



Major Incident Report – Correspondence Servers 9th
May 2005

Ref: CS/REP/216

Version: 1.0

Commercial In Confidence

Date: 09/01/2006

Document Title: Major Incident Report - Correspondence Servers online service 9th May 2005.

Document Type: Report

Release: N/A

Abstract: Report covering management, investigation, resolution, root cause analysis and corrective actions of Major Incident
PM0006760 / P10001523

Document Status: APPROVED

Originator & Dept: Mike Stewart - POA Customer Services

Contributors: *Peter Thompson – POA Customer Services*
Steve Parker – POA SSC
Tony Wicks – POA Customer Services

Internal Distribution: Carl Marx, Peter Thompson, Tim Roper, Peter Jeram, Colin Lenton-Smith, Mike Stewart, Tony Wicks, Peter Burden, Kevin Spence, Kathryn Simpson, Stacey Sutton, Liam Foley, Dave Baldwin, Richard Brunskill plus Contributors.

External Distribution: Dave Hulbert, Gary Blackburn, Beverley Dunn, Dave Smith, Clive Reed.

Approval Authorities: *(See PA/PRO/010 for Approval roles)*

Name	Role	Signature	Date
Dave Baldwin	FS POA Service Director		
Carl Marx	FS CS Service Management Manager		



0.0 Document Control

0.1 Document History

Version No.	Date	Reason for Issue	Associated CP/PEAK/PPRR Reference
0.1	09/05/2005	Initial Draft	
0.2	26/05/2005	Updated with comments and specific tables by Nikki Hawkins	
1.0	09/01/2006	Approval	

0.2 Review Details

Review Comments by:	23-FEB-2006
Review Comments to:	Mike Stewart

<i>Mandatory Review</i>	
FS CS Director	Dave Baldwin *
FS CS Service Management Manager	Carl Marx *
<i>Optional Review</i>	
FS Core Services Network Operations Manager	Dave Jackson
FS Core Services Site Support Manager	Paul Sandison
FS CS Services NT/Unix Support Manager	Andrew Gibson

(*) = Reviewers that returned comments

0.3 Associated Documents

Reference	Version	Date	Title	Source
PA/TEM/001	8.0	19/12/02	Fujitsu Services Document Template	PVCS



--	--	--	--	--

Unless a specific version is referred to above, reference should be made to the current approved versions of the documents.

0.4 Abbreviations/Definitions

Abbreviation	Definition
ADSL	Asymmetric Digital Subscriber Line
BAS	Broadband Access Server
BCM	Business Continuity Manager
CFM3	Network Support Team POA
DM	Duty Manager
FAD	Financial Accounting Division
HSH	Horizon System Helpdesk
IC	Inter Connect
LST	Live System Testing
PM	Problem Manager
PMDB	Problem Management Database
PO	Post Office
POA	Post Office Account
POL	Post Office Limited
RCA	Root Cause Analysis
SCT	Service Continuity Team
SMC	Systems Management Centre

0.5 Changes in this Version

Version	Changes
1.0	Approval



**Major Incident Report – Correspondence Servers 9th
May 2005**

Ref: CS/REP/216

Version: 1.0

Commercial In Confidence

Date: 09/01/2006

0.6 Changes Expected

Changes
Following reviews. None



0.7 Table of Contents

1.0	INTRODUCTION.....	6
2.0	SCOPE	6
3.0	MANAGEMENT SUMMARY	7
4.0	DESCRIPTION OF THE FAULT AND SERVICE FAILURE.....	8
5.0	SYMPTOMS AND BUSINESS IMPACT	10
6.0	INCIDENT MANAGEMENT	14
7.0	PROBLEM MANAGEMENT	17
8.0	CORRECTIVE ACTIONS	17



1.0 Introduction

This document details the initial incident that occurred within the live estate between 09:00 and 10:00 for a four minute period on Monday 9th May 2005 and the activities that were carried out for the remainder of the business day by the appropriate POA Service Management and support teams.

The incident that occurred within the live estate was a result of work activity at that time being carried out by support staff on the Correspondence Servers, which impacted online services with high transaction Reversals and Time Outs.

Service request E-0505090917 was raised at the time in relation to this incident on the HSH Powerhelp Helpdesk system.

2.0 Scope

This report covers the following areas in relation to the incident:

- Summary of Incident
- Service and Business Impact
- Recovery Timeline
- Root Cause
- Action Plan / Corrective Actions



3.0 Management Summary

The incident within the live estate occurred at the time that support staff were analysing the Correspondence Server volume capacity using the Riposte Volume Command.

At this present moment the root cause of the incident isn't known, as further investigations continue by the support teams.

The purpose of this activity on the Correspondence Servers was to ensure that there available spare disk capacity across the Correspondence Server disk volumes for the Riposte message stores.

The work activity was carried out as a precautionary measure as a result of an archiving job not completing following an event storm that occurred the previous week.

The Riposte Volume command has been used on numerous occasions before with no impact to service, as it simply displays the details of Riposte Volumes.

Running this command stops Riposte services running for a micro-second, and then starts (unlocks) the services again after the volume is taken. It is believed that in this instance, due to a bug in Riposte, the services did not start again. A peak (PC0120645) has been raised with Escher development.

This caused the first degradation to service for 4 minutes between 9 and 10am. An event storm with messages from the Correspondence Servers stated that services were locked. As timeouts reduced to normal levels, the support teams diagnosed that the event storms and locks had cleared.

At 1030am SSC started to be aware through events being generated that two of the clusters weren't harvesting messages from the counters.

Following further subsequent analysis by SSC support staff, in order to was made to restart the Riposte services on the impacted servers within clusters 1 & 2 respective following discussions between POL and POA senior management.

This decision was justified by a high number of calls being received into the service desk and by the increase in the abandoned rate (section 5 Table B) at the desk between 12:00 – 13:00.

The Riposte services were restarted at 13:05 on those servers in clusters 1& 2



As SSC continued to monitor the live estate it was clear that as the correspondence servers were in the process of synchronising the Riposte message stores within the clusters that this was having a performance impact on the live estate and timeouts were continuing to occur at the counters.

SSC formulated a maintenance plan to resolve the impacts on the live estate which was agreed to be implemented that evening from 18:45.

The work activity with regard to the maintenance plan concluded at 21:45 this involved restarting Riposte services on each of the clusters in order to ensure that the Correspondence servers were returned to normal operational status.

4.0 Description of the fault and service failure

The incident that occurred presently hasn't a determined root cause and occurred within the live estate whilst the SSC were analysing the Correspondence Server volume capacity using the RiposteVolume Command. The purpose of this activity was to ensure that there available spare disk capacity across the Correspondence Server disk volumes.

The work activity was carried out as a precautionary measure as a result of an archiving job not completing following an event storm that occurred the previous week.

The RiposteVolume command has been used on numerous occasions before with no impact to service, as it simply displays the details of Riposte Volumes.

Running this command stops Riposte services running for a micro-second, and then starts (unlocks) the services again after the volume in taken. It is believed that in this instance, due to a bug in Riposte, the services did not start again. A peak (PC0120645) has been raised with Escher development.

This caused the first degradation to service for 4 minutes between 9 and 10am. An event storm with messages from the Correspondence Servers stated that services were locked. As timeouts reduced to normal levels, the support teams diagnosed that the event storms and locks had cleared.

At 10:30am, the SSC balance monitor alerted that C12 responses for Clusters 1 and 2 were not being harvested. There were no other indications of service issues at this stage. Investigations isolated the fault to one server on each cluster which was processing 1% of

**Major Incident Report – Correspondence Servers 9th
May 2005****Ref: CS/REP/216****Version: 1.0****Commercial In Confidence****Date: 09/01/2006**

their normal transaction rate. NB if the servers had failed completely, then the counters would have automatically balanced across the remaining 3 servers. As the server was still running, this degraded the service to the counters connected.

From a customer perspective, the customers would have noticed timeouts but a subsequent transaction retry attempt would have been successful (section 5 table A2).

This Riposte fault has been seen on counters before but never at the data centre (PC0101472 refers - fix at S90).

As the calls into the Helpdesk rose through the morning, the number of calls waiting reached ~40 at 12:30, it was clear that the affected Correspondence Servers required reloading in order to restore these servers to normal operating speeds. I.e. the Riposte services were still locked.

A decision was made by Senior POL and POA Management following advice from the support teams to stop and restart the Riposte services on a server within cluster1, based on the number of transaction timeouts/events occurring via system management toolsets and the number of calls being received into the HSD.

At 13:05 one of the Bootle Correspondence servers had the Riposte services stopped and then restarted. Whilst this activity was taking place the numbers of calls waiting reduced to 0 at the HSD service desk. During that activity, the percentage of transaction timeouts was less than 2%

SSC staff continued throughout the remainder of the day to evaluate and monitor the live estate and to formulate an appropriate maintenance plan.

The maintenance plan was derived to perform the remedial activity to the Correspondence servers that required the Riposte services to be stopped and restarted outside of Core hours on specific servers within clusters 1, 2 & 3 respectively.

The activity plan and communication plan surrounding this activity was agreed by Service Management and support staff and approved through the OCP process (OCP11050).

This maintenance activity was required to return the Correspondence servers to normal operational mode and ensure that the overnight harvesting activity could take place.

As part of the OCP a decision was made to suspend harvesting until the remedial activity had taken place.



The maintenance activity started at 18:45 and concluded at 21:45 and at that point overnight harvesting was then initiated and subsequently completed successfully overnight via the Maestro Batch scheduler.

5.0 Symptoms and Business Impact

The following tables show the transaction profiles that have occurred for the four minute period for the previous four Monday's detailing success and failure rates.

Table A – Transaction table for business volumes

TIMED TOTALS		Transaction total	Successes	Fails	Fails %
LINK	11-Apr-05	2415	2411	4.00	0.17%
Alliance & Leicester plc	11-Apr-05	1,034	1,028	8	0.77%
card account at Post Office	11-Apr-05	35,990	35,956	34	0.09%
e-pay Ltd	11-Apr-05	1,376	1,376	0	0.00%
Streamline (Online Debit Card)	11-Apr-05	642	642	0	0.00%
LINK	18-Apr-05	3,990	3,988	2	0.05%
Alliance & Leicester plc	18-Apr-05	1,041	1,040	1	0.10%
card account at Post Office	18-Apr-05	34,558	34,531	27	0.08%
e-pay Ltd	18-Apr-05	1,516	1,516	0	0.00%
Streamline (Online Debit Card)	18-Apr-05	659	658	1	0.15%
LINK	25-Apr-05	2,548	2,543	5	0.20%
Alliance & Leicester plc	25-Apr-05	1,113	1,111	2	0.18%
e-pay Ltd	25-Apr-05	1,535	1,535	0	0.00%
card account at Post Office	25-Apr-05	37,006	36,974	32	0.09%
Streamline (Online Debit Card)	25-Apr-05	672	672	0	0.00%
LINK	03-May-05	4,466	4,458	8	0.18%
Alliance & Leicester plc	03-May-05	1,168	1,167	1	0.09%
card account at Post Office	03-May-05	32,747	32,718	29	0.09%
e-pay Ltd	03-May-05	1,475	1,473	2	0.14%



Major Incident Report – Correspondence Servers 9th
May 2005

Ref: CS/REP/216

Version: 1.0

Commercial In Confidence

Date: 09/01/2006

Streamline (Online Debit Card)	03-May-05	1,337	1,335	2	0.15%
LINK	09-May-05	2,890	1,039	1,851	64.05%
e-pay Ltd	09-May-05	1,088	561	527	48.44%
Alliance & Leicester plc	09-May-05	1,300	468	832	64.00%
card account at Post Office	09-May-05	39,853	14,333	25,520	64.04%
Streamline (Online Debit Card)	09-May-05	925	396	529	57.19%

Table A2

This table shows the failure rate for the five minute period after the initial service impact in comparison to the failures detailed in Table A.

Bank	Receipt Date	Transaction Volume	Successes	Fails	Fails %
Alliance & Leicester plc	09-May-05	1,210	1,204	6	0.50'
card account at Post Office	09-May-05	40,303	40,099	204	0.51'
e-pay Ltd	09-May-05	1,449	1,414	35	2.42'
Streamline (Online Debit Card)	09-May-05	691	673	18	2.60'
LINK	09-May-05	2,652	2,626	26	0.98'



Major Incident Report – Correspondence Servers 9th
May 2005

Ref: CS/REP/216

Version: 1.0

Commercial In Confidence

Date: 09/01/2006

Table B Help Desk Statistics:

HSD 30 Minute SLA Conformance																		
	ACD	Total	Totl	Avail	Aban	Abn	Abn	Aban	Aban	Aban	ACD	ACD	Ans	ACD	Ans	Ans	Ave	
	Call	Call	Call	Call	40	>40	> 40	20	20-	40-	> 60	Call	Infl	By	Call	30	< 30	
	Offer	Call	Aban	Ans	secs	secs	SLA	secs	secs	secs	Fwd	Out	NACD	Will	secs	SLA	Logged in	
							</=5%	secs	secs	secs				ReQ			>/=80%	
07:30	3	3	0	3	0	0	0.00%	0	0	0	0	0	0	0	0	3	100.00%	1
08:00	29	29	0	29	0	0	0.00%	0	0	0	0	0	0	0	0	29	100.00%	13
08:30	118	84	34	118	6	28	23.73%	1	5	7	21	0	0	0	0	57	48.31%	22
09:00	362	97	265	362	43	222	61.33%	16	27	45	177	0	0	0	0	13	3.53%	25
09:30	193	70	123	193	12	111	57.51%	5	7	16	95	0	0	0	0	14	7.25%	25
10:00	121	46	75	121	2	73	60.33%	1	1	8	65	0	0	0	0	4	3.31%	25
10:30	89	70	19	89	8	11	12.36%	3	5	2	9	0	0	0	0	31	34.83%	25
11:00	99	75	24	99	7	17	17.17%	3	4	1	16	0	0	0	0	22	22.22%	25
11:30	85	45	40	85	5	35	41.18%	1	4	4	31	0	0	0	0	6	7.06%	25
12:00	320	46	274	320	27	247	77.19%	10	17	18	229	0	0	0	0	4	1.25%	14
12:30	235	45	190	235	24	166	70.63%	12	12	8	158	0	0	0	0	2	0.85%	15
13:00	83	55	26	81	6	20	24.09%	1	5	4	16	0	0	0	2	30	37.34%	14
13:30	61	50	11	61	2	9	14.75%	1	1	1	8	0	0	0	0	26	42.62%	15
14:00	85	62	23	85	3	20	23.53%	0	3	5	15	0	0	0	0	20	23.53%	19
14:30	111	77	33	110	9	24	21.62%	5	4	3	21	0	0	0	1	15	13.51%	20
15:00	132	76	56	132	12	44	33.33%	3	9	6	38	0	0	0	0	13	9.85%	23
15:30	124	71	53	124	11	42	33.87%	7	4	7	35	0	0	0	0	12	9.68%	23
16:00	107	44	63	107	14	49	45.79%	8	6	11	38	0	0	0	0	7	6.54%	20
16:30	151	54	97	151	11	86	56.95%	4	7	9	77	0	0	0	0	3	1.99%	19
17:00	62	24	38	62	13	25	40.32%	12	1	2	23	0	0	0	0	15	24.19%	9
17:30	27	24	3	27	3	0	0.00%	3	0	0	0	0	0	0	0	24	88.89%	8
18:00	3	1	1	2	0	1	33.33%	0	0	0	1	1	0	0	0	1	33.33%	3
18:30	0	0	0	0	0	0	-	0	0	0	0	0	0	0	0	0	-	0
19:00	0	0	0	0	0	0	-	0	0	0	0	0	0	0	0	0	-	0
19:30	0	0	0	0	0	0	-	0	0	0	0	0	0	0	0	0	-	0
20:00	0	0	0	0	0	0	-	0	0	0	0	0	0	0	0	0	-	0
TOTAL	2600	1148	1448	2596	218	1230	47.30%	96	122	157	1073	1	0	0	3	351	13.50%	388



Table C Overall daily Stats NBX

09/05/2005	Transaction Volume	Successes	10/05/2005 08:38	FAILS
Deposit & other	1,031,243.00	993,774.00	96.37%	37,469
Cheque Deposit	2,119.00	1,874.00	88.44%	245
DCP	90,328.00	75,113.00	83.16%	15,215
ETU	77,155.00	67,858.00	87.95%	9,297
Withdrawal	775,413.00	747,616.00	96.42%	27,797
Combined	1,976,258.00	1,886,235.00	95.44%	90,023

E Top Up Transaction failures displayed in the table above occurred due to the interface between Fujitsu and EPAY. When authorised transactions timeout at the counter a reversal is automatically generated by the counter. Transactions were timing out due to the Correspondence server problems on this day. EPAY's systems have to match the reversal to the original authorisation request. The reversal has to get to EPAY within a time limit of 10 minutes.

E-pay advise that there is a high probability of a Reversal being successful if it arrives within 10 minutes of the original Sale Request. E-pay will attempt to reverse Transactions up to an hour after the original Sale Request, but advise that this is not likely to be successful. E-pay will not action a Reversal that arrives more than an hour after the original Sale Request.

Because of the problems with Correspondence server replication on this day a number of these reversals did not get through within the time limits.

Under normal circumstances, even if there has been a problem with a Correspondence server, the remaining correspondence servers in that cluster will ensure that the reversal is processed and handed off to the agents (of which there are also multiple instances) to pass to Epay.

In the case of the issues on 9th May there were multiple correspondence servers running slowly without failing and the reversals did not get to Epay within the time limits. (ie if the Correspondence server had failed in totality, the others in the cluster would have run normally and the timing issue would not have occurred). This is the first recorded instance of this sort of Riposte problem on multiple correspondence servers on the project.

The instance of this sort of issue is therefore very rare. In mitigation of the original issue SSC staff will no longer use the Riposte Volume Command during core hours until Escher have resolved the underlying bug.

One possible solution is to mitigate the financial exposure that Epay change their processing - however this would be difficult as a customer could be using the credit on the phone immediately upon leaving a Post Office branch.



For banking and debit/credit card transactions the process is that a counter requests a transaction, the financial institution accepts the transaction and the counter then confirms the transaction, which does not actually take place until the confirmation is received and can usually then be reversed up to 48 hours later. For an E-Topup transaction, Epay acts immediately on receipt of the request, issues an accept, and relies on the counter to reverse within 10 minutes.

Although it would be better for clients to operate the full financial transaction request/accept/confirm (RAC) model, and it may be possible to persuade Epay to change their process, in practice this may be very difficult.

Another option would be for Fujitsu, on early detection of a Correspondence server incident which could affect Epay transactions, to switch off the links to Epay to stop all transactions going through, however, this clearly would cause customer dissatisfaction and damage.

6.0 Incident Management

<i>Date & time</i>	<i>Action</i>	<i>Owner</i>	<i>Comments</i>
9/5/2005 09.25 hrs to 09.27	Contact from SSC to POA On-Line SDM. Information provided, stating SSC were seeing alerts. SSC unable to quantify the problem. SSC stated they would update further once they knew more.	Mike Woolgar	
09.28 to 09.32	Contact from HSH to POA On – Line SDM, stating that call E-GRO had been logged as an 'A' priority call, for On Line Banking issues and had been sent to SSC. They also expressed concern that POL SCT had been asking numerous questions to which they were unable to answer at this stage and went into detail. POA On Line SDM said that a call had already been received from SSC and they were already on the case. Similarly the questions and detail requested by POL SCT had slowed down this 'communication'. By taking longer than normal to discuss.	Mike Woolgar	
09.32 to 09.33	POA On Line SDM called Nikki Hawkins. A message was left giving as much detail as possible in		



Major Incident Report – Correspondence Servers 9th
May 2005

Ref: CS/REP/216

Version: 1.0

Commercial In Confidence

Date: 09/01/2006

	<i>relation to the incident</i>		
09.34	<i>POA On Line SDM returned a call outstanding to SMC(following a pager message) that referred to seeing alerts on all types of NBX transactions, and that SSC were investigating</i>		
09.35 to 09.39	<i>POA On-Line SDM called POL SCT. The line was engaged, but said if it was urgent please wait. POA On-Line SDM waited until cut of. This was repeated. Then POA On Line SDM contacted Beth Newton and informed her of situation and then finally managed to get hold of POL SCT at 3rd attempt.</i>		
09.40	<i>POA On Line SDM updated Carl Marx and Peter Thompson on way though to visit SSC who had a left a message that they had more details.</i>		
09.45 to 09.53	<i>SSC updated POA Online SDM that there were approximately 23,500 NBX transactions that had failed out of 62,000. They stated that this had been caused by them running a piece of software to 'manage space' on Correspondence Servers. They stated that this was supposed to be Non Intrusive, whereas it had in fact caused these problems</i>		
09.54	<i>POA On Line SDM updated Nikki .Hawkins</i>		
09.55	<i>POA On Line SDM updated POL SCT</i>		
9/5/05	<i>POA follow "across and up" communication to POL</i>	<i>POA Service Management</i>	<i>POA Service Management follow defined communication escalation process</i>
9/5/05 10:30	<i>SMC notified of high call rates to HSH and C12 responses not being harvested on cluster1 & 2</i>	<i>Pat Carroll/Steve Parker SSC</i>	<i>SSC continue to investigate</i>
9/5/05	<i>Discussion between POL and POA Senior Management takes place.</i>	<i>POA and POL management</i>	<i>Decision to stop and restart Riposte services</i>



Major Incident Report – Correspondence Servers 9th
May 2005

Ref: CS/REP/216

Version: 1.0

Commercial In Confidence

Date: 09/01/2006

12:30			on cluster1 in Bootle due to high call volumes to HSD and events generated by System Management Toolsets
9/5/05 13:05	MBCOR01 & MBCOR02 stop and restart Riposte services	Steve Parker – SSC	OCP11049
9/5/05 13:35	MBCOR01 & MBCOR02 returned to service	Steve Parker - SSC	Call volumes reduced to HSD and timeouts reduced.
9/5/05 16:30	Large delta marks observed on the Cluster Servers	Steve Parker- SSC	Riposte Message stores out of step within cluster
9/5/05 17:00	Maintenance Plan determined /OCP raised for activity	Maintenance to be carried out on Clusters1,2,3	Work to be carried out to ensure harvesting can be completed and message stores OCP11050
9/5/05 18:30	Communication Bridge Opened	POA Service Management/POA Support teams	
9/5/05 18:45	Maintenance Activity Plan commenced	Steve Parker SSC	Work commenced on clusters 1,2,3 ,stopping and starting Riposte Servers
9/5/05 21:30	Communication Bridge closed	POA Service Management / POA Support Teams	
9/5/05 21:45	Maintenance Activity Plan concludes	Steve Parker – SSC	All remedial activity to clusters 1,2,3 completed
9/5/05 21:50	Overnight Harvesting commenced	Andrew Gibson	Once the harvesting was initiated the Tip Files were produced by 00:30



7.0 Problem Management

Mike Woolgar and Mike Stewart were the On-Line Service Managers and Mike Stewart is the Allocated Problem Manager for this incident on the PMDB. A Cross Domain Problem has logged on the PMDB and the references are PM0006760 / P10001523

8.0 Corrective Actions

<i>Action to be taken</i>	<i>By Whom</i>	<i>By When</i>	<i>Status</i>	<i>Comments</i>
<i>Review of network configuration in relation to Correspondence Servers</i>	<i>Neil Preston</i>	<i>16/05/05</i>	<i>Completed</i>	<i>Analysis to determine if network layer and network performance contributed to the incident</i>
<i>Future use of the Riposte Volume Command to be governance by OCP</i>	<i>SSC</i>	<i>9/05/05</i>	<i>Completed</i>	<i>Governance</i>
<i>Progress Incident Call with Escher</i>	<i>SSC</i>	<i>9/05/05</i>	<i>Completed</i>	<i>The peak PC0120645 was raised and sent to Escher. Escher could not recreate the fault and neither could the SSC or Development in any of the test systems. As now the command is banned from being used in "Live" time operation, requires to be approved on OCP for Out of Operational Hours use and as we cannot test this in the Live Operation again, no further action will be taken on the Peak by either Escher or Development.</i>
<i>All minor changes to commands used or changes in the use of commands should be subject to OCP</i>	<i>All Support teams associated to the POA account</i>	<i>10/05/05</i>	<i>Completed</i>	<i>All changes are supported by OCP</i>