decision was made, the dates and the workload have changed. The workload defined here is pessimistic when compared with the expected rollout. The workload generated before Christmas '99 will be lighter than the expected workload with 8,000 outlets automated.

FUJ00078768

3.3 APPROACH TO PERFORMANCE

The requirements on the Pathway solution are very demanding, particularly during the roll-out programme where the system will grow from less than 200 outlets to almost 20,000 over less than 24 months.

Pathway therefore must deliver a well-engineered solution that is:

- predictable
- scaleable and
- manageable

A performance engineering approach must be applied to all key hardware, software, application and network components of the solution.

The approach to performance within Pathway has been structured to support the requirements of development and delivery. The four key strands to this are :

- capturing an understanding of the system inc. business volumes
- modelling inc. peak workload volumes
- performance testing
- performance management of the live system

3.3.1 Capturing an understanding of the system

The volume of work executed by the Pathway system will vary depending on a number of parameters including the time of the day and the day of the week.

A significant proportion of the workload comes from the payment of benefits most of which are available to be paid on Monday and Thursday. This has the effect of skewing the workload over the week with the number of benefit payments made on a Monday (the day on which retirement pension and child benefit is due to be paid) predicted to be approximately 50% more than those made on a Friday.

There are also seasonal variations to benefit payments, due to public holidays which typically fall on a Monday, and variations in EPOSS transaction volumes, such as at Christmas.

FUJ00078768

The modelling and performance measurement activities require accurate information about :

- peak workload volumetrics
- the hardware and software configured
- the implementation of services
- the resource usage of each function used by the services

Building a workload profile, and identifying the peaks in the workload, is key to understanding the load which the Pathway solution must be able to support.

3.3.2 Modelling

The Pathway resource models use:

- the configuration of the system
- the business load placed upon it
- the messages generated by each of the business functions
- the performance measurements from Performance Testing

as inputs to the modelling work (see diagram).

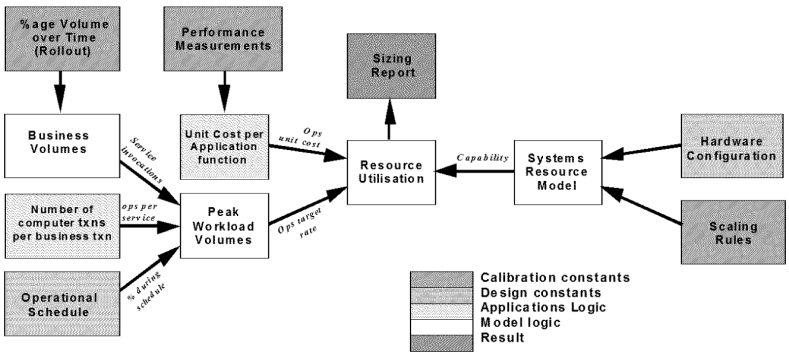

FUJ00078768

The models are initially implemented as *engineering models* that form part of the design evaluation process. As measurement data becomes available the models are refined and the impact of any changes assessed. Later as volumetrics and performance data becomes available from the live system, the models will be further refined. This process will, in the longer term, feed into the performance management precess.

3.3.3 Performance evaluation

Central to the effectiveness of the model are the measurements of resources (CPU, disc, memory etc.) used for each operation performed by the system.

The Performance Test team is an integral part of the process. This team designs and develops the workloads and scripts that are used to:

- collect performance data to calibrate the models
- evaluate the systems under load and
- demonstrate that performance critical components can support peak throughput

In a complex multi-layer solution like Pathway, volume end-to-end testing is not a practical proposition. Therefore to yield the information required, Pathway has devised a test plan covering:

- The performance evaluation of design options (prototyping/ design feedback)
- The calibration of modelling assumptions
- The measurement of components/subsystems under load
- The effect on performance of workload interaction

The performance tests have been constructed in such a way as to ensure that the:

- Input & output data flows
- Input & output concurrency and
- The work processed within the platform

are representative of the full system and in particular the input to the simulated environment must “look” the same to the *system under test* as the production environment.

FUJ00078768

3.3.4 Performance management

Whilst the Performance Test team can measure performance using representative workloads, the only true test of performance is the production service.

To ensure that the production system have the capacity to deliver the services and to identify potential *hot-spots,* a comprehensive set of performance monitoring tools will be delivered as part of the systems management framework for the system.

Tools will be developed to analyse and report on the data collected.

FUJ00078768

3.4 COUNTER SYSTEMS

3.4.1 Summary

The counter platform has already been scaled up significantly to a Pentium II, 400MHz, 64 Mb RAM, and 4.3Gb HDD. A further increase in memory (128Mb) and disk size (10Gb), prior to National Rollout has been approved. This is not to address performance issues, but rather to avoid potential costly and disruptive changes of the outlet hardware any earlier than absolutely necessary. It will, nonetheless, help to resolve any potential unforeseen performance issues that may arise in live running.

The counter platform has been significantly stressed. Performance problems have been revealed and dealt with. Overall the platform copes well with peak loads projected for NR2. The 20-counter test rig has successfully emulated the load on an office on

- The first day after the office is automated when a heavy load of OBCS 'foreigns' is predicted
- A peak on peak day at the end of rollout with multi-benefit by card in operation,

and has tested all of the back-office functions. The testing emulated 13 weeks of peak operation so that the message store grew to a size well beyond that of an operational of a 20 counter outlet. For 12 of the 13 weeks the peak on peak workload was run each day which stressed the counter system well beyond the volumes that would be processed in a live 20-Counter outlet. The message store was populated with a typical mix of messages as at full rollout.

This testing has demonstrated that there is ample headroom to spare for most customer facing operations. In particular, early fears that the gateway counter may suffer adversely under load have been shown to be ill founded. The 20-counter rig is being retained by Pathway and it will be used for stress testing future releases.

The 20-Counter tests included an activity to build a representative counter Message Store. This was achieved by running scripts on the counter terminals that simulated operations at the counter. As messages for more than 10 weeks were required, some of the scripts were speeded up by reducing the *think-type* time for each counter operation to reduce the elapsed time of the Message Store build process. This had the effect of increasing the processing rate to many times that which could be achieved in a normal outlet even on a peak on peak day. Both Riposte and the counter systems demonstrated that they could support this heavily over-scaled workload.

FUJ00078768

One problem area which will require management attention in the early stages of the Live Trial, is that of counter replacement (e.g. hardware swap out following failure). Here the tests and modelling together project that counter replacement would run too slowly for large outlets (greater than 6 counters) using simple Riposte bulk replication to reproduce a fully populated Message Store. This should not pose a threat for the existing outlets when they migrate from 1C to NR2, but could be problematic with the larger outlets which are rolled out shortly after the start of Live Trial. Initially even these would not be affected. Only when the messages accumulate to significant levels would the problem emerge (after about 4 weeks).

An alternative non-Riposte approach to introducing these large Message Stores has been developed and tested, under CP 1541. This will be intercepted prior to rollout of the new offices during Live Trial and so will avoid the problem entirely.

3.4.2 Performance Assessment

Section 2 of [Requirements] specifies the agreed set of tests. These consist of:

8.1 Installation/replacement of a counter terminal

8.2 Sizing of the counter terminal disks

8.3 Gateway counter terminal in a large outlet

8.4 ISDN links

8.5 Counter terminal non-functional performance

8.6 Printing receipts

8.7 Stock control and accounting operations within an outlet

FUJ00078768

3.4.2.1 *Installation/replacement of a counter terminal*

Test 2.1 was performed by the Performance Test team. The results of the test are reported in document [REP001]. This test revealed that for outlets with large Message Stores the planned method of establishing the Message Store, using Riposte bulk replication, was simply inappropriate. A new approach has been adopted (see CP1541) which avoids the use of Riposte bulk replication, using instead a simple file copy to generate the bulk of the Message Store. This new approach has been tested before it is released and on a 20-Counter outlet the time to restore the message store is <20 minutes. The 'old' technique will be used initially but CP1541 has been implemented and is in the final stages of testing. The revisions have been downloaded to the counters using Tivoli and will be enabled when the testing is complete and the change approved.

3.4.2.2 *Sizing of the counter terminal disks*

The discs in the counter PCs have grown in size from 1.6Gb through 4.3Gb to 10Gb [2]. The latest increase in size will be introduced at the start of national rollout and was introduced in order to reduce the possibility of a change becoming necessary in the future because of a new requirement.

The message store on the PCs is predicted to grow to less than 700Mb in the largest outlet. The message store on the counter PCs used in the 20-Counter test has grown to approximately 1Gb in size. The difference is due to the workload being run on the 20-Counter rig i.e. the peak on peak workload is run every day to stress the system. In a normal outlets the peak on peak workload would only occur on <5 days per year.

---

2 For single counter offices two 10Gb discs (one exchangeable the other fixed) will be configured.

FUJ00078768

3.4.2.3 20-Counter Testing

The 20-counter test was designed to demonstrate that a large outlet could support the heaviest load predicted and that all of the outlet back-office functions could support the load generated in a very large office with acceptable performance.

Tests:

2.3 Gateway counter terminal in a large outlet

2.5 Counter terminal non-functional performance

2.7 Stock control and accounting operations within an outlet

are all covered by the '20 Counter Tests' performed by the System Test team. This set of tests has completed one cycle during which the testing has been considerably expanded to provide a full stress test of a large outlet including:

- Installation of a new counter terminal or replacement of counter terminal or the disc in conjunction with peak daily operations at the other terminals.
- Measure the responsiveness of the Gateway Counter terminal during busy workload and interactions with the Correspondence Server.
- Measure the impact of Cash Account production on Gateway PC and during busy workload and interactions with the Correspondence Server.
- Measure Counter performance at busy times, transactions to include EPOSS, APS, BES encashments and OBCS encashments, foreigns.
- Measure the responsiveness of Stock control procedures, with 38 Stock units.
- Measure the impact of balancing 20 Stock Units simultaneously, rolling into new balance period and Cash week
- Measure how long it takes to produce a Trial Balance
- Measure how long it takes to roll the office balance into the next Cash Account period
- Measure the time taken to perform end of day processing with a full days activity recorded in the message store.
- Time report production in a busy outlet – including Daily/Weekly Stock Unit summaries, Daily office reports, Weekly office reports
- Measure timings for Administration tasks, including

FUJ00078768

- Postmaster log on/off
- Clerk log on/off
- Clerk Password change (including concurrent change)
- Suspend/Resume session
- Switch to Training session
- Add new users/Stock Units
- Allocate/Deallocate a Stock Unit
- Stock Unit enquiries
- Reports by Stock Units
- Postmaster Change of Password

• Concurrent Card swipes, barcode reads, foreigns across all 20 counters

For completeness the test team also ran the full set of ad-hoc reports to ensure that the performance of these operations was acceptable in a large outlet. The results of the tests show that these reports are produced in an acceptable time.

The results from the first cycle of tests are given in an interim report [20 Counter]. These conclude that the counter systems hold up well in all major respects under the peak on peak projected loads. Indeed these tests overstress the platform as the mode of operation, using automated data entry, operates much faster than would be the case with human operators.

Early tests revealed some functional problems in the areas of stock unit rollover and office CAP rollover when the system was operating under stress. The issues have been investigated and the problems corrected and the performance of these functions is more than acceptable. These problems aside, the Office Rollover process at these volumes has been timed at just 20 minutes.

For this testing the Gateway counter was also used as a normal counter position which is unlikely in a large office. Even with this additional load, processes on the Gateway counter were not significantly impacted by volume traffic on the slave counters and related ISDN traffic. Similarly, operation of the ISDN gateway function, and servicing of messages from the slave counters is not significantly impacted by processes on the Gateway counter.

FUJ00078768

3.4.2.4 ISDN links

The 20 Counter test (see §3.4.2.3) included testing of the ISDN link from the counters to the campus. The 20 Counter Day 1 test was representative of the load generated by OBCS 'foreigns' on the first Monday after a 20 Counter outlet is automated. The tests showed that the ISDN configuration in the Gateway PC could support the peak NR2 workload.

Beyond the high connection rate load generated by NR2, test 2.4 cannot be economically tested without using the live network. As no immediate risk arises from this area, it is planned that this test will be run during the Live Trial, when an appropriate 'quiet' period will be selected to conduct a bulk connection test of the ISDN network and the Routers (ee also ISDN NETWORK (§3.11)). This will stress concurrent ISDN connectivity to projected levels for full rollout.

Some outlets will be connected using Frame Relay at NR2. Frame Relay was not included in this testing programme as in operation it is very similar to ISDN with the exception that response times using Frame Relay will be reduced because there is no call set-up time. Riposte uses the same connection method for both ISDN and Frame Relay. Testing therefore focussed on the performance of the worst case connection method, i.e. ISDN.

3.4.2.5 Counter Test 2.5 to 2.7

Tests 2.5, 2.6 and 2.7 include a number of fixed overheads. Quite separately from the '20 Counter Tests' these have also been covered in part (for small outlets with low transaction volumes) by the Model Office Test. Early runs indicated that the timings for report production (balance reports, etc.) were not acceptable. This was traced to a build error, which had left the back office printers configured with their default factory settings. Once corrected, the timings obtained were reported as more than acceptable by the customer during the model office trial.

3.4.2.6 Printing receipts

The printing of receipts was an integral part of 20-counter testing. On most of the counters the printers were emulated but several real printers were configured for use during the tests.

All types of receipts were tested and printed. The test results showed that, even under heavy load conditions, printing commenced in less than one second (assuming the slip is in the printer).

FUJ00078768

The tests also showed that there are no delays in printing and the printer prints at its full rate even though the baud rate of the connection to the printer was reduced to improve reliability with the automated scripts (Note : the change in baud rate applies to the test only).

FUJ00078768

3.5 CORRESPONDENCE SERVERS

3.5.1 Summary

The Correspondence Server layer of the architecture has been extensively tested and modelled. A significant amount of the testing has used workload volumes and Message Stores that are representative of full rollout.

The tests and modelling conducted for the NR2 implementation of the Correspondence Server platform indicate that many of the Correspondence Server processes, including counter replication, are capable of operating at close to full rollout volumes. The single exception is when Riposte Archiving is run in parallel with other Riposte activity (see §3.5.2.4)

Testing initially concentrated on measurements of each service functioning separately to provide resource data for modelling. However, later tests (Tests 3.4 and 3.7) have emulated the Correspondence Server running the set of services that would run concurrently in live operation. The results indicate that the Correspondence Server can easily support the workload generated on the heaviest day during the rollout of the first 8,000 outlets, with significant headroom for expansion.

During the development of NR2, the Correspondence Server configurations have been upgraded with EMC^2 disc sub-systems that provide greatly enhanced resilience features. A by-product of the move to EMC^2 disc technology is a gain in performance because the EMC^2 disc sub-systems are configured with large, highly resilient caches that provide higher throughput for many Correspondence Server operations. In addition, the use of EMC^2 BCVs allows for a revised Correspondence Server backup and restore strategy that reduces the time a Correspondence Server node is out-of-service for backup. This will become more important in the future as the opening hours of the outlets lengthens.

For NR2, Pathway has also intercepted new Correspondence Server hardware. On each campus there are four pairs (master and backup) of Correspondence Servers providing fast failover in the unlikely event of a failure. The backup Correspondence Servers will use COMPAQ 4-way Intel Xeon servers. The additional power provided by new servers is not required at NR2, but the move will ensure that there is sufficient headroom and contingency for full rollout. The master Correspondence Servers will be upgraded later in the programme but well before the additional power is required.

FUJ00078768

It should be noted that all the tests conducted in this area so far have excluded the effects of ISDN delays and connect times. It is planned that this area will be explored thoroughly in future tests for volumes > 8,000 outlets, using sophisticated simulation techniques. In the meantime, whilst it is likely that there will be an effect on performance, it is relatively small. With so much headroom demonstrated this certainly does not constitute a risk at NR2 volumes. For further details of ISDN testing and Modelling see §3.11 - ISDN NETWORK

3.5.2 Performance Assessment

Section 3 of [Requirements] specifies the agreed set of tests. These consist of:

8.1 Replication from Counter System to Correspondence Server

8.2 Harvesting BES transactions from the Correspondence Server

8.3 Harvesting TPS transactions from the Correspondence Server

8.4 Concurrent Riposte Replication and Riposte Archiving

8.5 BPS Payment Authorisation Bulk Loader

8.6 Correspondence Server Riposte Archiving

8.7 Correspondence Server Workload Evaluation

3.5.2.1 Replication from Counter System to Correspondence Server

Test 3.1 was performed by the Performance Test team. The results of the test are reported in [REP002].

This test measured the rate at which Riposte messages could be replicated from the gateway PC in the outlet to the Correspondence Server, and the load on the Correspondence Server generated by many outlets performing this process concurrently.

The results of this test were not as expected. On investigation it was discovered that our understanding of the Riposte interactions between the Gateway Counter and the Correspondence Server was deficient and as a result the test had been incorrectly configured. The lessons learnt from this test were taken into account when preparing test 3.4, which was extended to cover the objectives of test 3.1. The results for 3.4 are reviewed in §3.5.2.4, which also covers the results of test 3.6, Riposte Archiving.

Further improvements in the method of simulating Counter to Correspondence Server replication were implemented in Tests 8.2 (see §3.10.2.2and Test 3.7 (see §3.5.2.7).

FUJ00078768

The tests that measured the peak hour replication load demonstrated that the Correspondence Server platform could just cope with replication for projected loads at full rollout volumes in the peak hour. This is well in excess of the demands the system will encounter for NR2.

3.5.2.2 Harvesting BES transactions from the Correspondence Server

Test 3.2 was performed by the Performance Test team. The results are reported in [REP003].

The test measured the rate at which BES messages could be harvested from a Correspondence Server, and the load on the Correspondence Server generated by one or more bulk harvester agents harvesting BES messages concurrently.

This test demonstrated that the Correspondence Server can cope comfortably with projected BES Harvesting loads up to and beyond full rollout volumes.

The testing method used in Test 3.2 did not include the cycling of the Bulk Harvester Agents which will take place in the live service every 15 minutes or so. However, as the number of messages that are generated during the heaviest hour of heaviest day during the rollout of the first 8,000 outlets (460K) take less than 10 minutes to harvest when one site is in operation, there is a significant amount of headroom remaining in the system. With so much headroom this does not constitute a risk at NR2 volumes.

See also BES Agent (§3.6.2.2) and BPS Host (§3.7.2.2).

3.5.2.3 Harvesting TPS transactions from the Correspondence Server

Test 3.3 was performed by the Performance Test team. It was combined with Test 4.8, and the results are reported in [REP010].

This test measured the Correspondence Server under the load that would be expected on the heaviest day in an average week at full rollout i.e. approximately 40M TPS messages would be harvested in total. The expected peak on peak rate at full rollout is approximately 53M messages and the TPS harvesting model [MOD019] predicts that on the heaviest day during the rollout of the first 8,000 outlets <15M messages would be harvested.

This test demonstrated that the Correspondence Server can cope comfortably with projected TPS Harvesting loads up to and beyond full rollout volumes.

See also TPS Agent (§ 3.6.2.1) and TPS Host (§3.7.2.4).

FUJ00078768

3.5.2.4 Concurrent Riposte Replication and Riposte Archiving

Test 3.4 was performed by the Performance Test team. The results are reported in [REP004].

This test combined Tests 3.1(Riposte replication) and 3.6 (Riposte archiving). For this test the deficiencies identified in Test 3.1 were removed so that the mechanism used by Riposte for replicating messages from the outlets to the Correspondence Server was representative of the live system.

During the peak hour at full rollout, approximately 3.6K messages per second will be replicated from the counters to the Correspondence Server, i.e. approximately 0.9K messages per second per Correspondence Server.

The tests that measured the peak hour replication load demonstrated that the Correspondence Server platform could support loads up to 1.1K messages per second. This is well in excess of the 0.2K messages per second per Correspondence Server that the model [mod024] predicts that the system will encounter at NR2.

However, the test also revealed a problem in running Riposte Archiving alongside general Riposte Replication, under stress. The problem was severe and has been referred to our supplier, Escher. As a result Pathway has changed the planned operational schedule for running Riposte Archiving, such that it will be run in dedicated time, not competing with other Riposte activity on the Correspondence Server. There is more than sufficient time to schedule this with NR2 volumes, and potentially enough to cope right up to full rollout volumes, though this position will need to be monitored beyond 8000 outlets. This approach entirely avoids the problem.

Test 3.4 was extended further in Test 3.7 (see §3.5.2.7)

FUJ00078768

3.5.2.5 BPS Payment Authorisation Bulk Loader

Test 3.5 (Correspondence Server) has been combined with Test 4.5 (Agent) and is being conducted by the Performance Test team. The results are reported in [REP005].

The capability of the CAPS Payment Authorisation Data loader to support full workload volumes was modelled early in the development cycle. The model identified that the performance of CAPS Payment Authorisation Data loading would be disc limited due to the large number of Riposte index updates and would not achieve the required level of throughput. A performance review resulted in CP1510 being raised. This CP reduced the number of indexes that would be updated by the load process. At the same time, the rate at which disc transfers can be performed has been significantly increased by the introduction of EMC^2 disc arrays for the Riposte Message Store. The EMC^2 arrays have large resilient caches that smooth the write traffic to the discs, increasing the maximum disc throughput of the Correspondence Servers.

The CAPS Payment Authorisation Data loader was measured using the expected load that will be generated on the heaviest day during the rollout of the first 8,000 outlets. The measurement test loaded 950K payment authorisations into a Correspondence Server. This test assumed that all of the high volume benefits would be paid by card. With NR2 rollout volumes up to 8,000 outlets, the model [MOD020] predicts that payment authorisation volumes will be <900K on the peak night even if multi benefit is applied i.e. approximately 220K payment authorisations per Correspondence Server. With Child Benefit only, the payment authorisation volumes are even lower. The test of a single Correspondence Server therefore more than covers our position for NR2 volumes.

The test demonstrated that 950K payment authorisations were loaded into a Correspondence Server cluster in 1.5 hours, including the indexing of all new messages in Riposte. The test measured bulk loader agents loading Payment Authorisations into both nodes of the Correspondence Server cluster concurrently. For 8,000 outlets the expected elapsed time to run the Payment Authorisation bulk loader is <30 minutes in single site operation.

FUJ00078768

3.5.2.6 Correspondence Server Riposte Archiving

Riposte archiving is the process by which messages that have expired (i.e. have exceeded the time a message should remain in the Message Store) are deleted from the Message Store. Each node in a cluster of Correspondence Servers is archived separately.

Because of the problems described in §3.5.2.4, the measurement of Test 3.6 was delayed pending the resolution by Escher of the problems identified in this test. The Performance Test team has produced an interim report [REP050] on the tests that have been run on Riposte internal archiving. Further testing will be undertaken before the start of national rollout.

An engineering model produced to evaluate the NR2 design identified that the proposed design would not enable the Message Store to be archived within the time available within the full rollout schedule. A design review with our supplier, Escher, identified a number of enhancements to Riposte archiving that will be implemented in the Riposte 6 release.

The release 5.4 implementation of archiving was tested using a Message Store fully populated with the messages, persistent objects, reference data, etc for 2,000 outlets [[3]]. The tests confirmed that each Correspondence Server could be archived in less than 10 hours per week. The Correspondence Servers will be archived in parallel. The NR2 schedule has allocated 15 hours to archive the Correspondence Servers.

[3] At 8,000 outlets, each of the four Correspondence Servers will support approximately 2,000 outlets

FUJ00078768

3.5.2.7 Correspondence Server Workload Evaluation

Test 3.7 represented the next stage of stress testing and builds upon Test 3.4. This test measured the load on a Correspondence Server cluster during the peak hour on the busiest day during the rollout of the first 8,000 outlets. This test combines the workload components from earlier tests into a stress test of the Correspondence Server i.e.:

- Replication from the counter systems to the Correspondence Server (Test 3.1 & 3.4)
- Replication of messages between Correspondence Server nodes
- BES Harvesting (test 3.2)
- Gathering the Audit tracks (Test 8.2)
- Processing OBCS 'foreigns' (Test 7.1)

Test 3.7 was performed by the Performance Test team. The results are reported in [REP007].

This test used the full NR2 Correspondence Server configuration (CP1498) including EMC^2 discs. The use of EMC^2 discs removed the 50Gb size constraint from the Message Store that had applied to earlier tests. This change permitted the construction of a Message Store with a full set of messages including persistent objects. The Message Store used for this test represented the volumes of messages generated by 8,000 outlets.

The first part was a confirmatory test for Riposte Replication only that demonstrated that the revised EMC^2 based platform does not degrade performance. The remaining tests were geared towards stressing the platform under mixed workloads and to confirm that Riposte performance does not suffer as a result of operating mixed loads under stress.

The tests demonstrated that even on the heaviest day at the end of the rollout of the first 8,000 outlets, the NR2 Correspondence Servers can support the peak hour workload without any degradation. In particular the test confirmed the viability of the EMC^2 based platform.

The results also confirmed that the current Pentium Pro 200Mhz Correspondence Server platform would not have sufficient power to support the full rollout peak on peak volumes. Pathway has already invoked the scalability plan for the Correspondence Server platform (see CP1814). This new platform will be stress tested during a future phase of performance testing.

FUJ00078768

3.6 AGENT SERVERS

3.6.1 Summary

The tests and modelling of the Agent Server platform indicate that it is capable of supporting up to full rollout volumes for most of the Bulk Harvester and Loader operations. The modelling of the TPS Bulk Harvester indicates that more instances of the TPS Bulk Harvester may be required to support single site operation and peak on peak load.

3.6.2 Performance Assessment

Section 4 of [Requirements] specifies the agreed set of tests. These consist of:

8.1 TPS Harvester Agent - Single Server

8.2 TPS Harvester Agent – Multiple Server

8.3 BES Bulk Harvester

8.4 OBCS Encashment Bulk Harvester

8.5 BPS Payment Authorisation Bulk Loader

8.6 Reference Data Bulk Loader

8.7 APS Bulk Harvester

8.8 Cluster Look Up Agent

8.9 OBCS Stops Bulk Loader

8.10 Change of Nominated Post Office

FUJ00078768

3.6.2.1 TPS Bulk Harvester Agent

Tests 4.1 and 4.2 were performed by the Performance Test team. Some time after the tests had been completed it was discovered that a number of the test cases had been incorrectly configured and some of the Agent processing (i.e. the Host TPS database accesses) was not included in the measurements.

However, TPS has been the subject of a number of tests, so the results from 4.1 that remain meaningful, were combined with the results from the design feedback tests, [REP044] and [REP045], to produce the agent model. The tests documented in [REP044] used a number of Agent servers running concurrently to load the SE70 Host. The results for test 4.1 are reported in [REP009]. The results from Test 4.2 all contained the fault so this report has been withdrawn. A re-test will be scheduled to re-run the test at some future point in order to confirm behaviour patterns for multiple servers with a full configuration. However, this work is not urgent as it relates to high volumes beyond those projected for NR2.

These results, coupled with modelling projections [MOD019], indicate that 16 Agent servers can easily cope with the projected NR2 loads of <15M messages. On the heaviest day during the rollout of the first 8,000 outlets, TPS harvesting will complete well before the 20:30 cut-off time. The agent model [MOD019] predicts that whilst the Agent servers remain moderately loaded during TPS harvesting (<40% cpu loading), the Host cpu is almost fully utilised by 64 TPS bulk harvester agents running concurrently. The limiting factor in throughput is therefore the Host and not the Agent.

However, when only one site is in operation the number of scheduled TPS Bulk Harvester processes currently scheduled is insufficient to process the heaviest projected number of TPS messages before the 20:30 cut-off defined in the NR2 Maestro schedule. Up to 6,500 outlets can be harvested before 20:30 with only one site in operation.

The combination of a site out with a peak workload night is unlikely to occur but Pathway has to plan for such eventualities. If such an event occurs, the cut-off time will be extended by a period of approximately 30 minutes. This will have no impact on the delivery of files to BA or POCL/TIP. Any impact will be taken by Pathway by delaying the delivery of the TPS output to the Pathway Data Warehouse.

FUJ00078768

In the long-term, to support workloads beyond that generated on heaviest day by 6,500 outlets with only a single site available, and without compromising delivery times to the Data Warehouse, Pathway is reviewing a number of options including:

- Increasing the number of TPS bulk harvesters per Agent server and
- Upgrading the specification of the Agent servers
- Implementing separate Maestro schedules for single and two-site working.

There is sufficient spare processor power in the NR2 Agent server configurations to support the number of harvester processes required to complete the task before the 20:30 cut-off up to 8,000 outlets.

Further evaluation of the Agent servers running TPS harvesting is planned. This testing will include testing using higher specification Agent servers against a NUMA-Q Host, as more TPS bulk harvester agents may be required to keep the NUMA-Q fully utilised.

See also TPS Correspondence Server (§3.5.2.3) and TPS Host (§3.7.2.4).

3.6.2.2 BES Bulk Harvester Agent

Test 4.3 was performed by the Performance Test team, and the results are reported in [REP011].

The BES Bulk Harvester agent operates throughout the working day (07:00 to 20:00). One BES bulk harvester agent runs on each of the 16 agent servers.

This test demonstrated (making projections from 3000 outlets up to 20000 outlets) that Pathway can cope comfortably with BES Harvesting loads up to full rollout volumes, peak on peak.

At NR2 the TPS harvesting model [MOD019] predicts that the number of messages that are generated during the heaviest hour of the heaviest day during the rollout of the first 8,000 outlets (460K) will take less than 10 minutes to harvest when one site is in operation. As there is little other activity on the Agent servers 09:00-17:30, the utilisation of the Agent servers whilst the BES bulk harvesters are running will be very low. The TPS harvesters are loaded at 12:00, but the volume of data harvested from the outlets that close before 17:30 is small and there is more than sufficient capacity to support concurrent TPS & BES Harvesting and interactive processes during the afternoon even on a peak day.

FUJ00078768

See also Harvesting BES transactions from the Correspondence Server (§3.5.2.2) and BPS Host Processing(§3.7.2.2)

3.6.2.3 OBCS Bulk Harvester Agent

Test 4.4 was performed by the Performance Test team, and the results are reported in [REP012].

The tests demonstrated that Pathway can cope with harvesting OBCS encashments up to the peak predicted OBCS volumes.

However, it should be noted that this test was based on OBCS dealing only with a maximum of 25% of the benefit load. If the system remains Child Benefit only, right through to full rollout volumes, then this assumption is broken, with OBCS accounting for more like 90% of benefit payments. The results still leave us comfortable under these circumstances at the 8,000 outlet level, but beyond that the schedule would need to be adjusted to allow OBCS processing to benefit from the missing BPS load.

At NR2 volumes with multi-benefit by card in operation the OBCS harvesting model [MOD018] predicts that the number of messages that are generated during the heaviest hour of heaviest day during the rollout of the first 8,000 outlets (1.25M) will take approximately 20 minutes to harvest when one site is in operation

See also §3.7.2.1 - OBCS Host Processing.

3.6.2.4 BPS Payment Authorisation Bulk Loader

Test 4.5 was combined with test 3.5 – see test 3.5 (§3.5.2.5).

See also BPS Host Processing (§3.7.2.3).

3.6.2.5 Reference Data Bulk Loader

Test 4.6 was performed by the Performance Test team. The results are reported in [REP014].

The Reference Data Bulk Loader loads changes to the reference data from the Host to the Correspondence Server. The volumes of data will be very variable and difficult to predict. The pattern of changes cannot be modelled. The pattern and volume of changes will be monitored during national rollout and the data collected will be used to re-work the model predictions and, if necessary, further testing will be carried out.

FUJ00078768

The Reference Data Bulk Loader requirement has been measured using the heaviest known workload during the rollout of the first 8,000 outlets, i.e. the loading of reference data into 100 newly automated outlets. This will generate more than 1.6M *core* and over 66K *non-core* items of reference data.

One copy of the *core* (data common to all outlets) data is loaded into a *virtual* outlet [4] on the Correspondence Server using the bulk loader agent. This set of data is then copied from the virtual outlet to each new outlet using a Replicator agent. The agent reads the 'Active Outlets Table' in the Host database to determine the outlets to which reference data is to be replicated. This agent runs a number of copies in parallel. The *non-core* or outlet-specific data is loaded from the Host directly into the outlet on the Correspondence Server.

The Reference Data loader model [MOD023] predicts that to load the reference data into 100 newly automated outlets on the Correspondence Server will take approximately 3 hours with only one site in operation. To download the reference data to the outlets will take approximately 45 minutes. The two processes overlap as the reference data is downloaded to an outlet soon after the new reference data is written into the message run for that outlet. This process will complete well before the 03:00 cut-off in the NR2 schedule, even with only one site in operation.

3.6.2.6 APS Bulk Harvester

Test 4.7 was performed by the Performance Test team. The results are reported in [REP039].

The APS Bulk Harvester transfers messages associated with automated payments from the Correspondence Server to an ORACLE database on the Host system. In the NR2 Maestro schedule, the APS bulk harvester runs after the TPS bulk harvester has completed.

The volumes of automated payments defined in the Workload Brief [BRIEF] amount to only 68 million per year, whilst current observed levels amount to about 150 million per year. The test was run using volumes approaching 200 million.

The APS bulk harvester test demonstrated that the agents can comfortably support projected APS harvesting loads up to full rollout volumes. Clearly if the AP volumes continue to rise then this position will need to be monitored.

[4] A *virtual* outlet is an outlet created for use by Riposte only.

FUJ00078768

At NR2 the AP loader model [MOD017] predicts that the number of messages that will be generated on heaviest day during the rollout of the first 8,000 outlets is approximately 380K. These will take less than 30 minutes to harvest when one site is in operation.

The tests indicated that beyond 8,000 outlets the elapsed time to harvest AP messages could be reduced by increasing the number of APS bulk harvester agent processes

3.6.2.7 Cluster Lookup Agent

Test 4.8 was combined with test 3.3. Test 4.8 was performed by the Performance Test team. The results are reported in [REP010].

The purpose of the Cluster Lookup Service is to enable the mapping of:

- Groups (i.e. Post Offices) onto Correspondence Server Clusters and
- Correspondence Servers onto NT machines

to be isolated from the agents.

Test 4.8 was designed to measure the time taken to re-establish the Cluster Look-up Agent following a failure. A test of the CLU Agent on a Correspondence Server with 5000 outlets demonstrated that the CLU Agent could re-establish itself in < 2 minutes. At NR2, the CLU Agent also should take < 2 minutes as there will be <2000 outlets per Correspondence Server. As CLU Agents run on all Correspondence Servers and Agents the loss on one agent will be minimal.

3.6.2.8 OBCS Stops Bulk Loader

Test 4.9 was performed by the Performance Test team. The results of OBCS Stops Bulk Loader testing are reported in [REP016].

The OBCS Stops bulk loader loads Stops from OBCS ORACLE database on the Host into the Correspondence Server. This process takes place overnight. There is also an interactive agent to handle ‘emergency stops’ which runs during the day. Most of the OBCS Stops will be processed by the bulk loader.

The OBCS Stops bulk loader Agents have been tested up to twice the volumes predicted for the heaviest day during the rollout of the first 8,000 outlets. The results indicate that the agents can comfortably support projected OBCS Stops harvesting loads up to the maximum volumes predicted for OBCS.

FUJ00078768

At NR2 volumes, the processing of OBCS Stops will take less than 5 minutes.

3.6.2.9 Change of Nominated Post Office Bulk Loader

Test 4.10 was performed by the Performance Test team. The results of the CNOMPOS are reported in [REP031].

CNOMPOS provides the capability of moving customers between outlets. Such changes can occur because:

- The customer requests the change at the Post Office
- The customer requests the change at a BA office or
- A Post Office has been closed (temporary or permanent).

CNOMPOS changes are loaded each night (Monday to Friday) by CNOMPOS bulk loader agents running on the Agent servers. All valid changes notified to Pathway by 20:00 should be available at the counter by 08:00 the following morning. However Pathway should notify CAPS of changes that have been made by 03:00. The schedule target for the completion of this task is therefore 03:00.

An engineering model produced to evaluate the NR2 design identified that the proposed NR2 design would be heavily I/O limited and would not be able to provide the turnaround required. The modelling work resulted in a re-evaluation of the design, and a revised design which minimised the amount of data moved between the 'old' and the 'new' post office was documented in CP1398. The first phase of the design change has been implemented at NR2. Beyond NR2 the design will be further refined to remove the requirement to move message data between outlets. This design will support the volumes predicted for full rollout.

The CNOMPOS model uses the test measurements of CNOMPOS that included moving customers between two outlets on the same Correspondence Server and between two outlets on different Correspondence Servers. Closure of an outlet was also tested.

On the heaviest day during the rollout of the first 8,000 outlets, the CNOMPOS model [MOD022] predicts that 37K CNOMPOS requests will be processed. The CNOMPOS model predicts that it will take less than 1 hour to process 37K CNOMPOS requests with single site operation. The NR2 Host processing model [MOD024] predicts that CNOMPOS will complete before 01:15 in single site operation.

FUJ00078768

3.7 HOST SERVERS

3.7.1 Summary

The Host platform has been extensively tested and modelled.

All of the testing focussed on demonstrating that a single 10-way Sequent SE70 server could support the Host workload at NR2 and at full rollout. At NR2 all of the Host workload (TPS, APS, OBCS and PAS/CMS) runs on one 10-way SE70, with the other SE70 providing the failover capability.

In order to ensure that the testing was as representative as possible a joint test team was formed from members of the:

- Pathway Performance Test team,
- ORACLE PAS/CMS design and development team and
- Sequent support team.

This testing team, known as the ORACLE Joint Working (OJW) team designed, set-up and ran a set of load tests on the Host system. This testing combined the stress testing of the Host applications by the ORACLE development team with performance testing.

The tests represented the key components of the Host overnight batch schedule i.e. TPS, PAS/CMS, OBCS and APS. The tests emulated the heaviest load that is predicted at three points in the rollout:

- 50% workload
- 100% workload on the peak day in an average week
- 100% workload on the peak day in an peak week

During the rollout the workload profile changes as cards replace books. Therefore a different message mix was used for the 50% and the 100% cases representing the profiles that are predicted by the workload volumes model [MOD026].

The 50% case ran OBCS, PAS/CMS, APS and TPS on a single 10-way SE70(see §3.7.2.5). The test demonstrated that the overnight schedule for the peak day with 8,000 outlets rolled out could be supported on a single SE70 server with spare capacity to cover operational contingencies. The results from this test confirmed the predictions from the engineering models that a single SE70 would not support full rollout volumes.

FUJ00078768

The 100% average week case ran PAS/CMS and APS on an SE70 to prove that these applications would scale to support the workload generated by 20,000 outlets (see §0). The test results demonstrated that a single SE70 could support this workload but with little spare capacity for operational contingency. All of the applications scaled well to the volumes generated by 20,000 outlets.

A 10-way SE70 could also support TPS volumes in excess of NR2 volumes (see §3.7.2.4) but the TPS model predicts that a single SE70 could not support the peak on peak volumes generated at full rollout.

As the previous two cases demonstrated, the SE70 reaches its limit beyond 8,000 outlets at peak, and Pathway will have to invoke the Host scaleability strategy before that limit is reached, in order to preserve the desired levels of contingency.

For the 100% peak week case the Host PAS/CMS batch workload was migrated to, and measured on, an early model in the NUMA-Q range (see §3.7.2.7). The database discs were disconnected from the SE70 and connected to the NUMA-Q without any changes. The transition of the application environment from SE70 to NUMA-Q was tested and this was achieved with minimal change. The test results demonstrate that the NUMA-Q platform will:

- provide sufficient power to support the full rollout volumes peak on peak and
- ensure that there is sufficient contingency in the overnight batch schedule to cope with operational incidents.

On the NUMA-Q platforms, further evaluation of key processes such as TPS Bulk harvesting is planned to ensure that the full power of the NUMA-Q systems can be exploited.

The disc sub-systems on the SE70 Host systems that support the ORACLE databases are high performance and availability Symmetrix systems from EMC^2. There is one Host Symmetrix system at each site. The Symmetric Remote Data Facility (SRDF) enables the databases to be mirrored across the sites with the minimum impact on performance. The Symmetrix systems were benchmarked [SRDF] for Release 1c and the tests showed that the impact on elapsed times was <5%.

FUJ00078768

3.7.2 Performance Assessment

Section 5 of [Requirements] specifies the agreed set of tests. These consist of:

8.1 OBCS Host Processing

8.2 BPS Host Processing

8.3 CAPS Payment Authorisation Data

8.4 TPS Host Processing

8.5 Concurrent Host Workloads (8,000 outlets)

8.6 Concurrent Host Workloads (20,000 outlets, average week)

8.7 Concurrent Host Workloads (20,000 outlets, peak week)

3.7.2.1 OBCS Host Processing

Test 5.1 was performed by the Performance Test team. The results of the OBCS bulk harvester test are reported in [REP031]. The results of processing the harvested messages on the Host are reported in [REP035].

The test measured the time for the Host to harvest and process a representative mix of OBCS Encashment & Administration messages. The test harvested 6M messages against the target for 8,000 outlets of approximately 1.1M (1.1M Encashments and 20k Administration) messages. The test was run at these volumes to ensure that the system could support greatly increased volumes if the rate of replacement of books by cards is different to that in the workload definition [VOLS]. The workload definition [MOD026] predicts that OBCS Encashments will peak at below 1.5M per day as cards replace books.

The OBCS Host processing model is divided into two parts:

- The OBCS Harvesting model (see [MOD018]) and
- The OBCS Host processing model (see [MOD024]).

The OBCS Host processing forms part of in the Host workload test (see §3.7.2.5).

FUJ00078768

The OBCS harvesting model is based on the results from the performance tests. For 8,000 outlets the model predicts that 1.1M OBCS (Encashment and Administration) messages will be harvested in approximately 20 minutes in single site operation. This is acceptable as OBCS harvesting is performed in parallel with TPS and APS Host processing and will complete before OBCS Host processing will be started by the Maestro scheduler.

3.7.2.2 BPS Host Processing

Test 5.2 was performed by the Performance Test team. The results of the BPS harvester test are reported in [REP032]. The results from processing the harvested messages on the Host are reported in [REP035].

The test measured the time for the Host to harvest the number of BPS Payment Authorisations (6.6M Encashments and 7.9M Payments) that would be generated on the peak day in an average week with 20,000 outlets rolled out.

The BPS Host processing model is divided into two parts:

- The BPS Harvesting model (see [MOD020]) and
- The BPS Host processing model (see [MOD024]).

The BPS Host processing forms part of the Host workload test (see §3.7.2.5).

The BPS harvesting model is based on the results from the performance tests. The test assumed that BPS harvesting was performed in a similar way to TPS and OBCS harvesting, i.e. a single harvesting operation at end-of-day. However, the BPS harvesting process runs in many batches throughout the day (07:30 to 20:00).

A BPS bulk harvester agent runs on each of the 16 Agent servers. Throughout the day, Maestro starts BPS bulk harvester on the Agent servers. The BPS bulk harvester agent harvests messages that are available to be harvested then closes. When a bulk harvester has closed, Maestro detects the closure and re-starts the bulk harvester. The cost associated with this additional processing has been estimated in the model and will be measured during a later phase of performance testing.

For 8,000 outlets the model predicts that BPS harvesting will take less than 10 elapsed minutes out of the available 1 hour to harvest the 460K messages generated in the peak hour.

FUJ00078768

3.7.2.3 CAPS Payment Authorisation Data

Test 5.3 was performed by the Performance Test team. The results of the Host bulk loader test are reported in [REP034]. The results from processing the payment authorisations on the Host are reported in [REP035].

Test 5.3 measured the cost on the Host of loading approximately 6.3M payments i.e. the number of payment authorisations generated on the heaviest day in a average week at full rollout.

The Payment Authorisation model [MOD020] predicts that on the heaviest day during the rollout of the first 8,000 outlets almost 900K messages will be loaded by the CAPS Payment Authorisation Data loader agent. This process takes place at the end of the overnight schedule and the data must be at the counter before the start of business i.e. 08:00. With only one site in operation the loading process will take less than 30 minutes and will finish before 05:00.

3.7.2.4 TPS Host Processing

Test 5.4 was performed by the Performance Test team. The results of the TPS bulk harvester test are reported in [REP031]. The results from processing the harvested messages on the Host are reported in [REP035].

TPS Host testing was undertaken using the workload generated by 10,000 outlets on a peak day in an average week. Testing of TPS on an SE70 did not exceed 10,000 outlets as early engineering models of the TPS workload indicated that a single SE70 platform could not support the full TPS Host workload up to full roll-out volumes. An early decision therefore was therefore required about where TPS would be hosted. Should it be:

- hosted on a more powerful Host than the SE70?
- run on a separate Host platform to PAS/CMS and APS?
- re-developed to run on a different platform?

The SE70 reaches its limit beyond 8,000 outlets at peak, and Pathway will have to invoke the Host scaleability strategy before that limit is reached to support the volumes predicted at full rollout peak on peak.

FUJ00078768

The TPS Host processing model is divided into two parts:

- The TPS Harvesting model (see [MOD019]) and
- The TPS Host processing model (see [MOD024]).

The TPS Host processing forms part of the Host workload test (see §3.7.2.5).

The window for harvesting TPS messages runs from 12:00 to 20:30 Monday to Saturday, but as TPS harvesting cannot be initiated until after the end-of-day marker has been received by the Host, most of the TPS harvesting will take place after 18:00.

The TPS harvesting model is based on the results from the performance tests. For 8,000 outlets the model predicts that on the SE70, harvesting will take approximately 1.5 hours with two sites in operation, but just over 3 hours with only one site working. This latter figure means that all of the EPOSS messages would not be harvested before the harvesters are closed at 20:30. This limitation is due to the number of instances of the TPS Bulk Harvester agent that are running i.e. 32 instead of the 64 with two-site operation. Pathway is reviewing the number of TPS Bulk Harvester agents and the power of agents required to support single site operation.

3.7.2.5 Concurrent Host Workloads (8000 outlets)

Test 5.5 was performed by the Performance Test team. The results of the SE70 50% workload test are reported in [REP035].

Test 5.5 measured a representative overnight batch cycle using 50% of the volumes generated on the heaviest day at full rollout. This test formed part of the ORACLE Joint Working (OJW) project that involved the ORACLE PAS/CMS design & development team, Sequent and the Pathway Performance Test team. The test revealed that for volumes representative of NR2, the Sequent SE70 Host could support the full overnight batch processing schedule with sufficient operational contingency for fails and re-runs.

The OJW team worked with the Pathway Maestro schedule team to validate and improve the overnight batch schedule for NR2 using the results from Test 5.5.

The workload model is based on the results from the performance tests together with the schedule start and finish points derived from the business model and the SLAs.

FUJ00078768

For 8,000 outlets the Host Processing model [MOD024] predicts that the overnight batch schedule will be completed by 05:00, even if only one site is in operation.

3.7.2.6 Concurrent Host Workloads (20000 outlets, average week)

Test 5.6 was performed by the Performance Test team. The results of Test 5.6 are reported in [REP036]. This test also formed part of the ORACLE Joint Working project.

Test 5.6 measured a representative overnight PAS/CMS and APS Host batch schedule using the volumes predicted for the heaviest day in an average week at full rollout. TPS is not included in this test as the test was designed to test PAS/CMS at high volumes. The test revealed that for volumes up to full rollout volumes (not peak on peak), the 10-way Sequent SE70 Host could execute the overnight batch processing schedule but there was minimal time left in the schedule for operational contingency.

Since this test was run, the projected CNOMPOS volumes have been revised upwards (see §3.6.2.9). With this additional load, the SE70 would not complete the overnight PAS/CMS schedule before the start of the next working day on the heaviest nights even in an average week. However, it is planned that the design of CNOMPOS will change at NR2^{+} to provide a performant solution for 20,000 outlets.

The SE70 can support more than 8,000 but less than 20,000 outlets running the overnight PAS/CMS and APS batch schedule. Pathway has therefore decided to implement its Host scalability strategy and replace the SE70 Host systems with Sequent NUMA-Q 2000 systems at or before 8,000 outlets have been rolled out.

A workload model for 20,000 outlets is planned.

FUJ00078768

3.7.2.7 Concurrent Host Workloads (20,000 outlets, peak week)

Test 5.7 was performed by the Performance Test team. The results of the NUMA-Q peak on peak workload test are reported in [REP036].

Test 5.7 measured a representative overnight PAS/CMS and APS batch schedule using peak on peak of full rollout volumes. The test revealed that the key components of the PAS/CMS overnight batch workload ran between 33 and 123% faster on the NUMA-Q system [5]. Pathway has therefore decided to implement the scalability strategy for the Host systems and replace the SE70 systems with NUMA-Q 2000 systems.

A workload model for 20,000 outlets is planned.

[5] The NUMA-Q system used in the tests was configured with 8-way 200MHz processors. The NUMA-Q system which will replace the SE70s is an 8-way 400MHz Intel Xeon system

FUJ00078768

3.8 DATA TRANSFERS TO/FROM HOST SERVERS

3.8.1 Summary

The interface between Pathway and its customers and suppliers involves the movement of large volumes of data both ways. Most of this data is transferred in bulk using the FTMS service or NFS. The interactive transfer of data to/from the Pathway system is reviewed in §3.9.

There are many services linking into Pathway but most of these services either support low volumes or are one of a set of similar services. A small subset of the key services was chosen for detailed testing and modelling. Each of these services carries large volumes of data and is subject to an SLA for the delivery of data.

The testing of the CAPS Batch interface has been performed jointly with ITSA and they have produced a report on the resources required to support the VME side of the interface. The report indicates that the measured load on the VME systems is less than had been expected.

The bulk data transfer services have been tested at volumes representative of those that will be achieved at full rollout which are considerably in excess of the volumes that will be generated on the busiest day during the rollout of the first 8,000 outlets. The test results were included in the NR2 workload model [MOD024], which indicates that all of the relevant SLAs can be achieved with sufficient headroom for contingency.

For NR2 the bulk data transfer services have been tested but not modelled. For full rollout this component of the workload will be modelled and the results will provide input into the workload model for 20,000 outlets. The OAS workload has not been measured. It is similar in profile to CAPS Batch but unlike CAPS Batch, the worst case i.e. the download of the national stops file, has already been completed on the live system.

3.8.2 Performance Assessment

Section 6 of [Requirements] specifies the agreed set of tests. These consist of:

8.1 CAPS File Transfer Interface

8.2 Host to TIP PC interface

8.3 TIP PC to PC file Transfer

8.4 MIS File Transfer

FUJ00078768

3.8.2.1 CAPS File Transfer Interface

Test 6.1 was performed as part of the CAPS Joint Interface project. The results of the tests are reported on in [REP037].

The Performance Test team and an ITSA team jointly produced the test plan and testing was run as a joint exercise. The ITSA team provided the VME components of the test and the message generator and the Performance Test team the Pathway PAS/CMS Host environment.

Files of data are sent to the CAPS Access Service (CAS) running on VME mainframes in the Benefits Agency. All files that are received by CAS up to 20:00 each evening are to be processed by Pathway as part of the overnight batch schedule. The transfer of data from CAS is scheduled to take place several times a day so only a small amount of data should be transferred to Pathway after 19:00. Each file could contain up to 40K transactions. In a peak on peak day 7.8M transactions will be transferred from CAPS to Pathway. Up to 8,000 outlets the volumes are less than 900K.

The projections derived from the measurements show that the cost of transferring data is lower than had been projected because only one process on the VME service is transaction based i.e. the costs are per transaction. The other VME services are file based.

The time to transfer a file with 40K payment authorisations was < 2minutes. The total elapsed time on the VME systems per day at full rollout has been estimated to be approximately 5 hours [CAPS], spread across 3 BA VME systems.

3.8.2.2 Host to TIP PC interface

Test 6.2 was performed by the Performance Test team. The results of the TIP transfer test are reported in [REP039]. Additional testing has been undertaken as part of TIP Interface Testing (see [VolStrat3]).

The TIP interface handles all of the EPOSS data being transferred from Pathway to the POCL TIP system. Up to 64 data files (plus a few small control files) are generated by the TPS Host processing suite and transferred to POCL. Each file contains the EPOSS messages for a set of outlets. On the heaviest day during the rollout of the first 8,000 outlets, approximately 12.5M messages are transferred to TIP. This increases to approximately 53M on a peak on peak day at full rollout.

Test 6.2 used a set of files generated by the TPS Host Processing test (see §3.7.2.4) to measure the time to transfer files of different sizes.

FUJ00078768

From the measurements it is estimated that the time to transfer the files to TIP will take less than 1.5 hours. However, most of the file transfer time is overlapped with the TPS processing which generates the files, as each file is available to be transferred to TIP as soon as processing of the file is complete. The files are required to be at TIP by 03:00, but the NR2 Host processing model [MOD024] predicts that they will normally be available before 23:00. This assumes that the TIP links are reconfigured from two 1Mbit links to one 2Mb link as the FTMS service cannot use the two 1Mb links concurrently (CP1741).

The testing in [REP039] used a low latency WAN link that is not representative of the live system. However, the testing in [VolStrat3] used the live WAN links from Bootle to Hulthwaite and this testing gave an accurate measurement of the transfer rate that can be achieved across the network.

3.8.2.3 TIP PC to PC file Transfer

Test 6.3 was performed by the Performance Test team.

For the assessment of test 6.3 See §3.8.2.2

3.8.2.4 MIS File Transfer

Test 6.4 was performed by the Performance Test team. The results of the MIS File Transfer test are reported in [REP040].

In §3.8.2.2 the interface from Host TPS to POCL TIP was measured. Host TPS has a second major datafeed, to the Pathway Data Warehouse/MIS system.

From the measurements it is estimated that the time to transfer the files to MIS using FDDI/100Mb Ethernet will take less than 45 minutes. However, most of the file transfer time is overlapped with the TPS processing that generates the files, as each file is available to be transferred to TIP as soon as it is closed by TPS. The Data Warehouse schedule requires that the files should be available on the Data Warehouse by 00:30, but the NR2 Host processing model [MOD024] predicts that they will normally be available before 23:00.

The results of the tests showed that the most efficient way of running the MIS file transfer is to run one transfer at a time and not to try running multiple concurrent file transfers.

FUJ00078768

The order in which MIS files are generated by TPS is under review. At present the Data Warehouse Upload process has to wait until all of the files have been received from TPS. If the order in which the files are generated is changed then the Upload can start as soon as the first file is available on the Data Warehouse system reducing the elapsed time of overnight processing on the Data Warehouse.

FUJ00078768

3.9 INTERACTIVE SERVICES

3.9.1 Summary

The tests conducted for the Interactive Services indicate that the system is capable of supporting the expected volumes that will occur on the busiest day at full rollout. Most of the load from the interactive services occurs during the 'on-line day' when the Host system is comparatively lightly loaded. The additional load generated by the interactive services on the Host in comparatively low.

It was therefore decided that the three key areas (see below) should be measured but not modelled.

3.9.2 Performance Assessment

Section 7 of [Requirements] specifies the agreed set of tests. These consist of:

8.1 OBCS 'Foreign' Transactions

8.2 PAS/CMS Helpdesk

8.3 CAPS On-Line

3.9.2.1 Test 7.1 – OBCS 'Foreign' Transactions

Test 7.1 was performed by the Performance Test team. The results of the tests are reported in [REP041].

The calculation of the number of OBCS 'foreign' transactions is dominated by the rate at which outlets are rolled out as the first time a book is presented at a newly automated outlet a 'foreign' transaction is initiated. This has the effect of 'registering' the book at that outlet and the outlet as the 'home' outlet for the NINO.

An engineering model produced to evaluate the NR2 design identified the need for greater throughput from the OBCS Agent in order to handle the heavy OBCS 'foreign' workload, and in particular the load generated by the automating of outlets. This model identified that the OBCS Foreign Agents were not capable of supporting the required peak OBCS 'foreign' throughput. To ensure that the system can deliver the throughput and responsiveness a revised design has been implemented at New Release 2 that introduces multi-threading OBCS foreign agents (see CP1118).

FUJ00078768

During the busiest week during the rollout of the first 8,000 outlets, on average 12 of OBCS 'foreigns' per second will be processed during the peak hour. The performance of OBCS 'foreigns' has been tested up to 18 per second per agent on one Correspondence Server. The response time for transactions entering the Correspondence Server through the campus 100Mb LAN was well below half a second. However, due to a limitation in the testing methods it was not possible to measure the response time at the Gateway PC i.e. including ISDN connection, catch-up and ISDN transmission times.

The 20-Counter evaluation test has measured the end-to-end cost of OBCS 'foreign' transactions on 20 counter Post Office running the peak on peak workload i.e. Christmas 2000. The mean response time for OBCS 'foreign' transactions (Scan OBCS Bar Code) at a counter is 3.3 seconds, including the gateway connection time and the ISDN delay.

The SLA states that the time taken for OBCS foreign transactions, (normal and fallback times) must not exceed an average of 28.1 secs, measured monthly over all outlets. The response time measured above is part of the overall figure.

The performance of BES 'foreign' transactions will be measured during a later phase of Performance Testing. However, the 20-Counter evaluation test has measured the mean response time for BES 'foreign' transactions (Swipe BES Magnetic Card) at 3.1 seconds.

This test also forms part of Test 3.7 (see §3.5.2.7).

3.9.2.2 Test 7.2 – PAS/CMS Help Desk

Test 7.2 was run as a joint test between the Performance Test team and the ORACLE PAS/CMS development team as part of the testing for Release 1c. The results of the tests are reported on in [REP035].

The Help Desk response times (at the 95th Percentile) were all sub-second for 400 users, using 15% of the Host CPU resources.

A Release 1c Help Desk client has a smaller foot print compared to a Release NR2 client, but on the Host Server side the same type of package functions/procedures are being used with added code. Between Release 1c and Release NR2 the Oracle core architecture (e.g. SQL*NET) did not change other than upgraded/patches, this means that the same type of messages are being passed between the client and server. The measurements taken at 1c are representative of the performance expected at

FUJ00078768

NR2.

The performance of the PAS/CMS Help Desk has been measured but not modelled.

3.9.2.3 CAPS On-Line

CAPS On-Line testing was performed as part of the CAPS Joint Interface project. The results of the tests are reported on in [REP041].

The Performance Test team and an ITSA team lead by John Stokoe jointly produced the testing plan and the test was constructed as a joint exercise. The ITSA team provided the VME components of the test and the message generator and the Performance Test team the Pathway PAS/CMS Host environment.

The maximum throughput at full rollout volumes is not expected to exceed 2 TP transactions per second. However to ensure that the end-to-end system could support higher throughput, if required, testing was performed up to 10 messages per second.

Up to 2 TP transactions per second, 99% of transactions had response times that were better than 3 seconds. The SLA states that > or = 95% of all on-line transactions have a < or = 3 seconds response time.

The 10 messages per second test was a functional test and not a performance test, so there was no response time target.

The performance of CAPS On-Line has been measured but not modelled.

FUJ00078768

3.10 SYSTEM FUNCTIONS

3.10.1 Summary

This set of tests was run to measure the performance of system functions so that sufficient time can be allocated in the operational schedule to perform the functions. There are no SLAs for the system functions.

3.10.2 Performance Assessment

Section 8 of [Requirements] specifies the agreed set of tests. These consist of:

8.1 PAS/CMS Backup and Restore

8.2 Correspondence Server Message Auditing

8.3 Correspondence Server Message Store Backup and Recovery

3.10.2.1 PAS/CMS Backup and Restore

Test 8.1 is being performed by the Performance Test team and ICL Outsourcing. The interim results of the tests are reported in [REP048].

The time to perform the daily 'hot backup' from the EMC^2 discs to the Sequent Host discs is approximately 75 minutes.

The 'hot backup' is performed at the end of the overnight schedule. On the heaviest night during the rollout of the first 8,000 outlets, the 'hot backup' will finish before 06:00.

3.10.2.2 Correspondence Server Message Auditing

Test 8.2 was performed by the Performance Test team. The results of the tests are reported in [REP043].

At NR2, all messages written to the Riposte Message Store are captured by the Audit Agent and written to an intermediate file on the Correspondence Server. This process is called 'harvesting the audit tracks'. This process takes place concurrently with other Riposte processes. Each intermediate file will hold 100,000 messages. When the limit on the number of messages has been reached a new file is opened and the previous file can be transferred to tape on the Audit Server. This process is called 'gathering the audit track'.

FUJ00078768

Test 8.2 was designed to collect data about the cost of running the harvester and gatherer whilst new messages are being written to the Message Store.

The test was run using the peak message volumes that will be generated during the peak hour at full rollout. This workload is more than three times the workload that is predicted from 8,000 outlets.

The measurements demonstrate that the Correspondence Server can support Correspondence Server Message Archiving running concurrently with the message replication loads that are predicted for the heaviest day during the rollout of the first 8,000 outlets.

This testing is being extended in Test 3.7 (see §3.5.2.7).

Correspondence Server Message Archiving has been tested at NR2 and the results used in the models.

3.10.2.3 Correspondence Server Message Store Backup and Recovery

Test 8.3 is being performed by the Performance Test team and ICL Outsourcing . The interim results of the tests are reported in [REP049].

The time taken to copy the message store from the EMC discs to the BCV in the EMC cabinet is approximately 75 minutes.

FUJ00078768

3.11 ISDN NETWORK

3.11.1 Summary

The Pathway ISDN network was modelled early in the design process. At the time the model showed that the network could support the full rollout volumes predicted. The latest engineering model predicts that the current ISDN network will support 4,000 outlets [6] with sufficient headroom for contingency.

What has changed? During this time the understanding of the network traffic generated by the workload has increased and this has had the effect of significantly increasing the number of datagrams the network has to support at peak times. At the same time the requirement to store network addresses in the ISDN Routers has increased dramatically in order to support the Tivoli systems management infrastructure and the performance of the ISDN routers in handling connections is a function of the size of the network address tables.

Pathway has put in place a major development plan to ensure that the ISDN network can deliver well in excess of the demand at any stage in the rollout programme. The first phase of the development is to increase the power of the ISDN routers by upgrading the processing unit. These developments will increase the safe limit of outlets that can be connected to 10,000 and will be intercepted by the end of live trial.

The ISDN network will support 20,000 outlets with the VPN developments at NR2+.

3.11.2 Performance Assessment

No specific performance measurement of the ISDN network has been done except to support the modelling.

It is intended that during live trial, the live trial Post Offices are used in a “test mode” to stress the ISDN network. This will allow the ISDN network to be loaded to a level higher than that expected at 20,000 outlets.

[6] The number of outlets that can be supported and which is stated in this section conforms to the headroom model described in the Scaleability Report [Scaleability]. It is not the maximum number of outlets that could be supported.

FUJ00078768

The proposal is to run a test of the ISDN links using the Live Trial system using a purpose written application that sits on the Gateway PC of each outlet in live trial. This application would be distributed to each of the Live Trial outlets prior to the test using the standard Pathway software distribution processes & procedures. The test would be initiated automatically. At a given time (e.g. on a Sunday) all 300 outlets phone the data centre.

As the clocks on the Gateways are synchronised, all 300 phone calls should arrive at the data centres together. The impact on e.g. the network routers and the load on the systems components can be measured. The call rate that would be generated by this test is higher than is expected from 20K outlets.

Extensive modelling using engineering models has been done of the ISDN network. This shows that the current ISDN routers can safely support 4,000 outlets. An upgrade to the router processors allows the current routers to support the rollout up to 10,000 outlets. This change has been committed in CP1807 for interception during live trial.

Pathway is investing in the development of the ISDN engineering models over the coming months. A sizing and modelling specialist, John Sewart, has joined the Pathway Performance Team to undertake this work.

The adoption of VPN in New Release 2^{+} will allow the current ISDN network to support 20,000 outlets. The programme of network modelling work will be extended to include VPN.

FUJ00078768

3.12 DATA CENTRE NETWORK

3.12.1 Summary

Early design work showed that the data centre network should not be a performance issue and that the busiest LAN segment [7] would be the one into the Correspondence Servers during the peak hour of the day.

Performance testing has shown that at 8,000 outlets this segment will be only 5% utilised.

3.12.2 Performance Assessment

There were no specific performance tests of the data centre network. Instead all performance tests report the network utilisation as part of the standard report.

The key report is [REP043] that tested Counter to Correspondence Server replication. During this test the maximum network utilisation was 17% at a message replication rate of 3.5 times that expected at 8,000 outlets. The test only had a single LAN connection to the Correspondence Server (simulating a failure on one of the LANs) and included the network load due to the Audit Server.

The data centre network has been tested at NR2 and the results used in the models.

[7] The site LAN is not a shared LAN. Each component is connected to a Gigabit switch across a point-to-point 100Mbit LAN. This architecture avoids collisions and permits each segment to operate at higher utilisations that a shared LAN.

FUJ00078768

3.13 WAN LINKS

3.13.1 Summary

The inter-site WAN link uses two 34Mbit E3 point-to-point connections. A separate 34Mbps E3 link is used for the Symmetrix disc mirroring system (SRDF) on the SE70 Hosts.

Performance testing has shown that at 8,000 outlets the inter-site WAN link will be 30% utilised [8]. This utilisation assumes that one of the two WAN links has failed.

3.13.2 Performance Assessment

There were no specific performance tests of the WAN link. Instead all performance tests report the network utilisation as part of the standard report. This allows the WAN link utilisation to be calculated.

The key report is [REP043] that tested Counter to Correspondence Server replication. During this test the maximum network utilisation on the LAN due to message replication was 7.5% at a replication rate of 5 times that expected at 8,000 outlets. This gives an expected network utilisation at 8,000 outlets of 1.5%.

The WAN link has 20% of the bandwidth of a LAN segment and has to handle the message replication traffic of 4 correspondence server clusters. This gives an expected utilisation of 30%. This assumes that only one of the WAN links is working.

The inter-site WAN network has been tested at NR2 and the results used in the models.

This gives an expected network utilisation at 8,000 outlets of 5%. For modelling purposes default scalebility plan will be invoked before the utilisation reaches 20% unless the CSP refines the default position.

[8] For point-to-point networks, utilisations up to 60% can be sustained.

FUJ00078768

3.14 DATA WAREHOUSE & MIS SYSTEMS

3.14.1 Summary

The performance of the Pathway Data Warehouse was tested in a joint project between the Data Warehouse architecture & design teams and OSD. This extensive project measured a full cycle of the Data Warehouse including end-of-week and end-of-month processing against the heaviest volumes which will be generated by 10,000 outlets. The results of the performance tests are recorded in Data Warehouse Technical Test Report [DW014].

Testing of the Data Warehouse is limited by the disc configuration on the Data Warehouse. The current disc configuration will support 55% of the predicted data volumes. Pathway has plans in place to upgrade the EMC^2 disc configuration on the Data Warehouse when required.

The performance tests demonstrated that the NR2 data volumes can be processed in the target window i.e. 00:30 to 05:30.

3.14.2 Performance Assessment

Testing of the Data Warehouse was divided into the four main phases:

Phase	Range of pre-populated data	Comments/Reason
Phase 1 - Upload	0 days	No prior data is required since the upload phase takes place in daily partitioned tables.
Phase 2 - Daily	0 days	No prior data is required since the daily consolidation phase predominately consolidates data into daily and weekly partitioned tables
Phase 3 - Weekly	6 days in current week	The weekly consolidation rolls up data from the current week; thus 6 previous days in the current week are required to give a representative time.
Phase 3 - Monthly	6 weeks of data in which the current month falls	The monthly consolidations select from up to 6 weekly partitions.
Phase 4 - casp01	6 weeks of data in which the current month falls	The casp01 process selects from up to 4 weekly partitions.
Phase 4 - vapp, capp, slam_int	0 days	The majority of the conformance processing operates on daily partitioned tables.
Phase 4 - Slam_upload	8 days after a full month	The SLAM application processes some data 8 days after a month end.

FUJ00078768

Performance models have been constructed for each of the 4 phases, and are summarised as follows:

- Phase 1 can be assumed to be near-linear (slightly worse) with an increase of 26 minutes for every 5,000 automated post offices that are rolled out. This allows us to predict the $NR2^+$ performance time to be 77+ 2 x 26 = 129 minutes (where 77 minutes is the time for the first 10,000 post offices rolled out).

- Phase 2 is a linear model showing an increase of 4 minute in performance time for every 5,000 automated post offices rolled out. Thus the $NR2^+$ times can be computed as 22 + 2 x 4 = 30 minutes (where 22 minutes is the time for the first 10,000 post offices rolled out).

- The normal daily operations in Phase 3 are constant over time, and independent of the number of post offices rolled-out.

- The end-of-week part of Phase 3 is well modelled by a linear relationship with an increase of 13 minutes per rollout of 5,000 automated post offices. The prediction for $NR2^+$ (20,000 post offices) performance time is 38+ 2 x 13 = 64 minutes.

- The end-of-month part of Phase 3 can be approximated by kN^x model (x = 1.44 & k = 0.000099). Using this model the performance time for 20,000 post offices is 154 minutes. The end-of-month at full-load is expected to be worse than the measured figure by a factor of approximately 2.5.

- Phase 4 is defined by a linear function of the number of post offices rolled-out. It is estimated that the performance time will increase by 151 minutes for every 5,000 automated post offices rolled out.

FUJ00078768

3.15 ROLLOUT BEAT RATE

3.15.1 Summary

The rollout model assumes that for Post Office rollout:

- National Rollout starts on 5th April 1999.
- The rollout rate in the first 6 weeks of National Rollout is 100, 100, 150, 150, 200 and 250 Post Offices respectively.
- The rollout is then 300 Post Offices per week until 22nd November 1999.
- Over the 7 week Christmas period the rollout is 150, 150, 150, 0, 0, 0, and 150 Post Offices respectively.

The rollout is then 300 Post Offices per week all Post Offices are rolled out.

All of the NR2 models use the workload Volume Model defined in [MOD026].

3.15.2 Performance Assessment

The rollout rate has a direct impact on many aspects of the system, including:

- The peak number of OBCS 'foreigns'
- The number of objects processed by the Reference Data Bulk Loader

as well as the more obvious measures like the number of calls across the ISDN network.

The workload model uses the expected case defined above to generate the workload volumes to ensure that the testing and modelling always uses the heaviest load that could be generated at any point during the rollout programme.

FUJ00078768

3.16 VECTOR SERVERS

3.16.1 Summary

Performance testing of the Vector Servers was undertaken at Release 1c and the results of these tests showed that the Vector Servers can support the workload up to full rollout.

Each BPS payment authorisation requires the Vector Server to issue a security vector or digital signature. The Vector Server has 10 million previously generated vectors.

3.16.2 Performance Assessment

Testing of the Vector servers formed part of the BPS Payment Authorisation bulk loader test (see §3.5.2.5). This test has demonstrated that 8 Vector Servers can generate the keys required by the BPS Payment Authorisation Bulk Loader agent on a peak on peak day in less than 6 hours. With both sites in operation this figure would be halved.

During the generation of vectors the CPU utilisation of the Vector server is 100%. However, this can run concurrently with the supply of vectors as this process uses very little Vector server resources and the supply is not effected by the vector generation process.

FUJ00078768

3.17 AUDIT SERVER

3.17.1 Summary

The performance testing and modelling of the Audit Server has focussed on the transfer of message data from the Correspondence Server to the Audit Server for long-term archiving. The messages are captured ('harvested') by an Audit Agent running on the Correspondence Server and stored in files on the Correspondence Server. The files of information are transferred ('gathered') by the Audit Server from the Correspondence Server using NFS.

3.17.2 Performance Assessment

Test 8.2 was performed by the Performance Test team. The results of the test are reported in [REP043].

The first part of the test was to re-run the counter to Correspondence Server message replication test previously reported on as part of Test 3.1 and 3.4 (see §3.5.2.4 for further details) using an enhanced version of the counter simulator to provide the most accurate representation of the load generated on the Correspondence Server by replication.

During the peak on peak hour at full rollout volumes the configuration used in the tests [9] could not support the replication from the counters and the Audit Agent operating concurrently. The maximum rate at which messages were replicated was reduced by 17% when the Audit Agent was run concurrently with replication.

The current configuration can support the heaviest workload generated during the rollout to 8,000 outlets. The load generated by NR2 has been tested in test 3.7 (see §3.5.2.7) when the Audit Agent was run concurrently with Riposte replication. At these volumes, no impact on the rate of replication was observed.

An analysis of the results from the peak on peak case indicates the problem is entirely due to the load on the CPU as the Correspondence Server is almost 100% utilised by replication only. Pathway has already taken steps to enhance the configuration of the Correspondence Server processor in-line with the Scaleability Strategy for Correspondence Servers [Scaleability]. and the I/O subsystem The 4-way COMPAQ servers with Intel Xeon processors will be included in future performance test plans.

[9] The Correspondence Server was a 4-way 200MHz Pentium Pro system with COMPAQ SCSI Disc Arrays.

FUJ00078768

During the development of the Audit Server architecture, a number of engineering models were constructed to assist with the sizing of:

- intermediate disc storage required on the Correspondence Server before the files are transferred to the Audit Server
- the communications link between the Correspondence Server and the Audit Server
- the tape sub-system on the Audit Server.

The output from these models was incorporated in the Audit Server architecture. The models do not form part of the formal model set as the performance testing described above is sufficient to provide Pathway with data to confirm that the interface is performant.

FUJ00078768

3.18 VIDEO BENCHMARKING

3.18.1 Summary

A benchmark of counter transactions took place to monitor service performance against contracted POCL service levels. The benchmark used video techniques to monitor and analyse the responsiveness of the system.

The results of the benchmark are documented in [VID011], [VID012], [VID013], and [VID014].

A subset of the scripts used in the benchmark has been included in the 20-Counter tests [20 Counter].

FUJ00078768

END OF DOCUMENT