FUJ00124014

Exhibit GJJ/1

Document Title:	Horizon Data Integrity
Document Reference:	ARC/GEN/REP/0004
Document Type:	Report (REP)
Release:	N/A
Abstract:	This document describes the measures that are built into Horizon to ensure data integrity. Note that it only covers Horizon and not HNG-X (Horizon Online).
Document Status:	Final Draft
Author & Dept:	Gareth I Jenkins
External Distribution:	

Approval Authorities:

Name	Role	Signature	Date
Suzie Kirkham	Account Manager		

Note: See Post Office Account HNG-X Reviewers/Approvers Role Matrix (PGM/DCM/ION/0001) for guidance.

FUJ00124014

0 Document Control

0.1 Table of contents

FUJ00124014

0.2 Figures and Tables

0.2.1 Table of Tables

None.

0.3 Document History

Version No.	Date	Summary of Changes and Reason for Issue	Associated Change - CP/PEAK/PPRR Reference
0.1b	02/10/2010	First Informal Draft. Changes from version 0.1a were marked in red (like this) with ~~strikeout~~ for significant deletions.	
1.0	02/10/2009	Version for release to Post Office.	

0.4 Review Details

Review Comments by :	02/10/2009
Review Comments to :	Gareth Jenkins
Mandatory Review	
Role	**Name**
Suzie Kirkham	Account Manager
Jeremy Worrell	CTO
Optional Review	
Role	**Name**
Guy Wilkerson	Commercial Director
LaToya Smith	Commercial
Amanda Craib	Head of Commercial, Retail, Royal Mail and Telcos
David Smith	Post Office
Issued for Information – Please restrict this distribution list to a minimum	
Position/Role	**Name**

(*) = Reviewers that returned comments

(†) = Reviewers that returned no comments

0.5 Associated Documents (Internal & External)

Reference	Version	Date	Title	Source
PGM/DCM/TEM/0001 (DO NOT REMOVE)			Fujitsu Services Post Office Account HNG-X Document Template	Dimensions
ARC/GEN/REP/0001			HNG-X Glossary	Dimensions

FUJ00124014

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

0.6 Abbreviations

Abbreviation	Definition
AP	Automated Payments
CRC	Cyclic Redundancy Check

0.7 Glossary

See also document ARC/GEN/REP/0001.

Term	Definition
Replication	The mechanism by which data is reliably copied between the local system and other systems (i.e. other counters, external storage in a single counter branch and the data centre)

0.8 Changes Expected

Changes
Review comments etc.

0.9 Accuracy

Fujitsu Services endeavours to ensure that the information contained in this document is correct but, whilst every effort is made to ensure the accuracy of such information, it accepts no liability for any loss (however caused) sustained as a result of any error or omission in the same.

0.10 Copyright

FUJ00124014

1 Purpose

This document is submitted to Post Office for information purposes only and without prejudice. In the event that Post Office requires information in support of a legal case Fujitsu will issue a formal statement.

This document is a technical description of the measures that are built into Horizon to ensure data integrity, including a description of several failure scenarios, and descriptions as to how those measures apply in each case.

Note that this document only covers Horizon. It does not cover HNG-X (Horizon Online).

FUJ00124014

2 Horizon Data Integrity

The Horizon system is designed to store all data locally on the counter's hard disk. Once the data has been successfully stored there it is then replicated (copied) to the hard disks of any other counters in the branch (and in the case of a single counter branch to the additional external storage on the single counter). Data is also passed on from the gateway counter to the Horizon data centre using similar mechanisms.

The replication process is designed such that should the data fail to be copied immediately (for example due to a failure on the local IT network within the branch or another counter being switched off or the branch being disconnected from the data centre), then further attempts are made to replicate the data at regular intervals until it is finally copied successfully. Once the data reaches the Data Centre a further copy is taken and added into the audit trail where it is available for retrieval for up to 7 years. Data in the audit trail is "sealed" with a secure checksum that is held separately to ensure that it has not been tampered with or corrupted.

Every record that is written to the transaction log has a unique incrementing sequence number. This means it is possible to detect if any transitions records have been lost.

While a customer session is in progress, details of the transactions for that customer session are normally held in the computer's memory until the customer session (often known as the "stack") is settled. At that point all details of the transactions (including any methods of payment used) are written to the local hard disk and replicated (as described above). It should be noted that double entry bookkeeping is used when recording all financial transactions, ie for every sale of goods or services, there is a corresponding entry to cover the method of payment that has been used. When a "stack" is secured it is written in such a way that either all the data is written to the local hard disk or none of it is written. This concept of "atomic writes" is also taken into account when data is replicated to other systems (ie other counters, external storage or the data centre).

The data for a stack will have been successfully secured to the local hard disk before the screen is updated indicating that a new customer session can be started. Note that although an attempt will have been made to replicate the data to an external system at this time, there is no guarantee at this point that such replication will have been successful. For example if there is a Network Failure followed by a Terminal failure there is a slight risk that transactions in the intervening period could be lost.

All data that is written includes a "checksum" value (known as a CRC) which is checked whenever the data is read to ensure that it has not been corrupted. Any such corruptions detected on reading will result in failures being recorded in the event logs which are held on the local hard disk for a few days for immediate diagnosis and also immediately sent through to the data centre where they are held for 7 years.

Any failures to write to a hard disk (after appropriate retries) will result in the counter failing and needing to be restarted and so will be immediately visible to the user.

Whenever data is retrieved for audit enquiries a number of checks are carried out.

1. The audit files have not been tampered with (ie the Seals on the audit files are correct)
2. The individual transactions have their CRCs checked to ensure that they have not been corrupted.
3. A check is made that no records are missing. Each record generated by a counter has an incremental sequence number and a check is made that there are no gaps in the sequencing.

FUJ00124014

3 Scenarios

It should be noted that these scenarios are all to do with equipment failures and these will always be visible to Fujitsu through the event logs which are retained.

3.1 A counter fails

When a counter fails, there are two possible scenarios:

- It can be successfully restarted
- It cannot be successfully restarted, so needs to be physically replaced

In each case the Data Integrity considerations are different and so are described separately below.

Once the counter has been restarted (regardless of whether or not it has been replaced) recovery may be carried out if recoverable transactions are detected on the counter. This is also discussed below.

3.1.1 The Counter is Successfully Restarted

In this case all the data that had been secured prior to the failure is still present on the counter and so is available for use. If the User is in any doubt as to whether a transaction had been completed or not prior to the failure they can use the transaction logs to confirm one way or the other.

3.1.2 The Counter is Physically Replaced

In this case there is no data on the local hard disk of the replacement counter. However, since the data should have been replicated to other counters in the branch (or in the case of a single counter branch to the external storage – which should have been physically moved to the replacement counter), then the data should be retrieved and copied to the new counter. If for some reason the data were not available locally in the branch, then it will be copied back from the data centre. This all happens automatically as part of the counter replacement procedure.

Note that the hard disks are encrypted so there is no danger of data protection issues once the old counter has been removed (or if it is stolen).

When a counter is physically replaced, there is a possibility that not all data has been successfully replicated to another system prior to the failure. In this scenario it is essential that the user confirms what the last successful transaction on that counter was, again by using the transaction logs.

3.1.3 Transaction Recovery

Some classes of transaction generate recovery data as they go along, so as to ensure that in the event of a failure between the transaction starting and the basket being secured, there is sufficient information available to enable the transaction to be recovered. On Horizon there are two separate mechanisms to cover different classes of transaction:

- Banking Recovery
- AP Recovery

Both these mechanisms are automatically invoked during Log On, should the system detect that there has been a possible failure. These are described below.

FUJ00124014

3.1.3.1 Banking Recovery

This covers credit card and debit card transactions and e-Top-Up transactions as well as online banking transactions.

A check is carried out to see if any incomplete banking style transactions (i.e. network banking, credit / debit card or e-Top-Up) exist in the transaction logs for that counter. An incomplete transaction is one where an authorisation request has been sent to the financial institution, and there is no corresponding completion message which is normally secured as part of settlement at the end of the Customer session.

In most cases, recovery information stored in the transaction log can be used to ascertain the outcome of the transaction being recovered and a suitable completion record is then recorded at the time of recovery. In some cases the user is prompted to confirm whether or not the transaction has completed successfully and the response from that prompt is used to generate the completion record.

3.1.3.2 AP Recovery

In the case of Automated Payments (AP), the user is asked if they wish to carry out AP recovery and they have the option of doing so immediately or leaving it until later.

If the user carries out recovery they will be asked about the last successful AP transaction (which can be seen from the branch copies of the AP receipts that are printed) and the system will then check to see if it has been completed in the system. If it has not been completed in the system, then the system will use the AP Recovery data stored in the transaction logs to ensure that all incomplete AP transactions on the counter up until the one specified by the user are completed at recovery time. To assist with this process, each AP transaction has a unique, incrementing sequence number which is printed on the receipt.

Fujitsu understand that these processes are defined in Post Office's Horizon User Guides.

3.2 A counter has a "Blue Screen of Death"

This is just a special case of a counter failure, so please see section 3.1 above.

3.3 There are package collisions on networks

The replication protocols used to copy details of transactions between counters and also between the gateway counter and the data centre ensure that the data is copied successfully. Should packets collide on the network (or should there be any other network issues such as the IT communications link failing) then the replication protocols will ensure that the data is re-sent. Such retries will continue until the data is finally successfully transmitted.