

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

Document Title: HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN

Document Type: High Level Test Plan

Release: Not Applicable

Abstract: This document details the approach to proving Business Continuity and the Integrity of the HNG-X Solution.

Document Status: DRAFT

Author & Dept: Steve Bansal/Alan Child

Internal Distribution:

External Distribution:

UNCONTROLLED IF PRINTED

Approval Authorities:

Name	Role	Signature	Date
Pete Dreweatt	HNG-X Test Manager		
Andrew W Thompson	Post Office Ltd (Test Manager)		
Tony Wicks	Business Continuity Manager		

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

(*) = Reviewers that returned comments



0 Document Control

0.1 Table of Contents

0	DOCUMENT CONTROL.....	2
0.1	Table of Contents.....	2
0.2	Figures.....	4
0.3	Tables.....	4
0.4	Document History.....	4
0.5	Review Details.....	4
0.6	Associated Documents (Internal & External).....	5
0.7	7
0.8	Glossary.....	11
0.9	Changes Expected.....	13
0.10	Accuracy.....	13
0.11	Copyright.....	13
1	INTRODUCTION.....	14
1.1	Principles.....	14
1.2	Definitions.....	15
1.3	General.....	15
1.4	Resilience and Recovery Testing.....	17
1.4.1	Data Backup and recovery.....	17
1.5	DR Testing.....	18
2	SCOPE.....	19
2.1	Features to Be Tested.....	20
2.1.1	Counter.....	20
2.1.2	Counter Business Applications.....	22
2.1.3	Branch Access Layer.....	26
2.1.4	Online Services.....	29
2.1.5	Batch Applications.....	34
2.1.6	Branch Database.....	35
2.1.7	Technical Network.....	37
2.1.8	Platform and storage.....	49
2.2	Other Area's.....	59
2.2.1	Storage Area Network (SAN).....	59
2.2.2	System Qualities.....	61
2.2.3	Estate Management.....	68
2.2.4	Time Synchronisation.....	70
2.2.5	Active Directory.....	72
2.2.6	Domain Name Service.....	73
2.2.7	Remote Support and Diagnostics.....	75
2.2.8	Branch Support Database.....	79
2.2.9	Aurora Console Access.....	83
2.2.10	Authorisation Services.....	84
2.2.11	File Transfer Managed Service (FTMS).....	87
2.3	Features Not to be Tested.....	89
2.3.1	(Hydra) 'Migration States' Integrity Testing.....	89

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE

3	RISKS.....	91
3.1	Hardware and Software Risk Issues.....	91
3.2	Planning Risks and Contingencies.....	91
4	QUALITY.....	92
5	APPROACH.....	92
5.1	General.....	92
5.2	DR Test Cycles.....	94
5.3	Test Case Analysis.....	94
5.4	Test Case Execution.....	94
5.5	Approach to Resilience.....	94
5.6	Approach for DR.....	95
6	ENVIRONMENTAL NEEDS.....	95
7	RESPONSIBILITIES.....	95
8	DEPENDENCIES.....	95
9	SCHEDULE.....	95
10	RESOURCES.....	96
11	ENTRY CRITERIA.....	96
12	EXIT CRITERIA.....	96
13	TEST PASS / FAIL CRITERIA.....	97
A	APPENDIX – MANUAL TESTING.....	98
A.1	Explicit Requirements.....	98
B.1	Counter Architecture.....	103
C.1	Branch Access Layer.....	104
D.1	Online Services.....	104
E.1	Batch Applications.....	104
F.1	Branch DB.....	104
G.1	Counter Business Applications.....	104
H.1	Technical Network.....	104
B	APPENDIX – AUTOMATED TESTING.....	105
C	APPENDIX – PLATFORMS/COMPONENTS.....	106
I.1	Servers.....	106



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

J.1	Storage.....	108
K.1	Network.....	108

0.2 Figures

Figure 1	BladeFrame resilient network architecture.....	46
Figure 2	Platform Definitions in HNG-X.....	51
Figure 3	Server Placements for Active/Active Services.....	53
Figure 4	DNS service resilience overview.....	74
Figure 5	Remote Access Framework.....	76
Figure 6	Services Spear.....	93

0.3 Tables

Table 1	SLTs for overall availability.....	62
Table 2	Branches and counter availability.....	62
Table 1	Maximum Target Recovery Time.....	65
Table 2	Service availability at DR.....	66
Table 3	Recovery mechanisms for essential Business Systems.....	67
Table 4	EM recovery and resilience requirements for each platform.....	Error! Bookmark not defined.
Table 5	BladeFrame Failure Effect and Action.....	72

0.4 Document History

Version No.	Date	Summary of Changes and Reason for Issue	Associated Change - CP/PEAK/PPRR Reference
0.1		Initial Release	
0.2		Incorporation of comments and updates	

0.5 Review Details

Review Comments by :	Wednesday 12 th December 2007
Review Comments to :	<div> <div>steve.bansal@</div> <div>GRO</div> <div>@</div> </div> <div> <div>RMGADocumentManagemen</div> <div>GRO</div> </div>
Mandatory Review	
Role	Name
HNG-X Solution Design	David Harrison



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

HNG-X Testing	Peter Dreweatt
Optional Review	
Role	Name
HNG-X Development	Graham Allen
HNG-X Test Design	Peter Robinson
HNG-X Test Design	George Zolkiewka
HNG-X System Test	Harjinder Hothi
HNG-X RV Manager	James Brett
TE Manager	Peter Rickson
HNG-X SV&I Manager	Sheila Bamber
HNG-X Programme Manager	Phil Day
Tester	Hamish Munro
HNG-X Testing	Peter Dreweatt
HNG-X Core Services	Ed Ashford
HNG-X V&I Manager	Peter Rickson
Test Manager	Lee Farman (POL)
Test Manager	Andrew W Thompson (POL)
Test Manager	Chris Young (POL)
Acceptance Analyst	Wayne Roberts (POL)
HNG-X Service Transition	Steve Godson
Security Architect	Jim Sweeting
Core Services	Pat Lywood
System Qualities Architect	Dave Chapman
Counter Architect	Jeremy Worrell
Branch Access Layer Architect	Wille Faler
Online Services Architect	Andy Williams
Batch Application Architect	Roger Barnes
Branch Database Architect	Nasser Siddiqi
HNG-X Joint Test Team	John Halfacre

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

(*) = Reviewers that returned comments

0.6 Associated Documents (Internal & External)



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

Reference	Version	Date	Title	Source
PGM/DCM/TEM/0001 (DO NOT REMOVE)			Fujitsu Services Post Office Account HNG-X Document Template	Dimensions
VAL/GEN/PRO/0003			HNG-X HLTP Definition Report	Dimensions
TST/GEN/STG/0001			HNG-X Testing Strategy	Dimensions
PGM/PAS/PRO/0005			Test Planning and Preparation	Dimensions
TST/GEN/WPD/0005			HNG-X Work Package for Non-Functional Testing	Dimensions
TST/GEN/HTP/0002			HNG-X Disaster Recovery/Business Continuity Testing – Scoping Document	Dimensions
ARC/PER/ARC/0001			HNG-X Systems Qualities Architecture	Dimensions
ARC/APP/ARC/0003			HNG-X Counter Architecture	Dimensions
ARC/APP/ARC/0004			HNG-X Architecture – Branch Access Layer	Dimensions
ARC/APP/ARC/0009			HNG-X Counter Business Applications Architecture	Dimensions
ARC/APP/ARC/0005			HNG-X Architecture - Online Services	Dimensions
ARC/APP/ARC/0007			HNG-X Batch Applications Architecture	Dimensions
ARC/APP/ARC/0008			HNG-X Branch Database	Dimensions
ARC/NET/ARC/0001			HNG-X Technical Network Architecture	Dimensions
ARC/PPS/ARC/0001			HNG-X Platforms and Storage Architecture	Dimensions
ARC/SOL/ARC/0001			HNG-X Solution Architecture Outline	Dimensions
ARC/SYM/ARC/0001			System and Estate Management – Overall Architecture	Dimensions
ARC/SYM/ARC/0003			HNG-X System and Estate Management Monitoring	Dimensions
ARC/SYM/ARC/0004			Remote Support and Diagnostics Topic Architecture	Dimensions
ARC/SYM/ARC/0005			HNG-X Estate Management Component Architecture	Dimensions



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

SVM/SDM/PLA/0003			HNG-X Business Continuity Test Plan	Dimensions
TST/GEN/PRO/0001			HNG-X Testing Process – Entry and Exit Criteria	Dimensions
TST/GEN/PRO/0003			HNG-X HLTP Definition Report	Dimensions
TST/GEN/PRO/0004			System Test Definition and Process	Dimensions
TST/GEN/PRO/0006			Process for Creating Data for inclusion in HNG-X High Level Test Plan Reports	Dimensions
TST/GEN/PRO/0010			HNG-X Defect Management Process	Dimensions
TST/GEN/PRO/0012			HNG-X - Non-functional Testing Process	Dimensions
TST/SOT/HTP/0003			HNG-X Performance/Stress High Level Test Plan	Dimensions
REQ/CUS/STG/0002			HNG-X Branch Exception Handling Strategy - Agreed Assumptions and Constraints	Dimensions
DES/APP/HLD/0006			HNG-X Generic Authorisation Services High Level Design	Dimensions
DES/INF/LLD/0016			HNG-X Backup and Recovery Low Level Design	Dimensions
DES/PER/HLD/0001			HNG-x Resilience and Disaster Recovery High Level Design	Dimensions
DES/SYM/HLD/0015			POA HNG-X Backup and Recovery High Level Design	Dimensions
DES/APP/HLD/0051			File Transfer Managed Service Delta HLD	Dimensions
DES/NET/HLD/0006			Domain Name System	Dimensions
DES/NET/HLD/0007			HNG-x Storage Area Network (SAN) HLD	Dimensions
DES/NET/HLD/0008			Data Centre LAN High Level Design	Dimensions
DES/NET/HLD/0009			HNG-X WAN HLD	Dimensions
DES/NET/HLD/0013			Time Synchronisation at HNG-X – High Level Design	Dimensions
DES/NET/HLD/0014			HNG-X Branch Access Network High Level Design	Dimensions
DES/PPS/HLD/0003			Active Directory HLD Design for HNG-X	Dimensions
DES/SEC/HLD/0001			Strong Authentication High Level Design	Dimensions
DES/SYM/HLD/0017			Remote Support Secure Access Server High Level Design	Dimensions
DES/SYM/HLD/0020			AURORA CONSOLE ACCESS HIGH LEVEL DESIGN	Dimensions
DES/APP/HLD/0023			Branch Support Database High Level Design	Dimensions
DES/GEN/STD/0001			Host Applications Design and Development Standards (HADDIS)	Dimensions
DEV/APP/LLD/0026			Host BRDB Capture of Instance Unavailability Low Level Design	Dimensions
TST/GEN/HPD/0003			HNG-X Work Package for Disaster Recovery Scenario Testing	Dimensions
CS/SIP/002			Business Continuity Framework	PVCS
TD/STR/007			FTMS Resilience and Recovery Strategy	PVCS



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

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

Abbreviation	Definition
[R1]	A request for authorisation.
ACE	Application Control Engine
ACL	Access Control List
ADSL	Asynchronous Digital Subscriber Line
APOP	Automated Payments – Out Pay
ASDM	Advanced Security Device Manager
ASM	Oracle Automatic Storage Manager
BAL	Branch Access Layer - The interface between Counters and the Data Centre – a set of Web Service end points
B-B credit	FC packets are stored in buffers within FC targets, switches & host in order to accommodate physical media delays. A buffer to buffer credit (B-B credit) indicates the packet is in a staging point but not committed to store. In distance solutions, B-B credits are used as a means to daisy-chain unrelated FC frames onto the same media, ensuring maximum performance in the use of the transmission media.
BCV	Business Continuance Volume (a clone or snapshot based local copy) - storage devices to which copies of working volumes are made in order to carry out business activities
BGP	Boundary Gateway Protocol
BIN	Binary Input (also used to describe the EMC Symmetrix configuration file)
CSM	Content Switch Module. A network device that allows incoming requests for service to be load balanced across a number of platforms.
DA	EMC Symmetrix Disc Adaptor
DCS	Debit Card System; Horizon service that supports payment by Debit Card
DMX	EMC Symmetrix storage device with a non-blocking Direct Matrix architecture for connection of any host to any storage volume (can be thought of as a distributed modular cross-bar architecture)
DR	Disaster Recovery
DRV volume	Special volume in EMC Symmetrix that allows the movement of data between physical disks by means of Symmetrix Optimizer
DWDM	Dense Wave Division Multiplex
EMC CC / ECC	EMC Control Center ... storage Management Application
ETS	Electronic [Phone card] Top-up Service
FA	EMC Symmetrix Disc Adaptor / port – used for host connectivity
FC	Fibre Channel (a storage communications standard which uses 8:10 bit encoding – at 1Gb/sec speeds, bandwidth is 100Mbytes/sec as 10 bits are used to represent 1 Byte)
FE	Fast Ethernet



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

FRU	Field Replaceable Unit – the parts which get replaced when a component fails (i.e.: if a whole controller board is a FRU, the whole board gets replaced if a single component fails on it)
FSBN	Fujitsu Services Business network
FWSM	Firewall Services Module
Gb	Giga bit
GE	Gigabit Ethernet
GK / gatekeeper volume	Special volume in EMC Symmetrix used as a conduit for system communications for use by SymAPI – uses SCSI protocol
GPRS	General Packet Radio Service
GRE	Generic Routing Encapsulation (specified in RFC 2784)
HLD	High Level Design
HNG-X	Horizon Next Generation
HSCSD	High Speed Circuit Switched Data
ICMP	Internet Control Message Protocol
IDS	Intrusion Detection System
IP	Internet Protocol
IPSEC	IP Security
ISDN	Integrated Services Digital Network
ISL	SAN Inter Switch Links used to connect fabrics together and for switch to storage port connections
JRE	Java Runtime Environment
L2TP	Layer 2 Tunnelling Protocol
LAN	Local Area Network
LC	Lucent connector – a type of hardware connector used for Fibre cabling
LCP	Link Control Protocol
LINK	the organisation responsible for branded and shared network of cash machines and self-service terminals of certain member banks and building societies in the UK, which enables services from one member bank or building society to be available at cash machines of all member banks and building societies.
LNS	L2TP Network Server
LUN	Logical Unit Number
MAC	Media Access Control
Mb	Mega bit; Mega bits per second
MPLS	Multiprotocol Label Switching
MSS	Maximum Segment Size
N+1	N instances are sufficient to handle the load.
NAT	Network Address Translation
NBS	Network Banking Service
NMS	Network Management System



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

NNM	Network Node Manager
NPS	Network Banking Persistence Service
NTP	Network Time Protocol
OBC	Outlet Business Change
OOB	Out of Band
OS	Operating System
P2P	Peer-to-peer
PAS	HNG-X Private Internet Address space (PAS)
PE Router	Provider Edge Router
PHU1	Portable Hardware Unit 1
PO	Post Office
POL	Post Office Ltd.
PPP	Point to Point protocol
PSK	Pre Shared Keys
PSTN	public switched telephone network
QoS	Quality of Service
RA	EMC Symmetrix Disc Adaptor / port – dedicated to remote replication
RAC	Oracle Real Application Cluster
RAC Model	Request, Authorisation, Confirmation; basic model for Transactions where initial Online Request [R] from Counter elicits Online Authorisation [A] from the Service Provider. Confirmation [C] of outcome of Transaction is sent in near real time from Counter to the Data Centre.
RADIUS	Remote Authentication Dial In User Service
RAID	Redundant Array of Inexpensive Disks, a.k.a. Redundant Array of Independent Disks
RDMS	Reference data management system
RDP	Remote Desktop Protocol
RFC	Request For Comments
RIP	Routing Information Protocol
RMAN	Oracle
RMGA	Royal Mail Group Account
RPO	Recovery Point Objective – after a disaster occurs, the point in time to which a DR system needs to recover to (i.e.: RPO=0 means that data needs to be recovered to the point immediately before the disaster occurred = last transaction before the disaster = no data lost)
RTO	Recovery Time Objective – after a disaster occurs, the time a DR system can take to recover full functionality (i.e.: RPO=0, RTO=1 hour, means that the application needs to be recovered in one hour, to the last transaction immediately before the disaster occurred.)
SAN	Storage Area Network
SCSI	Small Computer System Interface
SFS Volume	Special volume in EMC Symmetrix holding the Symmetrix File System (configuration)



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

SID	Session identity
SNTP	Simple Network Time Protocol
SP	Service Provider
SRDF / S	Symmetrix Remote Data Facility (remote copy) /S = synchronous
SRDF / S	Symmetrix Replication Data Facility / S = Synchronous
SRRC	Service Resilience and Recovery Catalogue. A document describing the failure modes, business impact, events raised, and recovery mechanism for each service.
SSC	Systems Support Centre
SSH	Secure Shell
SSL	Secure Sockets Layer
Abbreviation	Definition
ST	System Test
Symmetrix	EMC Range of storage devices
TACACS+	Terminal Access Controller Access-Control System plus
TCP	Transmission Control Protocol
TLS	Transport Layer Security Protocol
TNS	Transaction Network Services
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunications System
USN	Uniqueness Sequence Number
Vast Volume	Special volume in EMC Symmetrix (from Engenuity 71 code) that will hold the contents of DMX cache as it is de-staged to disk in case of a power outage in excess of 60 seconds – stored in the first drives of a DMX system (number of drives dependant on system cache configuration) Also used for Special Volumes in EMC CLARiiON to hold system configuration, for de-staging of cache contents in case of power failure, to hold copies of microcode for upgrade, for reserved use by CLARiiON layered software = the first 5 drives in a CLARiiON array
VIP	Virtual IP
VSAN	Virtual Storage Area Network
VSAT	Very Small Aperture Terminal
WAN	Wide Area Network
WD	Work Description
WWN	Manufacturers World Wide Name assigned the HBA (host) and storage devices (FA) or switch ports – akin to MAC address in Ethernet

[illegible]

[illegible]



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

	(Fujitsu Business Continuity Framework, CS/SIP/002 v12 18/8/06)
Disaster Recovery	<p>The HNG-X configuration includes two sites in order to provide business continuity in the event of a catastrophic failure. However, the configuration is much simpler than Horizon— only one of the sites (the primary site) operates as the Production site, whilst the other site remains in reserve as a Disaster Recovery (DR) site, only to be activated in the event of a catastrophic failure of the Production site. Again, each site is sized to accommodate full projected peak workloads</p> <p>The HNG-X System Qualities Architecture (ARC/PER/ARC/0001) defines the following possible triggers for DR :</p> <ul style="list-style-type: none"> • Critical failure to a documented system SPOF (primary storage, blade frame cabinet) • Failure of all components critical to providing service where no immediate replacement can be provided— this could either be for the entire service (switches, servers (e.g. active & standby)), or where it would not then be possible to provide full service (e.g. at a level of N-3 for an N+1 configuration) • Long term power failure to entire data centre • Disaster affecting part of the entire data centre (accident, fire, and flood). <p>Use of the DR Data Centre to run live systems may also be required in the event of a PCI forensic investigation requiring the Live Data centre to be quarantined</p>
Hydra	The deployed solution from the start of migration to the completion of migration. i.e. While Horizon and HNG-X co-exist.
Integrity	Within this document integrity is a framework term applied to not only the combination of resilience and recovery but also error monitoring.
IRE11	The site of the Production HNG-X Service
IRE19	The site of the Standby system for IRE11 which also doubles as a Test system
F1, F2, F3, F4	These are the names of the 4 one week long DR test slots planned for mid-2008
Platform	<p>A type of server hosting a business application or infrastructure service that is part of the HNG-X solution and hosted in the HNG-X data centres, a platform can have multiple instances and is built from a Platform Foundation</p> <p>(ARC/PPS/ARC/0001)</p>
Platform Foundation	The combination of HNG-X approved hardware and a HNG-X approved operating system for the purpose of hosting an HNG-X application, service or function in the HNG-X data centre; the platform foundation is provisioned through an automated process, the operating system has been modified to comply with HNG-X policies, e.g. Security Policy (ARC/PPS/ARC/0001)

0.8 Changes Expected

Changes

This document is the High Level Test Plan that represents the testing that will be required to meet our objective to ensure FJS hand over to PO Ltd, tested processes and facilities too incidents which could threaten normal business operations. The production of the HLTP is expected to be a two stage



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN



COMMERCIAL IN CONFIDENCE

process. The first stage will include the Scope and Objectives of the testing (all sections up to the Appendix). The second part to be listings / reports of the Planned testing that is to take place and **Wi** include extracts from Quality Center (where appropriate) in the appendix of this HLTP.

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

© Copyright Fujitsu Services Limited (xxxx). All rights reserved. No part of this document may be reproduced, stored or transmitted in any form without the prior written permission of Fujitsu Services.

UNCONTROLLED IF PRINTED



1 Introduction

1.1 Principles

- FJS will define an overall High Level Test Plan for Business Continuity and DR and be primarily responsible for carrying out testing (except where agreed with the Post Office and Joint Test Team)
- The testing is to ensure that the stated Requirements for Resilience and DR can be met. DOORS requirements are in [Appendix A](#)
- Tests, test designs and test results will be mapped against the DOORS requirements held in Joint Quality Center as the basis for Acceptance within the HNG-X program.
- As resilience is a quality of the system, all architecture and design documents need to state how they will meet the appropriate requirements
- where resilience is a feature of the manufacturers specification of a hardware or software component, and that component has been accepted by FJS as meeting the manufacturers spec on delivery, testing of resilience features will not be carried out unless FJS have tailored or amended the configuration etc for the purposes of HNG-X
- Resilience testing will include both the failover and the fallback scenarios
- Testing will be planned, designed and executed on an incremental basis, with each subsequent cycle of testing building on, but not duplicating, testing of previous cycles.
- Testing will focus on the vectors or variances between environment, infrastructure or applications. i.e. where a component has not changed in its configuration, coding or design since the previous layer of testing it will not be re-tested
- Resilience testing will be carried out throughout the HNG-X program focussed on components which have changed through configuration management. i.e. it is assumed that Resilience endures unless physical changes made to the infrastructure or component
- Resilience testing will not just be the ability to cope with a failure but includes that the failure results in an appropriate alert being raised
- Recovery testing will ensure that practised processes are in place to backup and restore all appropriate systems as dictated by the design.
- Resilience testing will be carried out separately to DR testing but may contribute to the subsequent DR tests through ironing out or proving any issues etc
- DR tests will be carried out through planned DR test cycles and the final DR test F4 (4th DR cycle) will be run against the entire infrastructure immediately prior to Live
- Results of tests will be reviewed by the Joint Test Team and successful completion, failure or variances reported to FJS and to POL's Business Continuity Manager to either a) confirm that resolution/rectifications is required or b) accept the risk that testing has either not been carried out or that specific tests have failed and are not to be rectified prior to acceptance
- FJS will be responsible for maintaining details of test results within Quality Center
- The depth of review of test results will be based on the combined Risk Impact + Likelihood prioritisation held in Joint Quality Centre (as per the HNG-X Testing Strategy). Requirements with a Priority of 1 (Highest) will be subject to in depth review, whereas those with Priority 5 (Lowest) will be subject to a 'light touch' check.

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE

1.2 Definitions

The organisation The Business Continuity Institute (<http://www.thebci.org/>) has a glossary giving a common definition of the various Business Continuity Management terms. However, these definitions post date Business Continuity activities within Horizon and do not always match up with our internal use.

Business Continuity Plans protect against the impact of unexpected interruptions in business as usual arising from any unforeseen event affecting the continuity of business critical processes, functions, premises, people, computer systems and supply chain (Royal Mail Group Business Continuity Planning Framework v7.0 May 2005)

At a working level, Post Office Limited and Fujitsu Services (Post Office Account) generally recognise the term Business Continuity as having three closely related components:

Resilience may be defined as the steps taken to avert a loss of service or disaster or reduce the likelihood of a disaster or loss of service.

Contingency may be defined as the interim processes and procedures adopted during the loss of service.

Recovery may be defined as the business and technical arrangements to restore a lost system or service and manage the process of reversion to normal processing and full resumption of service.

For V&I testing stream Business Continuity is strictly limited to the resilience and recovery of the HNG-X and the changed Horizon components. For Disaster Recover (DR) we will work with the Business Continuity manager testing specific areas of concern, as well as the main Business Continuity test plan. ([SVM/SDM/PLA/0003](#))

Disaster Recovery relates to the elements of the hosting solution which provide hosting of the solution from a secondary data centre in event of a catastrophic failure at the primary data centre, with very little or no loss of data ([DES/PER/HLD/0001](#))

Resilience (Testing) relates to elements of the hosting solution which provide tolerance to faults within the primary data centre; one of the main design goal being that any single failure will not prevent the application from continuing to work within the primary data centre, Doors ID [ARC-492](#) ([DES/PER/HLD/0001](#))

Recovery (Testing) may be defined as the business and technical arrangements to restore a lost system or service and manage the process of reversion to normal processing and full resumption of service. ([CS/SIP/002](#))

Estate Monitoring relates to events and alerts that are raised on the estate. If a problem is not detected, then all the plans are not going to be put into effect. The timelines of detection and the response to any alarms raised are part of the overall impact to the customer of the service outage ([DES/PER/HLD/0001](#))

1.3 General

This document is the High Level Test Plan (HLTP) that represents the testing that will take place for the Business Continuity and Integrity testing stage, in accordance with the HNG-X Testing Strategy.

It seeks to fulfil the objectives as set out in both the document, HNG-X Work Package for Non-Functional Testing ([TST/GEN/WPD/0005](#)) and HNG-X Work Package for Disaster Recovery Scenario Testing ([TSTGENWPD0003](#)).

The objectives are stated as HNG-X Work Package for Non-Functional Testing:

Integrity Testing (Resilience and Recovery)

- *Devise a set of tests that need to be defined by taking inputs from the analysis of the potential points of failure, test objectives (based on risk prioritisation) and identifying system, configuration and infrastructure changes.*



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

- *To carry out the tests throughout the project and not just at the end of the testing lifecycle.*
- *Needs to take Disaster Recovery requirements, plans and testing into consideration.*

The Disaster Resilience model for HNG-X is based on one production (IRE11) and one standby (IRE19) Data Centre. Resilience within the Data Centre is achieved by duplication of servers, either load-balanced through CSM's or in active/standby configurations. Services that are load-balanced are deployed on N+1 server, temporary loss of one server does not affect business capacity. Services that are deployed in an active/standby pair support either the whole estate or a defined partition of the estate within the single active instance; these services are designed to support rapid fail-over.

Standardised platforms and Blade Frame technology for virtualisation will allow for rapid replacement of a failed server processing unit – another Blade picks up the task and restarts the server. The overall solution design minimises disruption to Post Office trading.

This is further expanded in the HNG-X Testing Strategy (TST/GEN/STG/0001).

The objectives are stated as HNG-X Work Package for Disaster Recovery Scenario Testing:

- Evaluate and discuss with the Infrastructure Team how the DR system will be constructed. It is the role of the Infrastructure Team to plan and build both the Data Centre Environments but as one of the main users of the Disaster Rigs we need to define the requirements the test teams will have on this environment.
- To liaise with CS to define a process on how to backup the Test Rigs on a daily basis. The present process for the Horizon project is that the SPTS routinely backup the Test Rigs on a daily basis. As the Test Rig will be reconfigured into running the live system within 3 hours in the event of a disaster, backing up the test environment will be very important as will to restoring it back to the test role when the problems have been resolved.
- To plan how to test the DR scenario along with all the other teams that need to be involved. In the event of a disaster to restore the environment will need involvement from more than one party. Therefore we need to plan with other teams how the DR plan will be put into action and how we test that plan. There maybe key areas that Test can check independently of the main plan to try and resolves issues before a 'Big Bang' approach with other teams
- To plan how DR testing will be executed producing HLTP's. Some areas to consider are:-
 - Testing of specific components and procedures that are part of the DR Scenarios
 - Testing of DR scenarios within the Fujitsu domain
 - Testing of the DR scenarios with connections to third parties
- To plan how to restore the DR environment to its testing role once the live site has been restored, in particular defining the requirements on this operation, e.g. time frame (unless this is already defined in the contract). Procedures and guidelines will need to be drawn up on how to action this. This will require the test unit to make sure all processes and procedures are in place and function according to the requirements.



1.4 Resilience and Recovery Testing

Most of the Resilience and recovery testing will not be part of a separate testing stream, but performed as part of the testing of the individual component. For those components where no evidence of Resilience and recovery testing can be found then tests will be need to be devised and run resulting from the raised defect. Adopting this approach it can be seen that an extensive test/document inspection phase is required relying on Quality Center being up to date.

There is a risk that using this approach will result in there not being adequate time to both test components and resolve any issues if it is discovered that Resilience testing has been missed, then this would either a) delay the program or b) mean going live with areas of untested resilience which could cause business loss subsequently.

It is essential that Quality Center be up to date. If it is not up to date then it will not be possible to determine what tests have been performed. There should be very few tests that do not result in Quality Center being up to date within a few minutes of the test completing. Volume testing would be one example where for instance significant analysis may be required to ensure that the test has been passed.

ARC/SYM/ARC/0001 and ARC/SYM/ARC/0003 describe the overall approach to Estate Management, and specifically to monitoring the estate. **If a problem is not detected, then plans will not be put into effect. The timelines of detection and the response to any alarms raised are part of the overall impact to the customer of the service outage**, although specific service level agreements may break the response down into smaller units that are more easily measured.

1.4.1 Data Backup and recovery

Backup and Recovery Design is covered in DES/SYM/HLD/0015 and DES/INF/LLD/0016.

Always remembering that it is critical to test Recovery.

Data backup and recovery strategy is detailed in the POA HNG-X Backup and Recovery High Level Design paper (DES/SYM/HLD/0015)

Various other solutions are identified within DES/SYM/HLD/0015 to cover migrated Horizon systems and the Hydra state of Horizon/HNG-X co-existence

In order to allow continued validation of fixes or extensions to functionality the Test System must be able to perform all backups and restores by all methods under control of the scheduling system and monitored by an analogue of the SYSMAN system.

Recovery testing is two phase. For systems that recover from a local dump (RMAN and SQL-Server) it is possible to test this recovery independent of the backup infrastructure. Similarly the method of BCV image or image set recovery is effectively being tested every time a test rig is being reset.

Many applications such as DRS have not changed since Horizon, and share a recovery mechanism. It is considered adequate to show a sample recovery to prove the generic processes still hold.

In general testing should separate testing of the backup infrastructure and generic processes, which is what is provided by this design, from specific application recovery which is covered by the design for that application, and which may include schedule recovery following the corruption.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

1.5 DR Testing

The HNG-X System Qualities Architecture ([ARC/PER/ARC/0001](#)) defines the following scenarios in which DR would be used:

- Critical failure to a documented system SPOF (primary storage, blade frame cabinet)
- Failure of all components critical to providing service where no immediate replacement can be provided– this could either be for the entire service (switches, servers (e.g. active & standby)), or where it would not then be possible to provide full service (e.g. at a level of N-3 for an N+1 configuration)
- Long term power failure to entire data centre
- Disaster affecting part or the entire data centre (accident, fire, and flood).

The HNG-X Resilience and Disaster Recovery HLD ([DES/PER/HLD/0001](#)) section 2 states:

In the event of a catastrophic failure, that is, a failure which renders the primary data centre unable to host the solution in a commercially viable manner, the hosting of the entire solution will move to the secondary site (a “site failover”).

To support this, all servers, network switches, routers, firewalls, storage and supporting infrastructure at the primary site will be duplicated at the secondary site, and will operate in manner where they are permanently ready to fail over. This places limitations on the use of such components for testing, and it may be necessary to deploy dedicated test equipment at the secondary site where use of DR equipment is prohibited by such requirements.

Following a site failover it is assumed that any issue or failure at the primary site will be resolved, and that the hosting of the solution will move back to the primary site (a “site failback”). A failback is disruptive to the branch service, and failback will always be a planned event that aims to minimise the service outage of failback. Such service outages do not normally count towards SLA targets. With the increasing likelihood of 24 x 7 counter operations it may be much more difficult than in Horizon to agree the timing for failback.

There are currently planned to be 4 (four) DR slots named F1, F2, F3 and F4. These slots will be used to incrementally test the DR capability and resolve any issues. The end result needs to be confidence that the planned full DR trial scheduled for Q4 2008 will succeed.

The planned content of the slots are as follows:

F1	Network Failover
F2	Network Failover and Issues Resolved, Rig Integrity
F3	Failover, Rig Integrity and Issue Resolved, Failback
F4	Failover, Rig Integrity, Failback and Issues Resolved

The low level details of these slots will be determined after discussions with Tony Wicks, POA CS Business Continuity Manager, IS Operations in IRE11 and other interested parties. For Horizon the Business Continuity Test Plan has a series of subtests. The intention is that we use the equivalent tests as developed as part of the HNG-X Business Continuity Test Plan [SVM/SDM/PLA/0003](#) rather than develop a separate set of tests.

For this activity the process, procedures and scripts for DR should be in place supplied by IS Operations in IRE11. However, previous experience has shown that there will be considerable effort needed to both correct and in some case write these. Where possible assistance will be given but ownership of this is with IS Operations in IRE11.



2 Scope

This document is concerned solely with Business Continuity (Disaster Recovery) and Integrity (Resilience & Recovery) testing related aspects of the solution and does not cover other non-functional requirements. Specifically Performance (Volume) testing, will be a separate discrete testing stage and subject to its own HLTP ([TST/SOT/HTP/0003](#)). However, it is likely that there will be some cross-over between testing stages as the work is to be carried out on the V&I test rig. This will need to be co-ordinated during test execution planning. For example, it may be necessary for some Business Continuity tests to be performed under full load and therefore appropriate for these to be conducted during the Performance Testing stage.

There is an issue regarding the scoping of the activity. Originally it was considered that the scope of the testing was:

- 1 Check all of the components for resilience and an absence of single point failure
- 2 Test and document the new active/standby DR configuration

However, for the DR it is clear that there is an existing agreed Business Continuity activity that includes documented procedures for Horizon. This is already being updated for HNG-X and it is difficult to see how this can be improved on. The Business Continuity activity is managed by the POA CS Business Continuity Manager and the intention is to liaise with the POA CS Business Continuity Manager. This document contains the tests as highlighted in [TST/GEN/HTP/0002](#) and in HNG-X Systems Qualities Architecture ([ARC/PER/ARC/0001](#)). These will need amending in the light of the discussions with the POA CS Business Continuity Manager.

The other area is the resilience/integrity of the components – the intention is that this will be made up of the components in the branches and all components in the Data Centre. This is again a problem area as the resilience of the branch components and some of the Data Centre should be tested as part of Unit and System (SV&I) testing. Therefore, the activity will be in two parts:

- a. Investigate the testing of the components and see whether or not resilience is or has been tested
- b. For those components not tested then scripts will be developed and the components tested

The components are broken down as per the architecture documents:

- i. Counter Architecture [ARC/APP/ARC/0003](#)
- ii. Branch Access Layer [ARC/APP/ARC/0004](#)
- iii. Online Services [ARC/APP/ARC/0005](#)
- iv. Batch Applications [ARC/APP/ARC/0007](#)
- v. Branch Database [ARC/APP/ARC/0008](#)
- vi. Counter Business Applications [ARC/APP/ARC/0009](#)
- vii. Technical Network [ARC/APP/NET/0001](#)
- viii. Platforms and Storage Architecture [ARC/PPS/ARC/0001](#)

The testing of these areas will be based on the Recovery and Resilience design of these components as stated in the Architecture and Design documents associated with each area.

Some sections will appear to duplicate material that has been presented earlier. This is intentional, and is designed to make the description of each component easy to read, there will also be a natural overlap.



2.1 Features to Be Tested

Ensuring that data centre components are resilient to a single component failure, Door ID [ARC-492](#)

Resilience and Recovery features of HNG-X will be tested or proved. This could be by an actual test or possibly by another appropriate approved method of acceptance (further details of acceptance methods in [section 5.1](#)). If the test is by another test stream then it will be made clear where the details of the testing can be found within Quality Center.

Given the move to the 2 Data Centre model some features of Horizon will need to be tested. Only those features that have changed will be tested. For instance if virtualisation is chosen then the resilience of the virtualised system will be tested. However, if it is, say, just a server move, rather than an upgrade, then the Resilience will not be tested. Until the design of the migrated system is clearer it is not possible to give details of the testing required but it is likely to be substantial.

Testing of the process of transferring connections between the respective HNG-X and the EDG data centres will need to be covered in HNG-X testing. DR testing will need to include testing of links from the HNG-X DR environment to the Live EDG with two-way data transfer, as well as all remaining non-EDG data transfers

2.1.1 Counter

The following text is taken from the Recovery and Resilience section of HNG-X Counter Architecture ([ARC/APP/ARC/0003](#)). The tests are outlined in [Appendix B.1](#) and text in green is used to clarify what may be taken as ambiguous statements.

Counter failures tend to have a more limited impact of service availability than failures in the central Data Centres. Nonetheless, it is important to consider counter failure scenarios, and make adequate provision for remedying or working-around them. This section considers a number of possible failure conditions affecting counters, high level test cases are then created ready for migration into Quality Center.

For more information, readers are referred to HNG-X System Qualities Architecture ([ARC/PER/ARC/0001](#)).

2.1.1.1 Counter Hardware Failure

If permanent or intermittent failures occur in counter input or output devices such as printers, monitors, keyboards and scanners, they will be noticed quickly by Post Office staff attempting to use them, who will be trained to log helpdesk calls. Engineers will attend the branch to repair or replace the.

This should be no change from Horizon in that the process users use to report hardware problems will continue to be used.

Intermittent failures affecting internal subcomponents such as hard disk drives may not be visible to users, but should result in events being written to the Windows event logs, where they can be picked up by SYSMAN and subjected to event correlation, to see if a pattern of failure is developing. It will therefore be necessary to ensure that events at the counter can propagate up through SYSMAN3 and be detected.

HNG-X Counters connect independently to the Data Centre, so the failure of a single counter usually only affects one user and only until a repair or replacement can be affected.

This should be no change from Horizon in that the process users use to report hardware problems will continue to be used.

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

2.1.1.2 Counter Software Failure

If any software component experiences a pattern of software failure in testing, then root cause analysis will be undertaken and a fix applied. This should be no change from Horizon in that the process users already use to report software problems will continue to be used.

Unfortunately, large-scale systems integration projects occasionally uncover patterns of software failure in Commercial Off-The-Shelf products, which cannot be reproduced or fixed by their vendors. If such problems are found, then alternative software will be considered, or additional SYSMAN or ClearDesk tasks will be incorporated to reduce their impact.

Fatal software failures experienced in the field will be recognised by SYSMAN's monitoring agents, and events will be created, which will trigger alerts to operators who can intervene to diagnose the problem.

Non-fatal software failures experienced in the field are also likely to be reported in Windows event logs or in proprietary log files. Whilst these may not immediately have any effect on the service, they present a risk of more significant associated problems occurring later on. To reduce this risk, a wider set of event types will be configured to alert operators during pilot deployments of the HNG-X application on both the NT and XP operating systems, so that non-fatal errors can be analysed during these periods, and a decision can be made about whether to alert on these non-fatal events during full live operations too.

2.1.1.3 Counter Data Failure

No transaction data is held locally on the Counter's hard disk drive, so there is no requirement for a data recovery process from failed Counter PCs.

All data stored locally on counters is used by at least one software component, so if software components discover data integrity or corruption problems in their data they will write error or warning events to their logs, which will then be forwarded to operators as for "Software Failure" (see earlier).

2.1.1.4 Failure of other Counters

In multi-counter branches there are some circumstances where the failure of one counter could cause an outage to the auxiliary functions on other counters at the same branch:

- Where the failed counter was acting as a local software distribution staging post.
- Where the failed counter was running CNIM2 in active mode to monitor the branch router.
- Where the failed counter was sharing a report printer.

The Counter High Level Design consider how best to minimise the effect of the failure of one counter on other counters at the same branch.

2.1.1.5 Network Failure

All HNG-X branches will have their own locally-installed Branch Router, which will detect WAN connection failures, and switch to an alternative connection type without the need for users to restart their application sessions. Depending on their duration, WAN failures are classed as "transient", "longer" or "permanent", and the impact of each type of failure on the user depends on the type of transaction

**COMMERCIAL IN CONFIDENCE**

they are trying to undertake. For more details readers should refer to *HNG-X Branch Exception Handling Strategy- Agreed Assumptions and Constraints*.

A network failure on the LAN is much less likely because of the relative simplicity of the network infrastructure, but could occur in a single counter, in the Branch Router, or in the Ethernet hub or cabling. If the LAN fails, all business transactions for all affected users will be prevented, so users will report the problem to the helpdesk who will deploy an engineer to repair or replace the faulty hardware. For the branch there should be no change from Horizon. The current process for reporting hardware problems will continue to be used.

2.1.1.6 Branch Router

DES/NET/HLD/0010 Branch Router HLD

The branch router is an 'of the shelf' router. Its primary purpose is to connect counters to the FSBN network.

It has no high-availability features built-in. If the router were to fail, an engineer will visit the site and install a new one.

There are multiple ways for the router to connect to the FSBN: ADSL, ISDN, PSTN, GPRS (3G). If the primary connection method were to fail, then the router will automatically connect via an alternative method.

2.1.1.6.1 ADSL

The router connects to one of six LNSs

The LNS peer with the FSBN and the C&W MPLS network

2.1.1.6.2 ISDN/PSTN

ISDN/PSTN connections dial in to C&W routers at diverse sites round the country.

2.1.1.6.3 GPRS/3G

Two GGSNs (Gateway GPRS support node) in different locations each connected to different FSBN POPs

2.1.1.7 Server Failure

Counters connect via the Branch Router to a number of sever platforms. All such platforms are deployed in redundant configurations, such that a single server failure has no significant impact on counter operations. This will be covered later as part of server failure.

2.1.1.8 Data Centre Failure

If the live Data Centre fails, Counter connections to Data Centre-resident servers will be broken.

No reconfiguration of the counter will be required in order for reconnect to the secondary Data Centre, though users will need to restart their sessions and re-authenticate.

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

2.1.2 Counter Business Applications

The following is from HNG-X Architecture – Counter Business Applications ([ARC/APP/ARC/0009](#)).

Each counter operates independently of other counters in the same branch.

The counter business application relies on a resilient network connection provided by the branch router. A counter cannot continue to trade if there is no network connection to the data centre. The resilience of the network connection is covered within other architecture documents.

2.1.2.1 Counter Business Applications Software

The counter business applications need to be robust and reliable.

All component design and implementation must consider the error conditions that can arise and the way that the business application can handle the error. The scale of the system is such that errors that are “only 1 in a million” will occur several times per week.

The software must be resilient to hardware errors in attached peripherals.

Any exceptions in lower level components need to be trapped and handled within the calling software. Errors need to be logged so that support staff can analyse the problems encountered, however error reporting must consider the potential load on support systems if all counters report errors at the same rate. Some filtering logic will be applied to avoid excessive error logging.

The overall development guideline is that the business application does not break and if it does it does it gracefully! This should be covered by functional testing and will be checked by an inspection of Quality Center rather having separate tests.

2.1.2.2 Branch Peripheral failures

The business application provides alternative means of data entry for input of data where the counter peripheral has failed: Bar code reader, magnetic card reader, scales.

Note that for some transaction types, the availability of an alternative means of data entry (e.g. fallback to PKE) is controlled by reference data.

The counter business application provides a print preview mechanism to display the intended print layout on screen. The Clerk must manually transcribe the data to produce a manual receipt / report.

There is no fallback mechanism for the counter printer when printing vouchers (such as Postal orders).

There is only one back office printer per site. It is connected to a single (fixed) PC, except in larger offices where it will be network connected. The business application provides a print preview facility to enable user to complete essential business processes in the event that the back office printer is not available.

In the case of touch screen and keyboard failures, the general design target is that everything that can be done via the touch screen can also be performed via the keyboard and vice versa. There are however, occasional exceptions to this rule, for example not all characters can be input via the keyboard.

In the case of the Rates board, the fallback mechanism is to print a copy of the current rates via the back office printer.

For PIN pad failures, there is a potential to fallback to swiped transactions from chip and PIN. However, this only applies to card payment (and then only as long as permitted by EMV). Failure of the PIN pad, effectively removes the capability to perform Banking transactions at the counter. If no other counter is available with a working PIN pad (e.g. Single Counter Office) then Banking transactions will not be available at that branch.

**COMMERCIAL IN CONFIDENCE**

This should be covered by functional testing and will be checked by an inspection of Quality Center rather having separate tests

Other cases of peripheral failure are effectively the same as failure of the counter base unit, e.g. VDU. For the branch there should be no change from Horizon. The current process for reporting hardware problems will continue to be used.

2.1.2.3 Branch Exception Handling

This section relates to exceptions that can occur within the branch due to failures in hardware or other software components outside the business applications layer.

The main scenarios described are around the failure to communicate successfully between the counter and the data centre systems – in particular, failure to write transaction data to the Branch Database.

The connections to the Branch Access Layer do not use sticky sessions. There is significant resilience within that layer and there is a high expectation that a retry of a message will succeed. This should be covered by functional testing and will be checked by an inspection of Quality Center rather having separate tests

The HNG-X Branch Exception Handling Strategy- Agreed Assumptions and Constraints

(REQ/CUS/STG/0002) provides additional detail. The business process surrounding such exceptions is summarised below

2.1.2.4 Online transactions

The counter will connect to the data centre for online transactions during customer sessions.

Where the transaction is recoverable, a recovery record will be written to the branch database prior to the connection to the authorisation service. Not all online transactions are recoverable – e.g. DVLA (which is a read only interface and conveys no state), is not a recoverable transaction.

There are further forms of online transactions which simply lodge a recovery record and / or an audit event, but require no further processing within the Data Centre: for example, prior to label printing within Postal Services transactions.

The counter will manage the connections to the data centre, timing out the communications based on configured timeout values, and following business processes agreed for the individual transactions.

2.1.2.5 Settlement

A counter needs to be able to connect to the Data Centre systems at the end of each customer session.

Recovery records will be updated at settlement time to indicate that they are now complete and do not require further recovery.

The counter will manage the connections to the data centre, timing out the communications based on configured timeout values.

In the case of settlement transactions, the counter will automatically retry to write the transaction data automatically a configurable number of times. It will also prompt the Clerk to decide if further retries are needed. Failure to write the transaction data will result in a forced log out.

**COMMERCIAL IN CONFIDENCE****2.1.2.6 Reporting and Other Data Access**

Several counter business applications use the Branch Database to store persistent business data in addition to the Transaction data written in the Settlement transaction. Examples are administration functions such as User, Stock units, pouch details and bar codes allocated to pouches, and Track & Trace. The reporting applications retrieve the reports data from the branch database, and also write "cut off" markers to delineate parts of the reporting data that have already been processed.

The business application will query the Branch Database via the Branch Access Layer for data. Where data is to be written back to the branch database, the applications will do this in a similar way to Settlement transactions. The business application will use recovery records as appropriate where the resilience characteristics dictate that the results of a specific counter operation must not be lost.

2.1.2.7 Recovery and Recovery records

The format of the data written in the recovery records is transaction specific, but there will be some required fields such as User Session identifier, Branch / terminal identifier and Transaction Identifier.

The application must construct the recovery records according to the business rules, for example Banking and DCS recovery records need to be PCI compliant.

Recovery records will be updated at settlement time to indicate that they are now complete and do not require further recovery.

- It is anticipated that there will be some form of local session level checking to ensure that all recoverable transactions started within the customer session have been completed – and the recovery process started if there is any not completed. [This will check against errors at application level].

The recovery process will be performed on the original counter after a failed session.

The Recovery process will be driven by the counter application. The business application will obtain from the Data Centre the list of recovery records outstanding, together with the associated recovery record details.

For some online services (Banking, Credit / debit cards and E-top Ups), the counter will enquire from the authorisation agent the outcome of the original authorisation request.

2.1.2.8 Log on after failed sessions

The recovery process will be performed after the first logon at the original counter that failed.

The system will detect a user logging on after a failed session, either any user on the original terminal or the original user on a different terminal – however, the current assumption is that recovery will be outstanding until a user logs on to the original terminal.

- A more general recovery process could be examined again during the design stage. Consideration would need to be given to constraints in the Branch Database, the ability to detect race conditions to prevent parallel recovery sessions, and the impact on the audit trail.

The Branch Exception Handling Assumptions and Constraints CCD provide additional detail – particularly on the new concept of a "Recovery Receipt".

2.1.2.9 Log on when no network connection

The system should detect when there is no network connection and make the user aware prior to presenting the log on screen. The precise form of the user dialogue will be agreed during the design stage.



2.1.2.10 Audit records

One of the major audit requirements is that there are no gaps in the audit trail. The solution adopted is that when the write of the audit record has not been confirmed, the application will re-attempt to write an audit records a number of time (under user control). The record contents must not change between each attempt to write the data.

However, if the user confirms that the attempt must be abandoned, then the user session will be terminated by a forced log off. This behaviour must be considered within the design of individual business transactions, and appropriate use of recovery records used if the application could reasonably carry on in certain failure scenarios, or where the counter application needs to record specific state changes that must not be lost.

UNCONTROLLED IF PRINTED



2.1.3 Branch Access Layer

The following is from HNG-X Architecture – Branch Access Layer ([ARC/APP/ARC/0004](#))

For resilience purposes, the system will need to be able to run services in relative isolation from each other. The degradation or even loss of one service should not adversely impact any other services.

Resilience is addressed in the general architecture by virtue of the split into online service routing and the local data access services architectures. The online service routing architecture more closely addresses resilience.

Recovery does not need to be addressed in much detail, as the architecture is stateless, and any state stored is held in a database. Any recovery after failures such as corrupted files etc is best addressed through a full redeployment of the HNG-X applications. This should be covered by functional testing and will be checked by an inspection of Quality Center rather than having separate tests

The physical resilience architecture is based on an $n+1$ server's architecture. This means that the loss of a single server is not disastrous; however it does mean that the maximum concurrency of the system will be impacted in proportion to the loss of hardware to total hardware. For instance, if we have 10 servers, and one fails, it will mean a loss of 10% of the capacity in terms of maximum concurrency and load. A test is required to ensure that the loss of one server does not degrade the service to the extent that it fails the SLAs.

2.1.3.1 Start-up

On start-up, the BAL will need to automatically be able to start individual Services automatically.

2.1.3.2 Branch database connection fail-over

The branch database connectivity, as detailed in section 2.4.7 of [ARC/APP/ARC/0004](#), will need to have intelligent routing, testing and fail-over capabilities for the BRDB connectivity. This means that the database connectivity will need to do the following things:

- Test the connection.
- If a connection fails to connect to a BRDB node altogether, change BRDB node.
- Handle total failure in a managed fashion.

2.1.3.3 Online Service Routing fail-over

The Online Service Routing architecture will need to be able to handle failure of backend systems in a managed way. This means that it will need the following characteristics:

- Handle time-outs.
- Handle retries.
- Handle alternative network routes.
- Handle complete failure.

2.1.3.4 Other considerations

As the servers are stateless, the fail-over and disaster recovery characteristics of the Blade Frames will be sufficient. These considerations are outside of the scope of the BAL.

**COMMERCIAL IN CONFIDENCE****2.1.3.5 Branch Access – Hydra**

The following is drawn from HNG-X Resilience and Disaster Recovery HLD ([DES/PER/HLD/0001](#)).

2.1.3.5.1 Correspondence Server

Correspondence Servers host the Riposte Message Store distributed database, of which each counter is also a member.

The total message store has been manually split into four clusters, each of which contain approximately one quarter of the branch estate (both in terms of size and performance).

Each message store is supported by four servers, known as neighbours, two at each site. In Horizon one server at each site uses EMC storage to allow BCV backups, and the other is on Compaq RAID array in case of EMC failure. In HNG-x one server will use DMX-A and the other will use DMX-B.

This is analogous to Branch Database resilience, except that Riposte has all four members active, whereas Branch Database is active/standby at the primary site with a failover delay to the secondary site. This reflects the higher Horizon availability requirement for PAS/CMS which is now defunct.

Active/active. No failover.

If a rebuild is required there is a procedure for recovering the message store either from a backup or by replication from a surviving neighbour.

2.1.3.5.2 Generic Agent

The agents run services (confusingly also known as agents) which allow messages to be passed between Riposte and the back-end databases. There are three different classes, online, bulk load, and bulk harvest.

There are also streams running in the daytime maestro batch schedule, such as pouch delivery, which use the bulk load agents to turn LFS into a sort of online system with high latency, so it is not simple at Horizon to look at a stream and say whether it is batch or on-line.

EACRR is used to track the agents in the pool, and make sure that one (and only one) of each type is running. In practice most back end systems are able to cope with multiple agents, as the agent recovery is typically a reharvest which generates duplicate input anyway.

Stateless. No complex failover just restart.

Recovery by reprovisioning.

2.1.3.5.3 KMS

Key Management Service. A SQL-Server database running on Windows NT4 in a special security domain. The database and any required file store are on the Key Management Server S: drive which is on EMC SRDF replicated storage. BCVs are not used, SQL-Server does a dump to disk and areas of the S: drive are backed up.

One of the most important uses for these keys is to encrypt the counter message store. The complexity of doing this and the need to recover transactions from "dead" counters that may not have replicated to the correspondence servers has been one of the big drivers for HNG-x.

The KMS has a hardware random number generator, which mandates that it be outside the BladeFrame unless a software equivalent is approved (the company that makes the hardware RNG has stopped because it believes the software one is better, but the approval process is complex).

Resilience: fail over to secondary.



2.1.3.1.4 VPN

The VPN servers are really part of the network infrastructure, but rather than being hosted on appliances they are physically separate NT4 servers that run active/active.

Active/active. No failover.

Recovery by reprovisioning.

2.1.3.1.5 RADIUS servers

This controls login to the VPN service by the counters.

There are several types of RADIUS server for the various types of connection e.g. ADSL, ISDN, FRIACO, Branch Router, and there is peer resilience amongst each type of server, but at Horizon not N+1 resilience following loss of a site.

The performance balance is weighted to ISDN at the moment, but this is expected to have moved considerably towards Branch Router by the time of data centre migration.

UNCONTROLLED IF PRINTED

2.1.4 Online Services

The following is from HNG-X Architecture – Online Services ([ARC/APP/ARC/0005](#))

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

The term "online services" is misleading – all Counter access to the HNG-X Data Centre will now involve an online service of some form. Online service in this instance means required during a Counter transaction and not provided directly by the Branch Access Layer.

Online services are invoked from the Counter via the Branch Access Layer, which may contribute to the functionality of the service e.g. by writing recovery data to the Branch Database.

2.1.4.1 Session recovery Overview

HNG-X must cope with [perceived] failure to deliver a session's settlement data to the Data Centre and subsequent cancellation of the session. Session recovery has to take account of the difference in behaviour of transactions with and without an on-line component, so a summary is included here. See [REQ/CUS/STG/0002](#) for further details.

Transactions written to the session basket are either cancellable or recoverable. Some transactions are cancellable up to a point before becoming recoverable. One class of recoverable transaction is where communication with a third-party is used for authorisation, causing a change of state and implicit transfer of funds between accounts based on that transaction. Another class is where the transaction creates an item of value (e.g. prints a parcel label) and/or acquires a resource (e.g. allocates a bar code that is not reusable). To allow recovery following a cancelled session, a recovery record must have been written to the Branch Database for each recoverable transaction in that session.

Where network problems prevent a Clerk from completing a session, he or she can cancel the session. Depending on the point in the session at which the cancellation occurs, settlement data for the session may have already been committed. The Clerk follows a defined process for cancelling cancellable transactions and settling recoverable transactions, printing a Recovery Receipt and any Customer Receipts required (which may include void receipts for cancelled transactions). The Clerk is forcibly logged off at the end of this process.

The Data Centre tracks logon and logoff for each counter. During logon at a Counter, the Data Centre can detect if that Counter was previously logged on but that a normal logoff did not occur. In that situation the Data Centre can invoke the recovery process at the Data Centre and Counter. The recovery process attempts to bring the Data Centre/External Client view of the session state in line with the Counter's view. For example, if the settlement data had been committed, that data might contain transactions that the Counter regards as cancelled, and the system would then need to reverse out those transactions.

Any recovery records for an interrupted session can be retrieved by the Counter. A recovery record may simply identify the transaction, requiring a further request to the Data Centre to retrieve the transaction status, or it may contain all the transaction details necessary for recovery, such as the transaction amount. No sensitive data is stored in a recovery record.

Recovery records are written to the Branch Database by the Branch Access Layer. Writing of the record is triggered by the presence of an item in the XML of the service request with a well-known tag. The recovery item contains at least a Transaction Recovery identifier but will also include any data that the application needs for recovery. Recovery records are marked as completed when the transactions to which they refer are committed to the Branch Database.

This should be covered by functional testing and will be checked by an inspection of Quality Center rather having separate tests

2.1.4.2 Identifying transactions

Each Branch is assigned a unique numeric identifier, and each Counter within the branch is identified by a unique numeric logical position – if a Counter is physically replaced its logical position does not change. These two numbers are sufficient to identify the source of any message sent to the Data

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

Centre. A globally unique identifier to be used in a message can be generated by combining the branch and counter identifiers with a unique value from a Counter specific sequence.

The Branch Database holds a sequence number, known as the Uniqueness Sequence Number, or USN, for each Counter in the estate. The USN to use at a Counter is supplied at logon, and the Counter tells the Data Centre what the [updated] USN is as part of settlement.

When a recovery record is written, it contains at least the Branch and Counter identifiers and the current USN, and the USN is incremented. These three fields together make up a unique *Transaction Recovery Identifier*.

If a session is terminated abnormally, the Data Centre will not have the latest USN used by the Counter, but at the next logon it can determine the highest USN that has been used in a Transaction Recovery Identifier, and can therefore provide the Counter with the next USN for it to use. Note that the Counter may have allocated a higher USN than this during the abandoned session, for some purpose other than recovery, but only USN's that have actually been written to the database need to be unique at the start of the next session.

The Authorisation Agents require a correlation identifier to tie together the Request, Authorisation and Confirmation messages that make up an online transaction. The Transaction Recovery Identifier could have been used for this purpose but the Transaction Processing System expects the correlation identifier to be expressed as a string with a particular syntax, thus:

pp-bbbbbb-cc-nnnnnnnnnn-uu

This form of identifier is known as an HTxnNum. The format is common between Horizon and HNG-X, but the derivation of the fields is different between the two systems:

Pp	is a fixed prefix; 00 for HNG-X, 44 for Horizon
bbbbbb	is the Branch Identifier
Cc	is the Counter Identifier
nnnnnnnnnn	is the USN in HNG-X, and a Riposte Message Number in Horizon
Uu	is a fixed suffix for HNG-X, with the value 1; for Horizon, it is a value in the range 1-99 that must be combined with the Message Number to give a unique value.

Leading zeroes in the variable fields are suppressed, giving identifiers such as 00-49934-2-787239-1.

This should be covered by functional testing and will be checked by an inspection of Quality Center rather having separate tests

2.1.4.3 Retrieving RAC Transaction Statuses

When performing recovery for a transaction in the DCS, ETS or NBS domain the Counter may need to retrieve the status of the transaction from the NPS to answer such questions as "did the [R1] reach the external client?", "was the transaction authorised?", and so on. The data stored in a recovery record for one of these domains must provide sufficient information about the transaction to allow the status

**COMMERCIAL IN CONFIDENCE**

request to be generated. In particular, the recovery data must include the Routing Gateway and Agent Hash values that identify the Authorisation Agent to which the [R1] was sent. Because of the need to partition the NPS for efficient access, the Agents may require additional transaction identification over and above the Transaction Recovery Identifier. Such additional data will be identified in the High Level Designs for the Authorisation Agents.

This should be covered by functional testing and will be checked by an inspection of Quality Center rather having separate tests

2.1.4.4 Failure of an Authorisation Agent

A logical Authorisation Agent is implemented as an active/standby pair of agents running on separate platforms within the active Data Centre. Both agents in the active Data Centre start in standby. Heartbeats exchanged through the NPS are used to decide which agent becomes active and when and if the standby agent should take over.

When an Agent instance fails all outstanding transactions being processed by that instance are abandoned and the agent terminates. No attempt is made by the replacement instance to recover work (e.g. from the NPS). Failure is detected at the Counter by timeout. Where the Counter generates a transaction reversal following a timeout the transaction status held in the NPS determines whether an explicit reversal needs to be sent to the external client.

The Authorisation Agents depend on other system components to monitor their state and restart Agents that fail. As for normal service start, heartbeats exchanged through the NPS are used to decide which instance within the logical agent becomes active.

The Authorisation Agent for LINK acts as a server for the external service and is accessed via a CSM. Only the active Authorisation Agent for LINK advertises itself to the CSM as available. On fail-over, LINK reconnects and the CSM routes connections to the newly active instance. The other Authorisation Agents act as clients of their external services. On fail-over the newly active agent establishes new connections to the client.

An active agent, will not failover on loss of its connections with the external client. The resilience built in to the network is such that the standby agent is unlikely to fare better, particularly as both active and standby agents are in the same Data Centre.

Change of state from active to standby and vice-versa is notified to all BAL components connected to the agent.

2.1.4.5 Network Configuration Errors

The protocol between the Authorisation Agents and the BAL, described in [DES/APP/HLD/0006](#), allows each component to check that it is connected as expected. The messages that flow through the Authorisation Services identify their expected target, allowing rejection of misrouted messages.

2.1.4.6 NPS failure

The NPS is implemented as two Oracle instances, with the Agents maintaining permanent connections to both instances¹. If the Agent's preferred instance fails the Agent continues operating using the other instance, while trying to re-establish the failed database connection. Failure of both instances leads to termination of the Agent. Other system components should detect this and perform a controlled restart.

¹ More correctly, the critical threads maintain connections to both instances, while less critical threads switch between Oracle instances as necessary.

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

Where an Agent has failed over to its non-preferred NPS instance it will, if possible, switch back to using its preferred instance overnight.

2.1.4.7 Guaranteed Reversals

The business rules for an Authorisation Agent may require guaranteed delivery of reversals.

If an agent sends an explicit reversal to its client and fails to get an acknowledgement it will repeat the request up to a configured retry limit.

To ensure that reversals sent from the Counter are not lost, two routes are used for their delivery. A Counter requests a reversal explicitly through the BAL, and it includes the reversal in the session's settlement data. When the session is settled the reversal is committed to the Branch Database. A near real-time process transfers reversals from the Branch Database to the NPS from where they are polled by the agent. The transaction status in the NPS is used to discard duplicate reversals.

Note that no recovery record is written for the reversal request. A recovery record will have been written for the original authorisation and will remain in the Branch Database until the settlement data is committed. If the session is cancelled the recovery record for the original authorisation request is sufficient to allow the recovery process to determine the transaction status and generate any necessary reversal. The number of cases where a Counter can get through to the BAL with a reversal request but the request fails to reach the Authorisation Agent does not warrant an additional recovery record update.

This should be covered by functional testing and will be checked by an inspection of Quality Center rather having separate tests

2.1.4.8 Web Service Failure

The internal and training Web Service applications run within Interstage Application Servers. Interstage monitors the processes and restarts them in the event of failure. For resilience there are N+1 instances of each Web Service, running on separate platforms, where N instances are capable of handling the load from the whole estate. This architecture assumes that N=1; performance measurements are required to confirm the number of instances actually required. The CSM between the BAL and these Web Services monitors their availability and routes traffic accordingly.

In addition to the Interstage monitoring, health checks allow Systems Management components to monitor service availability and hence trigger remedial action such as restarting the service.

This functional part of this failure should be covered by functional testing and will be checked by an inspection of Quality Center rather having separate tests.

A separate test will be required to ensure that the performance can still be met with the failure of a single instance.

2.1.4.9 Disaster Recovery

The Disaster Resilience model for HNG-X is based on one live and one standby Data Centre. During normal operation the standby Data Centre is used for testing.

The stand-by Data Centre is a copy of the live one. The systems at the DR site are capable of providing the same functionality (capacity, performance, resilience and backup) as the primary site. Note that the use of Blade Frame technology means that there is only one set of logical servers i.e. when failing-over Blade-based servers to the DR site the same logical server images execute on the DR Blade as were executed on the live Blade. See [ARC/PPS/ARC/0001](#) for further details of Blade technology.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN



COMMERCIAL IN CONFIDENCE

Systems crucial to POL operation should have no data loss; storage for these systems will be synchronously replicated from the active to the DR Data Centre so that any transaction committed will be persisted on the storage at both Data Centres.

Both live and standby Data Centres contain the necessary Networked Security Processors. The NSPs are accessed via virtualised addresses to enable local access at each site. (Note that additional separate NSPs need to be available at the standby site for use in testing).

UNCONTROLLED IF PRINTED

2.1.5 Batch Applications

The following is from HNG-X Batch Applications Architecture ([ARC/APP/ARC/0007](#))

The Batch Applications are host database applications which run under Solaris in the Horizon system. These systems provide the delivery links to POL, POL clients, and POL service providers as well as the deliveries from these back to branch.

These systems provide and receive branch data through the branch database via The Branch Access Layer. Specialist batch transfer processes running on the database servers transfer the data via database links to / from these systems.

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

Data for all the host database systems will be replicated via a synchronous link to the second Data Centre. This guarantees that no transactions will be lost. Failover will be manually controlled on servers. There will be a standby Solaris server at each site that can take over in the event of a failure of the Solaris host. Failover to the DR failover site is an all or nothing operation. All the requisite software will be in place at the failover site and in the event of failover, Oracle instances for each of the host servers will be brought up at the failover site.

Data that does not require such a high level of protection and availability will be hosted on more cost effective secondary storage. Where required this data will be replicated to the second Data Centre via an asynchronous link or a scheduled replication mechanism.

Historical and audit data will be placed on dedicated Centera storage arrays and the contents are replicated to the second Data Centre.

Both Data Centres will contain all the appropriate management systems to allow for the management of all storage platforms from either Data Centre.

The NPS Host is run on an Oracle RAC server with load balancing between two nodes. Failure of one node is not necessarily fatal. Applications would only be re-homed outside core hours following recovery after a failure.

APOP is also run under Oracle RAC, but is only run on one node, hence would be required to failover to a standby node or wait for recovery of its node on Blade failure.

NPS and APOP are run on Bladeframe servers which have the ability to quickly switch to an alternative Blade.

Resilience has been built into the schedule of the applications where possible e.g. to cope with possible late or non arrival of data files (e.g. Link Rec files).

Remote access for POL is to be via two application servers which will be under Windows 2003 running on Blades giving extra resilience.

For NPS and APOP the DR site does not have separate code images but shares the image on the SAN.

See section 6 of HNG-X Solution Architecture ([ARC/SOL/ARC/0001](#)) for further detail.

2.1.6 Branch Database

The following is from HNG-X Branch Database ([ARC/APP/ARC/0008](#))

HNG-X stores branch data in a new centralised Branch Database in the Data Centre.

The Branch Database replaces the Riposte Message Store used by Horizon.

The database uses Oracle version 10gR2 and it uses an Oracle Real Application Cluster (RAC), which runs the database over multiple nodes (servers).



2.1.6.1 Hardware and software failure

Oracle RAC provides a high degree of resilience. If the hardware or software on one node fails, the surviving nodes carry on running. BAL redistributes the workload of the failed node on to the surviving nodes. The workload distribution is through re-routing of messages based on metadata stored in the Branch Database. The workload distribution is even to ensure that the surviving nodes remain well balanced.

The node configurations chosen ensure that the system can run at its design limits if a single node fails.

2.1.6.2 Data corruption

It is possible, but extremely unlikely, that the data written to disk is corrupted. However, Oracle will identify this type of data corruption and raise an alert. In most cases, data corruptions can be resolved without affecting services. In rare cases, the recovery of the entire Oracle database may be necessary. This can take many hours, as data recovery is from backups and then archive logs have to be re-played.

To avoid this situation, and ensure uninterrupted service to Post Office the system keeps an additional local copy of the data. The local copy is maintained using Oracle Data Guard. This acts as a standby database. By using a standby database, restoration of trading is achieved within minutes. For performance reasons Standby Database and the Branch Database must be in the same BladeFrame cabinet. The standby database uses the same nodes as live database while recovery of the live database takes place. This delivers similar performance to live. After the live database has been restored, switch over to live from standby only happens outside core Post Office hours.

Standby Database and the Branch Database use separate storage arrays within in SAN.

DN: The high-level design should provide a list of recovery scenarios including the steps the DBA needs to follow to recover from the situation.

2.1.6.3 Disaster recovery

The DR Data Centre has an Oracle RAC configuration that is identical to the live data centre. Both live and DR Data Centres have identical copies of business data. Synchronous replication between the two Data Centres helps in maintaining identical copies. No business data is lost if the primary Data Centre suffers a catastrophic failure.

2.1.6.3.1 Steams

This section is taken from [DES/PER/HLD/0001](#)

Oracle Streams is used to propagate data from a source database, often a reporting or data warehouse type system. This allows large queries to be run in near real-time on the target database without impacting the performance of the source database.

In HNG-x the Branch Database will use Streams to send data to the Branch Support database. The consequence for resilience and DR is that a Data Guard failover to the Branch Standby must continue to replicate via Streams to the Branch Support database.

This is described in detail in the Branch Database Recovery HLD

2.1.6.4 Branch Database Unavailability

The following details are taken from the Host BRDB Capture of Unavailability Low Level Design [DES/APP/LLD/0026](#).

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

The Branch Database (BRDB) will reside on a four node Oracle Real Application Cluster (RAC). Database access will generally be via a specified Node/Instance name (such as BRN1, BRN2 etc.), rather than via a Service (where the particular Node/database instance used for the connection is transparent to the user).

To balance the load across the nodes, a bespoke mapping process will be used to control which database Instance each Post Office outlet will use. When an Instance becomes unavailable, a mechanism is required that will allow subsequent processes that would have used the Instance that is now unavailable, to be routed to one of the remaining, available Instances. This document provides the Low Level Design (LLD) for this mechanism.

It should be noted that the availability of a database instance is affected by the availability of the Automated Storage Manager (ASM) and the node itself.

UNCONTROLLED IF PRINTED

2.1.7 Technical Network

- 2.1.8** Every discrete server that connects to the network will have at least two NICs. Each NIC will connect to a different network switch. The NICs will be configured in an Active/Passive configuration (not load balanced)

Discrete Windows and Linux servers will utilise Broadcom 57xx NICs, and run in Broadcom SLB [Smart Load Balance] mode. This is a feature of the Primergy RX300 platform rather than Windows or Linux.

2.1.9



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

The following is from HNG-X Technical Network Architecture ([ARC/NET/ARC/0001](#))

This section describes the capability and technique to handle a defined set of exception conditions within the context of the business availability requirements.

2.1.9.1 Outline Approach

This is based on reuse from Horizon as far as possible, an exception being SSL offload, the following table summarises the techniques;

Component	Resilience mechanism
Branch Access Network	<p>Network Paths</p> <p>The network is engineered to avoid single points of failure in the “high order part of the network” by use of multiple components and links.</p> <p>A Routing protocol is used between Network components in the HNGX domain and those in the service provider domain to maintain optimum network paths.</p> <p>Network Paths from Branch</p> <p>Wireless WAN (GPRS/EDGE/3G) will be used as a backup to ADSL. The Branch Router will determine when ADSL has failed and switch to the backup network. Once the ADSL service is restored, the branch router will verify this using a traffic test and switch back.</p> <p>Single Points of failure</p> <p>For those parts of the Branch network over which POA has design responsibility, single points of failure in the high order part of Branch network are avoided. High order is intended to mean any network component which concentrates traffic (statistically or in absolute terms) from a significant number of post office branches.</p> <p>Since there are components in the service provider cloud which represents single points of failure. For example in the BT ADSL service, there are over 100 devices on which the branch sessions terminate (the exact number is not provided by BT). The failure of any such device will result in no service to about 1/100 of the branches (assuming a reasonable spread). Such single points of failure represent a trade-off between network service costs and availability. These trade-offs are made and agreed prior to service introduction with the customer and service management. Mitigations are agreed with the service provider to minimise both the downtime and the likelihood of downtime.</p> <p>Application endpoint</p> <p>The Application at the branch targets a Virtual IP address which represents a service. This Virtual IP address is created on the ACE and the ACE selects a working application server to forward the TCP connection. This decision is based on regular probing by the ACE of the Branch Access Layer Servers for Application availability.</p> <p>Therefore the branch application does not need to locate a functioning application server. Rather this is provided transparently by the network.</p>
Access Tier Firewall	<p>These are configured as an Active / Standby pair with TCP connection state replication between the pair. When the Standby Firewall determines that the Active Firewall has failed then it will (by design) take over with no loss of Application connectivity due to loss of state.</p>
ACE used as layer 4 Load	<p>At each HNGX Data Centre there are 2 Catalyst 6513 switches each with one ACE blade. The ACE blades are configured in an active standby pair with state replication</p>



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

Balancer	<p>between them. When the Standby ACE determines that the Active ACE has failed then it will (by design) take over with no loss of Application connectivity due to loss of state.</p> <p>Note that the 6513 are on separate power phases.</p>
Distribution Tier Firewall - FWSM	<p>At each HNGX Data Centre there are 2 Catalyst 6513 switches each with one FWSM blade. The FWSM blades are configured in an active standby pair with state replication between them. When the Standby FWSM determines that the Active FWSM has failed then it will (by design) take over with no loss of Application connectivity due to loss of state.</p>
LAN Connectivity	<p>Application servers</p> <p>These are configured with two physical network interface cards that are “teamed” to create a single logical interface. Each such card is connected to a separate Catalyst 6513 switch port. This arrangement means that the failure of any one of {network interface card, Catalyst port, Catalyst switch} does not result in loss of LAN connectivity. Failover takes less than 2 seconds.</p> <p>Next hop for application server</p> <p>Application servers do not participate in Routing protocols. The next hop gateway is determined by the having two possible Gateways running VRRP (Virtual Router redundancy protocol). In the event that the Active Gateway fails, the standby will take over this role with no loss of Application connectivity.</p>
SAN Extension	<p>There are at two DWDM provided Fibre Channel services using separate components and over two fibre circuits. There is at least 5m between the fibres at all times. The storage array determines which path to use.</p>
Inter Campus IP traffic	<p>There are two DWDM provided 1 Gigabit Ethernet services using separate components and over two fibre circuits. There is at least 5m between the fibres at all times. A Routing protocol is used to determine a working path.</p>
SSL offload	<p>This is performed within the ACE blade.</p> <p>When a standby module takes over the functionality of the active module, the existing SSL sessions are lost. New SSL sessions are established on the standby (now active) module using the same configuration available on the active module.</p>
Catalyst 6500	<p>At each HNGX Data Centre there are 2 Catalyst 6513 switches in an Active / Active arrangement within the Distribution Tier. The partial or full failure of one of these results in the other Catalyst being used as described elsewhere in this table – for example LAN Connectivity.</p> <p>Each Catalyst functions as an independent layer 3 switch and other layer 3 devices will select the functioning Catalyst based on interior Routing protocols.</p> <p>Note there are also 2 Catalyst 6500 switches in the Access Tier. The same considerations as stated above apply.</p>

2.1.9.2 Component Targets

The following list summarises the availability targets for different classes / types of network components. The term repaired is defined to mean service restoration as far as all other components are concerned. An example of service restoration would be finding an alternative path around the failed device via Routing protocols.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

- Single Layer 3 component failure in Core of Network repaired within 30 seconds.
- Single Layer 2 component failure in Core of Network repaired within 30 seconds.
- Failure of Active ACE or FWSM repaired within 20 seconds with no TCP/IP connection loss. SSL sessions will need to be re-established as these would be lost.
- Complete failure of Catalyst 6500 repaired within 60 seconds.
- Loss of LAN connectivity (cable / port / interface card) repaired within 10 seconds.
- Failure of Active Outer Firewall repaired within 25 seconds with no TCP/IP connection loss.
- Failure in Branch "High order Network " repaired within 210 seconds (based on standard use of BGP timers)
- Detecting that an ADSL service in a Branch is unusable will take no more than 60 seconds for clean failures and those where the PPP interface comes down based on LCP probes – interval 10 seconds). Switching to a backup service will take place within 5 seconds. Therefore the maximum period for which there is no network path is about 65 seconds. The application will need to explicitly control TCP timeouts to avoid long blocking periods due to exponential back off.

Note that the current application design assumes a timeout in the region of 30 seconds, with one automatic retry.

2.1.9.3 DR

As far as the network is concerned, there is a single Active network and single address space at both the Primary site and Secondary site. It is the case that whilst the Secondary site is supporting testing then;

- The Primary site traffic is mainly Production traffic with some Test traffic
- The Secondary site traffic is mainly test traffic with some Production traffic.
- Separation of the Production and Test traffic classes is maintained through a variety of mechanisms.
- The resources in shared platforms such as Catalyst 6513 available to Test traffic classes will be limited to avoid the very low risk of Test impacting Production.
- Branches use Virtual IP addresses (VIP) to target services. Since the network is Active / Active, each branch will simply select the Data Centre proving the service and detect this based on IP Routing since the VIP is advertised into the Access network. This is how Horizon supports web services today.

Consequences of this approach are that:

- There will be minimal change in the network when DR is invoked. The changes will be limited to closing down Test traffic streams between the Secondary Data Centre and external sources and also ensuring that external traffic sources are directed to the Secondary Data centre. platform and storage

2.1.9.4 HNG-X Wide Area Network

The following BCP/DR testing considerations are drawn from the HNG-X Wide Area Network High Level Design DESNETHLD0009



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

The WAN, both target and interim is designed to have no single point of failure. Network connectivity to IRE11 and IRE19 will operate in an active/active state)

Although each of the new data centres has a single FSNB CE router, resilience through triangulation is provided by the intercampus VLANs over DWDM fibre. Handoff routers for third party connections will be similarly provided as single routers triangulated between sites. The intercampus WAN links and provision of VLANs is documented within the LAN HLD

Network devices are deployed in pairs for resilience (with the exceptions previously mentioned) and will be mounted within separate racks and have separate power feeds from UPS

Devices interfacing with equipment that cannot operate dynamic routing protocols such as firewalls will use VRRP to provide a resilient gateway.

2.1.9.5 Data centre LAN

The following considerations are drawn from the HNG-x Resilience and Disaster Recovery High Level Design (DES/PER/HLD/0001)

The intercampus link is a high speed fibre link between the primary and secondary data centre hosting sites, comprising two redundant, diversely routed fibre links which are DWDM multiplexed to form a number of usable logical links. The DWDM end points are separated by at least 5m at each site. There is a detailed description in [DESNETHLD0007](#).

Over each of the diverse links there will be two 4GB Fibre Channel and two 1GB Ethernet links

The link may be used in a number of ways during normal steady state operation:

- Real-time replication of storage traffic from the primary site SAN to the secondary site SAN
- Network traffic between the two sites, for example, copying of backups to the secondary site for restoring onto test systems

The intercampus link also may be used in a number of failure scenarios:

- In the case where the FSNB link into the primary site fails, it will be possible to route network traffic via the secondary site and then over the intercampus link
- In the event of site failover but where storage is still available at the primary site, the link will be used for replicating SAN traffic in the opposite direction (i.e. from secondary to primary site).

The following BCP/DR testing considerations are drawn from the Data Centre LAN High Level Design [DES/NET/HLD/0008](#)

Section	Design Text
2.1.1 Fibre Services	2.3.5.6(d) Between both HNG-X data centres there are a pair of fibre optic cables. The radial distance of each of these is < 100 km and the two fibres are kept separate along their runs with no common interconnection points.
3.1.5 Data Centres	Under normal operating conditions, IRE11 primarily offers the Production HNG-X service and IRE19 primarily offers the Test service. Production components of the solution are active at both IRE11 and IRE19 simultaneously. The network is therefore considered active at both data centres,



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

	<p>even though IRE19 is primarily offering a test service. This stance is further enhanced by the need for rapid failover from IRE11 to IRE19 in a disaster recovery situation.</p> <p>The direct distance between IRE11 and IRE19 is 6.303 kilometres (3.917 miles). Via the main roads, the distance is 15.8 kilometres (9.82 miles). The radial distance of the fibre services between the two sites is unknown at this time. The separacy of the fibre services is still to be established. The requirement is for a minimum separacy of 10 metres.</p>
3.1.8.5 LAN Topology Bladeframe	There is one production and two multimode BladeFrame in each IRE11 and IRE19 data centre. (see diagram 11)
3.1.8.5.1 Bladeframe Resilience	See details below
3.1.9.5 (Core MSFC)	The Core MSFCs are deployed as a resilient pair. Both MSFCs participate in the OSPF routing protocol and are with other Core MSFC at the other data centre and with ASA firewalls at the local data centre. The core MSFCs at both data centres are attached to OSPF Area 0 and multiple subordinate transit areas for different access layer connectivity. Note that all inter area traffic must pass through the MSFCs to the other access layer areas or the other data centre
3.1.9.6 (ASA Firewalls)	ASA Firewalls are deployed in an Active / Standby configuration. The firewalls share common IP addressing and in the event of failover the interface IP address migrates from the failed server to the standby server. Servers connected to an ASA DMZ will use the ASA DMZ interface address as their default gateway.
3.1.11 Virtual Services	The ACE use Route Health Injection to insert a static VIP address into the MSFC routing table when a back-end service is available. The VIP is redistributed into the IGP by the MSFC. The same or a different VIP may be used at both data centres depending on the requirement. Because of the production / DR status of the data centres, the VIP is only advertised from the Production data centre. In a DR situation, the VIP advertisement is disabled at IRE11 and enabled at IRE19.
3.1.14 Inter D/C Connectivity	<p>There are two fibre paths, between the IRE11 and IRE19 data centres, nominally named North and South. The North and South fibres are diverse as defined in section Error! Reference source not found.. Data services are divided across North and South to provide resilience to a failure of either fibre. The collapsed core and distribution layer switches are connected to North and South as are the Access layer switches. If North or South fails, both the Core and Access layers will continue to operate via the remaining fibre path.</p> <p>North carries 2 * 1Gb Ethernet between Core switches A and Access Switches A. South carries 2 * 1Gb Ethernet between Core switches B and Access Switches B. There are therefore a total of 2 * 1Gb connections carried on the North fibre and 2 * 1Gb connections on the South fibre</p>
3.2.1 Availability & Resilience	<p>The following key principles apply to network resilience:</p> <ul style="list-style-type: none"> No Single Points of failure exist within the data centre network The inter data centre WAN service is resilient with no single points of failure



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

	<ul style="list-style-type: none"> There is a single WAN CE router at each data centre. <ul style="list-style-type: none"> Resilience for high speed WAN connections (>8Mb/s) is achieved by triangulation through both data centres and across the inter data centre WAN. <p>Resilience for low speed WAN connections (<8Mb/s) is achieved by providing resilient services into each data centre. This is to avoid a single failure of such a service preventing a successful DR transition</p>
3.2.1.1 Power	<p>All network devices that support multiple power supplies shall be connected to separate uninterruptible power supply (UPS) power distribution units (PDUs).</p> <p>The DWDM devices will be separate in the data centres (minimum separation 10m) and use separate power supply phases.</p>
3.2.1.2 Physical Proximity	<p>Network devices are provisioned in a resilient pair. The resilient pair may be within a single data centre or spread across both IRE11 and IRE19 data centres in the case of access layer triangulation.</p> <p>Resilient pairs within a data centre are mounted in separate racks</p>
3.2.1.3 Physical Interfaces	<p>Servers have at least two physical interfaces that constitute a single logical interface. The server is connected to either the access layer switches or the core switches. Servers may not be connected to both access layer and core switches as this may bypass network security. One physical interface is connected to one 6513 switch and the other interface to the other 6513 switch. A logical interface is typically associated with a single IP address. A VLAN containing the two physical interfaces is trunked between the switches to provide a single layer two domain. This domain may be served with other servers and network equipment</p>
3.2.1.4 Network Cabling	<p>External network cabling, providing a resilient service, entering the data centre shall have a minimum separation of ten metres (10m). Within the data centre the resilient cables must terminate on separate patch panels and / or devices in separate racks.</p> <p>Network cabling shall conform to the following standards:</p> <p>Copper</p> <ul style="list-style-type: none"> UTP Category 5e² ANSI/EIA/TIA 568B (≤ 100m) <p>Fibre</p> <ul style="list-style-type: none"> multimode 850 nm / 62.5 micron (≤ 220m)
3.2.1.5 Next Hop Redundancy	<p>Routers run the VRRP protocol to provide next-hop redundancy. VRRP is the preferred protocol. Where VRRP is not available, HSRP may be used.</p> <p>ASA firewalls, wherever possible, use a serial failover cable between the two firewalls. Where it is impractical to use a cable (due to distance limitations) a failover VLAN may be used.</p> <p>FWSM firewalls use dedicated VLAN interfaces over which the failover protocol runs</p>
3.2.2 Data centre Failover	<p>IRE19 is normally in Test mode; that is, it primarily hosts the test environments. In the event that DR is declared, all live services must failover from IRE11 to IRE19.</p>

² Note, Category 5e and 6 cables can store high levels of static electricity because of the dielectric properties of the materials used in their construction. The cables must be grounded, to discharge any static electricity present, prior to connection to interface cards.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

	<p>To facilitate the instantiation of the production service in IRE19, the test environments are isolated from the production network to prevent any interference with production activity. The network reconfiguration required to implement the isolation are a mixture of procedural and automated activities. The procedural activities require the Networks operational support team to follow a script to ensure all steps are followed. The automated activities are scripts on the Network Management Systems that are invoked to perform discrete tasks; i.e., shutdown interfaces, disable VLANs, alter routing configuration etc.</p> <p>The intention is to ensure that no services are advertised out of IRE11 and all live services are advertised out of IRE19. Services are never advertised from both data centres concurrently</p>
4 Acceptance and Testing	<p>The network is highly resilient due to the multiple provision of redundant hardware. Every component of redundancy must be considered as part of the acceptance and testing strategy. In principle, double failure is not accounted for. In the case of double failure, reliance is placed upon the maintenance / renewal schedule of the equipment.</p> <p>The following should be considered in component network designs as part of acceptance and testing:</p> <ol style="list-style-type: none"> 1. Physical infrastructure <ol style="list-style-type: none"> a. What components will be lost due to: <ol style="list-style-type: none"> i. Power failure ii. Cable failure iii. Module failure iv. Unit failure b. What component offers redundancy for the failed component? 2. Physical server <ol style="list-style-type: none"> a. Interface redundancy <ol style="list-style-type: none"> i. Does the server failover to use the alternative layer 2 switch when the primary switch fails? ii. Is the layer 3 next hop preserved / available after failover? b. Server redundancy <ol style="list-style-type: none"> i. Does the service hosted on the server failover to the alternative server within the same data centre or to a server at another data centre? 3. Timing <ol style="list-style-type: none"> a. Is failover timely? i.e., does the server / service / network failover within the prescribed limits for the solution component? b. Does spanning tree converge within acceptable limits? c. Does routing converge within an acceptable time?



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

	<p>4. Routing</p> <ul style="list-style-type: none"> a. Are the OSPF DR and BDR router configured properly? <ul style="list-style-type: none"> i. i.e., the DR and BDR routers are pre-determined by use of the OSPF priority setting. b. Is there any load balancing? <ul style="list-style-type: none"> i. Routing is deterministic; load balancing is to be avoided. I.e. the path is pre-determined by use of the OSPF interface cost setting. c. Does the routing converge as expected? d. Is the convergence optimal? <ul style="list-style-type: none"> i. i.e., is dog-legging between data centres avoided? e. Is the routing deterministic? <ul style="list-style-type: none"> i. i.e., is a well defined routing path followed under all failure conditions? ii. Is it possible to pre-determine the routing path to be used under all failure conditions? f. Is there any route leakage? <ul style="list-style-type: none"> i. i.e., routes being incorrectly redistributed ii. route visibility outside of the routing domain – i.e., route leakage from one VPN to another, between 3rd parties etc. <p>5. IP addressing</p> <ul style="list-style-type: none"> a. Is the correct IP address advertised from the correct data centre? <ul style="list-style-type: none"> i. i.e., VIPs, LANs etc. b. Is the IP addressing summarised as expected?
--	--

The BladeFrame network connectivity consists of:

Type	Use
Management port (eth0)	100Mb dedicated interface for access to the PAN Manager software on the cBlades.
pBlade	BladeFrame Processor Blade. Physical component used to provide pServers
pServer	BladeFrame Processing Server. A virtual processing server composed of physical and virtual hardware resources (CPU, memory, disk etc). I.e., consists of a number of pBlades.
©Copyright Fujitsu Services Ltd 2007 Commercial In Confidence Ref: TST/SOT/HTP/0006 Version: 0.2 Date: 19-Oct-07 Page No: 49 of 108	

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE


cBlade	BladeFrame Control Blade. Physical component used to interface IO between the BladeFrame internal network and the external network. The PAN Manager software (used to configure the BladeFrame) runs on the cBlade. Load balancing and fail-over policies are configured on the cBlade. Each cBlade has a 100Mb management interface and eight 1000Mb external network interfaces. Redundant cBlades provide resilience.
vEth	BladeFrame virtual Ethernet interfaces connected to pServers. The PAN Manager software is used to connect vEths to vSwitches.
rEth	BladeFrame redundant Ethernets. Two or more physical NICs from different cBlades providing resilience to failure. A vSwitch rEth is similar to a traditional switch uplink port. Note the uplinks may carry default VLAN or 802.1q encapsulated traffic.
vSwitch	BladeFrame virtual instance of a layer 2 Ethernet switch that spans pBlades and cBlades. Used to connect pServers together in an LPAN, LPANs together and pServers and LPANs to external network equipment. vSwitches may not be connected to other vSwitches. Routing between vSwitches is performed at layer 3 by a dedicated internal pServer or an external router.

Note that vSwitches cannot communicate directly and therefore there is no possibility of layer 2 loops and no requirement for the Spanning Tree Protocol to be enabled on the BladeFrame interfaces of the external switches.

Communication between vSwitches is via an external router.

The default cBlade network interface configuration is "PAIRED". PAIRED specifies that one cBlade NIC in the rEth pair is active and the other cBlade rEth NIC is idle. The idle NIC is only activated if failover is required, for example if the active external switch or cBlade fails.

Figure 1 BladeFrame resilient network architecture



2.1.9.6 Branch Access Network

The following BCP/DR testing considerations are drawn from the Branch Access Network High Level Design [DES/NET/HLD/0014](#)



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

Section	Design Text
4 Design	The network is designed to provide high availability for all areas of the network where branches are aggregated. Single points of failure being restricted to individual branches and their associated local access media (although most branches will benefit from backup access methods leaving the Sarian router as the single point of failure).
4 Design	The two data centres will operate in an active/DR configuration with IRE11 as the normally active site, and all branch traffic will be steered towards IRE11 under normal operation using routing metrics. Although the data centres operate as active/DR, the network will operate in an active/active configuration at all times. Under normal operation, the DR site (IRE19) will be used for testing, and all test traffic will be steered towards this site. Under invocation of DR, all test traffic will be ceased
4.6.2 Branch Router Mgmt	The branch routers are managed by a Sarian system known as ROSS. The ROSS platforms are currently located within the Wigan and Bootle data centres and will be re-provided in IRE11/19. The C&W VPN for ROSS management will be extended to include IRE11 and IRE19. Operation is expected to remain unchanged for HNG-X. Further detail can be found in the Branch Router HLD (DES/NET/HLD/0010).
4.8 Disaster recovery	<p>The Horizon data centres at Wigan and Bootle operate in an active/active configuration. For HNG-X however, the two new data centres in Northern Ireland will operate in an active/DR manner, with IRE11 as the normally active site. IRE19 will be used as a test facility under non-DR conditions. Although applications and services from the data centres will operate as active/DR, the network will operate active/active at all times.</p> <p>Live traffic will be steered towards IRE11 under normal circumstances using BGP attributes. The local CE router (and handoff router where applicable) will be preferred, and the path will be deterministic. Traffic will not be load balanced across parallel paths. Test traffic will be steered towards IRE19 in a similar manner. Support staff will have connectivity to either site.</p> <p>Failure of data centre WAN equipment on the preferred path (local CE and/or local Handoff router) will result in traffic re-routing via the equivalent router in IRE19 and the intercampus LAN. Failover will be dynamic with convergence dependent on the routing protocol in use.</p> <p>Invocation of DR is a manual process that is likely to take around two hours to conclude. Network failover to DR does not need to be dynamic and will use scripting wherever possible to manage the changeover.</p>
6.3 Resilience	<p>The WAN, both target and interim is designed to have no single points of failure. Network connectivity to IRE11 and IRE19 will operate in an active/active state (although the applications and services provided by the data centres may operate as active/DR).</p> <p>Although each of the new data centres in Northern Ireland has a single C&W CE router, resilience through triangulation is provided by intercampus VLANs over DWDM fibre.</p> <p>Network devices are deployed in pairs for resilience (with the exceptions previously mentioned), and will be mounted within separate racks and have separate power feeds from an uninterruptible power supply. Service recovery expectations are:</p> <ul style="list-style-type: none"> • Failure of Layer 3 component where OSPF routing used: < 30 secs



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN



COMMERCIAL IN CONFIDENCE

	<ul style="list-style-type: none">• Failure of Layer 3 component where BGP routing used: < 210 secs• Failure of Layer 2 component: < 30secs <p>Failover performance to alternate media within Sarian router is covered by Branch Router HLD</p>
--	--

UNCONTROLLED IF PRINTED

2.1.10 Platform and storage

There are different classes of server. These are described in the Platforms & Storage Architecture.

- BladeFrame, which is SAN-attached



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

- Discrete Windows 2003 R2 servers Fujitsu-Siemens Primergy platforms (principally RX300)
- Discrete Linux RHEL 4.0 servers on Fujitsu-Siemens Primergy platforms (principally RX300)
- Discrete Solaris 9 Servers on various Fujitsu-Siemens PrimePower platforms (principally PW450)

Discrete Solaris 10 Servers on various Fujitsu-Siemens PrimePower platforms (principally PW250, PW650 or T1000)

A platform is only permitted to be outside BladeFrame if:

- It has some hardware that is not provided by standard BladeFrame e.g. serial cards in Aurora server
- It is not based on Intel architecture e.g. SPARC Solaris
- It requires direct SAN attachment e.g. systems with SYMCLI
- Some other business justification exists

These services will be run with one instance in a discrete server in each data centre. This is undesirable, as these servers are difficult to move from a Test domain to a Production domain, and they tend to multiply rapidly which has a detrimental impact on running costs and system complexity.

These systems are inherently more prone to outage due to component failure, and the "server explosion" may be further exacerbated by the need for a local N+1 resilience model. Even for active/active systems due consideration must be taken of the need for continued resilience and service availability AFTER the loss of a site.

The following is from HNG-X Architecture - Platforms and Storage [ARC/PPS/ARC/0001](#)

Due to the inconsistent use of the term 'Platform' throughout the Post Office Account a new term is introduced to describe the combination of operating system and hardware. The combination of an operating system and hardware will be referred to as 'Platform Foundation'. The moment this generic Platform Foundation is being modified by changing operating system parameters or configurations for the purpose of deploying an application or service to be hosted on the Platform Foundation, it becomes a platform.

2.1.10.1 Platform Strategy for HNG-X

- Publication of a Definitive Hardware List (DHL)
 - New platforms can only be deployed on approved hardware that is listed
 - DHL contains sufficient hardware choices to accommodate all platform requirements present and future
 - DHL will be maintained and refreshed on an ongoing basis
- Publication of a list of approved and supported operating systems for HNG-X
 - Solaris 10
 - Red Hat Enterprise Linux 4 AS
 - Windows 2003 R2
- Abstraction of hardware and OS from platform definition to enable update/refresh of platform foundations independent of application changes

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

The HNG-X platform strategy changes the way a platform is designed. Instead of the Horizon model where the application drives the selection of hardware and operating system, in HNG-X the application owner will have to choose from one of the pre-approved Platform Foundations. This may result in hardware exceeding the minimum requirements for an individual application, but will reduce the overall support costs over time. In the same way the supporting software packages are selected from standard components published in the Definitive Software Library (DSL) wherever possible.

The following diagram illustrates the HNG-X approach.

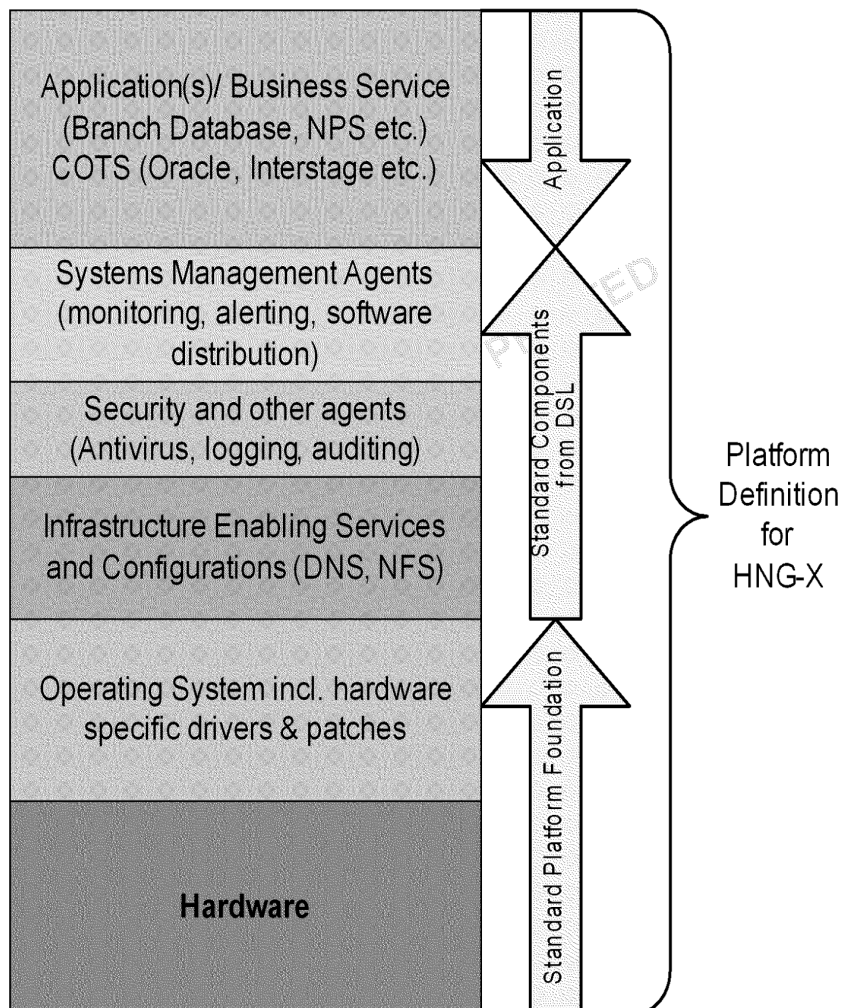


Figure 2 Platform Definitions in HNG-X

2.1.10.2 Hardware Selection for HNG-X

In the context of the platform strategy the hardware has to fulfil the following requirements:

- Provide a reliable and sustainable foundation for all HNG-X business applications, business, infrastructure and systems management services

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

- Enable and support the HNG-X resilience and disaster recovery requirements
- Provide sufficient maintenance and hardware support cover from supplier to fulfil contractual obligations for HNG-X
- Supportable by Fujitsu Services Core Services
- Must fit into the Fujitsu Services supplier strategy

A detailed analysis of the HNG-X disaster recovery service level targets (SLT) in relation to the HNG-X application and data centre architecture has resulted in the following findings which subsequently influenced the hardware choices for HNG-X:

- The Network Banking service will have to be available to the Post Offices Branches within two hours of a disaster being declared.
- The Network Banking service relies on a number of other services and applications to perform its business function, most importantly the Branch Database and the Branch Access Layer (BAL). This provides the ability to settle banking transactions and any other transactions that only require these components.
- To enable Post Office Branches to trade the branches will have to be able to communicate with the data centre which requires that, apart from the physical network, support systems such as Radius servers are available
- The minimum number of servers and support systems required to enable Network Banking to be available to the Post Office Branches within two hours of the declaration of disaster situation is in the region of 35 to 45 servers
- In order to recover the Network Banking service within the required timelines a number of complex and sophisticated applications and systems need to be recovered at the same time, each with its own recovery mechanism

Added to the Network Banking Service Level Targets are the following constraints:

- The data centres are operated in an Active/Disaster Recovery (DR) mode
- While the DR data centre is not used for disaster recovery purposes it will be used for testing
- The data centres are operated remotely and no personnel is available locally in the data centre to aid the disaster recovery process

Based on the aforementioned requirement that the hardware has to enable and support the resilience and disaster recovery requirements for HNG-X, the Bladeframe system was chosen as the platform for the majority of business critical applications and support systems.

2.1.10.3 Active / DR data centres strategy and server placement

The strategy for the HNG-X data centres is to provide a single fully resilient data centre that exclusively handles all business applications and services that form HNG-X. Only in the case of a major disaster will the second, identically configured data centre be used for business application and services. In order to enable the fail over in a disaster situation a number of infrastructure and security services have to be already active in the second data centre prior to invoking the disaster recovery procedures. Examples for such servers are remote access systems, security and authentication systems and infrastructure relevant functions such as storage management. These servers and services are implemented as discrete

**COMMERCIAL IN CONFIDENCE**

servers outside the Bladeframe environment and run in an active configuration in both data centres at all times. The following diagram depicts this configuration. The individual platform designs determine whether a server is exclusive to one data centre or part of the active/active configuration.




Figure 3 Server Placements for Active/Active Services

2.1.10.4 HNG-X Storage Architecture

HNG-X storage comprises multiple physical tiers of storage arrays with different qualities. The data hosted on these storage arrays are grouped into four storage service classes and the classes are commissioned on the appropriate hardware tier. The storage service classes are described in detail in the following sections.

A minimum of two storage arrays for live data is deployed in each data centre alongside the existing Horizon Archiving arrays on EMC Centera. The enterprise class storage array is used to provision data in a highly protected and available way, including synchronous replication to a remote data centre inline with storage services class 1. It is also used to provision other service classes as long as the capacity is sufficient and the midrange system cannot provide a more cost-effective solution. The midrange system hosts service classes 2 - 4 (defined in the next section) where it provides a cost-effective alternative to using the enterprise class storage system. Each array has an identical equivalent in the second data centre and replicates data to the second system either synchronously or asynchronously depending on the application needs.

Although the storage arrays do not contain a single point of failure in themselves, the storage solution as a whole is a single point of failure. Due to the extremely high cost of duplicating the storage arrays in each data centre, this has been accepted as a risk and is in line with the systems qualities architecture for HNG-X. The data is protected against a catastrophic data centre failure by replicating the contents of the storage arrays, either synchronously and/or asynchronously across to identical storage arrays in the second data centre.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

2.1.10.4.1 Data Classification

The data classification for HNG-X focuses on the Recovery Time and Recovery Point Objectives (RTO & RPO) for the individual application data. All data hosted on Horizon was analysed and the owners of the data were interviewed to determine which data need to be hosted by the HNG-X storage solution during the migration from Horizon to HNG-X. An attempt was made to classify all data to be stored as part of the HNG-X solution and the findings are as follows:

- The negotiated service level targets for HNG-X are very detailed for some business services and not very exact for others
- Two main systems and a number of smaller systems have an RPO of zero, translating into a requirement to ensure that no data will ever be lost
- Most data has an RPO greater than zero but it is very difficult to determine the exact value
- RTO requirements vary widely between the application data, from minutes to days

Four main groups of data with similar characteristics have emerged and were combined into the following classes of storage service:

1. Storage Service Class 1

Storage Service Class 1 supports critical applications and databases that need high performance and replication (RPO = 0, RTO ≈ 0).

This service class will replace the existing Horizon Symmetrix storage arrays. Some of the business functionality will move to service levels 2 and 3.

On this storage service level class will reside the most business critical data for the Post Office. This data is captured in the new "Branch DB" database. This tier will also host storage for a few additional Oracle databases supporting "Online Systems" as well as the Oracle data guard standby databases for the Branch Database,

2. Storage Service Class 2

Storage Service Class 2 supports critical applications and databases that need high performance and/or replication, but have extended recovery objectives (RPO=0, RTO>24hr).

This service class will replace the existing Horizon Symmetrix storage arrays insofar as some of the business functionality that has extended recovery time objectives will move to Service Class 2.

On this service class will reside data which is critical data for the Post Office but which can be recovered with less demanding requirements and/or is performance sensitive. This class supports the production SAP system for the Post Office (POL-FS), and Data from the Horizon Riposte platform (correspondence servers) during the migration phase into HNG-X. Both SAP and Riposte data are performance sensitive. This class does need to host OS boot volumes.

3. Storage Service Class 3

Storage Service Class 3 supports other production databases, Quality Assurance (QA), test – asynchronous replication or no replication (RPO > 0, RTO > 24hr).

This service class will support the remaining HNG-X production databases and other data from the Horizon platform during the migration phase into HNG-X. This service class will provide capacity for recovery databases, i.e.: capacity for QA and test systems for the Post Office.

Replication of the data is to be configurable (synchronous or asynchronous) allowing to cater for both low RPO as in the case of Batch systems, and to cater simply for the purpose of ease of recovery in case of site DR fail-over and fail back.

4. Storage Service Class 4



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

Storage Service Class 4 supports Support databases, near line file system storage – Asynchronous replication or no replication (RPO > 24hr, RTO > 48hr).

This service class hosts the remaining data supporting the HNG-X platforms.

Replication of the data is to be asynchronous allowing to cater for ease of recovery in case of site DR fail-over and fail-back.

5. Storage Service Class 5

Storage Service Class 5 supports Data with long term, regulated retention. Regulatory reports, SOX Records, email, etc. – may require replication.

6. Storage Service Class 6

Storage Service Class 6 supports Backup and Restores / Synthetic Full Backups – VTL with replication capabilities

7. Storage Service Class 7

Storage Service Class 7 supports Tape for off-site storage

2.1.10.5 Recovery Time Objectives by Storage Service Class

The required recovery times in the four relevant storage service classes can be subdivided into local recovery requirements in the event of an application failure and remote recovery in the event of a data centre disaster. The local and remote recovery objectives for the different classes are listed below:

2.1.10.5.1 Storage Service Class 1

Branch DB

1. Local Recovery in case of database corruption (and corruption replicated to remote DC)
 - a. Recovery Time Objective (RTO) = 1 hour from local split mirror backup
 - b. Recovery Point Objective (RPO) = time of last split mirror backup (maximum 24 hours)

(Upon declaration of Branch DB data corruption, the Branch DB functionality will be provided by an Oracle standby database. RPO = up to 2x lag between production and standby database, RTO = 30 minutes. This acknowledges that recovering from database corruption can be a lengthy and complex process which is not storage dependant)

2. Disaster Recovery in remote data centre from point of disaster declaration
 - a. RTO = 1 hour
 - b. RPO = No data loss

Online System Databases

1. Local Recovery in case of database corruption (and corruption replicated to remote DC)
 - a. Recovery Time Objective (RTO) = 2 hours from local split mirror backup
 - b. Recovery Point Objective (RPO) = time of last split mirror backup (maximum 24 hours)
2. Disaster Recovery in remote data centre from point of disaster declaration
 - a. RTO = 2 hours
 - b. RPO = No data loss



COMMERCIAL IN CONFIDENCE

Standby Branch Database and the Standby Online Systems Databases

1. Local Recovery in case of database corruption (and corruption replicated to remote DC)
 - a. Recovery Time Objective (RTO) = 1 hour from local split mirror backup
 - b. Recovery Point Objective (RPO) = time of last split mirror backup (maximum 24 hours)
2. Disaster Recovery in remote data centre from point of disaster declaration
 - a. RTO = 1 hour
 - b. RPO = No data loss

2.1.10.5.2 Storage Service Class 2

POL-FS

1. Local Recovery in case of database corruption (and corruption replicated to remote DC)
 - a. RTO = 24 hours from local split mirror backup
 - b. RPO = time of last split mirror backup (maximum 24 hours)
2. Disaster Recovery in remote data centre from point of disaster declaration
 - a. RTO = 24 hours
 - b. RPO = No data loss

Horizon Riposte

1. Local Recovery in case of database corruption (and corruption replicated to remote DC)
 - a. Recovery Time Objective (RTO) = 2 hours from local split mirror backup
 - b. Recovery Point Objective (RPO) = time of last split mirror backup (maximum 24 hours)
2. Disaster Recovery in remote data centre from point of disaster declaration
 - a. RTO = 0 hours – Application cluster is active on both data centres
 - b. RPO = 0 – Application based host replication

2.1.10.5.3 Storage Service Class 3

Batch Processing Solutions

1. Local Recovery in case of database corruption (and corruption replicated to remote DC)
 - a. Recovery Time Objective (RTO) = 2 hours from local split mirror backup
 - b. Recovery Point Objective (RPO) = time of last split mirror backup (maximum 24 hours)
2. Disaster Recovery in remote data centre from point of disaster declaration
 - a. RTO = Hours
 - b. RPO = No data loss for some

SAP POL FS Test and QA

1. Local Recovery in case of database corruption (and corruption replicated to remote DC)
 - a. RTO = 24 hours from local split mirror backup
 - b. RPO = time of last split mirror backup (maximum 24 hours)

**COMMERCIAL IN CONFIDENCE**

2. Disaster Recovery in remote data centre from point of disaster declaration
 - a. RTO = none
 - b. RPO = none

2.1.10.5.4 Storage Service Class 4

Recovery Objectives for this service class is as follows:

1. Local Recovery in case of database corruption (and corruption replicated to remote DC)
 - a. RTO = > 48 hours
 - b. RPO = time of last split mirror backup (maximum 24 hours)
2. Disaster Recovery in remote data centre from point of disaster declaration
 - a. RTO > 48 hours
 - b. RPO = TBD

2.1.10.6 Storage Arrays

The previously described four storage service classes impose different requirements on the storage hardware, and the resilience and availability features that the hardware has to provide. It can broadly be broken down into two main sets of qualities

Enterprise Class Storage

Data from Service Classes 1, 2 and 3 will be hosted on enterprise class storage arrays, capable of delivering both an extreme high level of data availability and performance. The storage array itself (at each data centre) will contain no single point of failure. The array performance must be scalable with the addition of hardware elements (controller boards, host connection ports, remote replication ports, cache, etc.) should this prove to be necessary.

The following are characteristics of the storage systems:

- Each discrete component of the storage array can be exchanged without taking the storage array offline. This (replacement) operation will have no noticeable impact on performance.
- Each sub-component / FRU is monitored for fault conditions
- Major sub-components / FRUs are monitored for pre-fault conditions
- Multiple fault monitoring sub-systems per FRU
- Disk shelf implements electrical drive isolation / bypass at a hardware level to isolate electrical faults in drives
- Disk shelf implements FC port bypass at a hardware level to isolate faulty drives
- Continuous integrity checking of data on drives / end-to-end data integrity checking
- Ability to seamlessly move data within storage array to compensate for disk contention "hot spots"
- Consistency technology
- Identical storage arrays to be deployed at both sites (Active – Passive configuration for HNG-x)
- Synchronous remote replication

**COMMERCIAL IN CONFIDENCE**

- Remote copy protected with same drive type and RAID level as production source
- Local copy clones protected with RAID protection (RAID 5 expected)
- Ability to take local split mirror backups or equivalent copies with no performance impact

Midrange Class Storage

Data that has less demanding requirements and does not make use of all the features offered by the enterprise class storage systems will be hosted on midrange class storage arrays, capable of delivering data in a secure and available manner. The storage solution will contain no single point of failure.

The following are mandatory characteristics of the storage systems:

- Major sub-components / FRUs are monitored for pre-fault conditions
- Disk shelf implements FC port bypass at a hardware level to isolate faulty drives
- Identical storage arrays to be deployed at both sites (Active – Passive configuration)
- Asynchronous and synchronous remote replication
- Local copy clones or snapshot area protected with RAID protection (RAID 5 expected)

The service classes are provisioned out of either physical platform. Where only one platform type can meet the requirements of the service class, this platform type is the only possible way to provide this service class (e.g. storage class 1). The enterprise class storage array is used to provision service classes 1, 2 and 3 to make use of the available capacity.

2.1.10.7 Recovery and resilience

The HNG-X solution mandates that no data centre shall contain a single point of failure. The only exceptions to this rule in the platform and storage domain are the storage arrays, due to the exorbitant costs of duplicating the storage environment locally. The storage arrays are fully resilient in themselves and do not contain single points of failure.

Discrete servers outside the Bladeframe have either a stand-by server in case the hardware fails or are clustered where high availability is required. The only exceptions are the Audit servers; the design for the Audit servers does not require a data centre local resilience facility other than re-provisioning the Audit server in case of a hardware failure. Where resilience is inherent in the application through the deployment of multiple instances (e.g. Active Directory Domain Controllers or DNS servers) a pool of stand-by servers is available to rebuild or recover the failed server. All servers, services or infrastructure functions that are located in the outer ring in diagram and are necessary to access the data centre and to invoke the disaster recovery procedures are run in an active/active configuration across the two data centres with multiple instances of each server in both locations. In the event of a catastrophic loss of one data centre at least two active instances of these servers remain in the second data centre.

Servers inside the Bladeframe environment are protected against hardware failure through a pool of stand-by servers. Should an individual blade server fail it will automatically be restarted on a spare blade either inside the same Bladeframe or in a different Bladeframe.

The storage arrays are configured to replicate their data according to the storage service classes and the application requirements in regular intervals or constantly via a synchronous link. Should a storage array irrecoverably fail, a site disaster will be declared and the entire data centre function will be moved to the second data centre, and the servers and services will be restarted using the replicated data.

In the event that a disaster has been declared the active components housed in the outer ring as described above will be used to access the second data centre and invoke the DR procedures. This will be co-ordinated with the necessary changes to the network. After test access has been disabled by the network the Bladeframe systems and the necessary discrete servers will be re-configured to use the boot images and the application data that were replicated to the secondary data centre during normal

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

operation. The reconfiguration is largely automated using specially developed scripts that will prepare the storage arrays and then attach the servers to the replicated data. The entire process is designed to take no longer than 60 minutes to allow for application and service, testing and verification prior to enabling branch trading.

UNCONTROLLED IF PRINTED

2.2 Other Area's

2.2.1 Storage Area Network (SAN)

To enable successful site failover to occur, both sites need to be reasonably consistent. That is to say, the amount of data loss due to a site failover needs to be minimised. The major area where this may occur is in the SAN replication.

Processes need to be in place to ensure that any change that happens on the primary site also occurs on the secondary site. The "lag" must be carefully monitored, and alerts should be raised in case the lag grows beyond an acceptable threshold. This is generally an alert which is raised by the storage system itself.

2.2.1.1 Storage

A number of classes of storage are proposed in ARC/PPS/ARC/0001. Applications which require zero data loss on failover must specify suitable storage, which in this case is EMC DMX with synchronous SRDF enabled. Applications which can recover following a failover where some amount of data has been



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

lost, either because the loss of the data is not significant or because the recovery may be effected from journals or upstream systems, may use a lower storage class. Add ref to storage HLD

Storage presentation is very important. LUNs that are replicated between sites need to have the same SCSI IDs so that the BladeFrame at the failover site can recognise the storage without need for reconfiguration.

In addition, to allow the workload of any pBlade to be run from any BladeFrame, each LUN that is required by any pServer will be presented to the cBlades in every BladeFrame.

2.2.2 Network

Servers will appear on the secondary site with the same IP as they had on the primary site.

2.2.3 All VLANs that are required by any pServer will be trunked to all of the BladeFrames

The following details are taken from the HNG-X Storage Area Network (SAN) HLD DES/NET/HLD/0007:

Section 4.1	There will be two SAN Directors deployed at each HNG-X data centre. Both power and FC cabling will be fully resilient
Section 4.1.2.4	Power Supplies : If a power supply fails, a single power supply is sufficient to power the entire system
Section 4.1.2	The supervisor module is a hot-swappable module. In a dual supervisor module system this allows the module to be removed and replaced without causing disruption to the rest of the system
Section 4.1.2.3	MDS 9509 Crossbar Fabrics : The system will not experience any disruption or any loss of performance with the removal or failure of one supervisor module
Section 4.1.2.6	MDS 9500 Supervisor 2 Module: Provides fully redundant operation. Each chassis has 2 supervisor modules for resilience. The control engine operates in active/standby mode
Section 4.1.2.1	In the event of a crossbar failure, the standby channel on the remaining crossbar becomes active, resulting in an identical amount of active crossbar switching capacity
Section 4.1.2.2	FSPF-based multipathing : In the event of a switch failure, dynamically reroutes traffic
Section 4.1.2.4	System Fans: The tray is designed with 1:1 redundancy. The system can sustain a multi-fan failure with no negative effect. Up to 4 fans can fail before the systems are affected. The entire fan tray is hot-swappable. The system can run for up to 30

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

	minutes without a fan tray installed
Section 4.1.4	Virtual SANS : The second level of isolation (and resilience) will be achieved by having a second set of connectivity devices and VSANs at each site
Section 4.1.4.1	The isolation of tape storage network from disk storage network eliminates an application-level impact due to failures such as power rest of a tape library unit
Section 4.2.1	Wavestream (IRE11-IRE19) : All channels are protected using dual fibre routes for resilience
Section 12.2	Testing Connectivity – Test 1 (disable SAN ports), Test 2 (Pulling FC cables), Test 3 (Introduce transient FC errors using ANUE equipment), Test 4 (increasing latency using ANUE equipment)

2.2.4 System Qualities

The HNG-X System Qualities Architecture ([ARC/PER/ARC/0001](#)) document defines availability and recovery times. These should form the basis of testing in the HLTP

The HNG-X system satisfies or provides solutions for the 3 main aspects of system qualities as follows:

- Capacity/Performance
- Availability. The system has been designed to be highly resilient, having Single Points Of Failure avoided wherever possible for all components within the solution, using standard industry practices. Where it is not cost effective to avoid Single Points of Failure, these will be itemised. In order to protect against data corruption, persistent data will be backed up, in accordance with the requirements for the system.
- Disaster Recovery. The systems at the secondary data centre will provide full functionality (capacity, performance, resilience and backup) so that should DR be invoked and it is required to move the entire service to the Secondary data centre, the same service provided by the Primary data centre (in terms of capacity, performance and local resilience of the systems) would be provided at the Secondary data centre. Additionally it is required that systems crucial to POL operation should have no data loss; therefore for the DR solution, storage for these systems will be synchronously replicated to the Secondary data centre so that any transaction committed will



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

also be persisted on the storage both at the Primary and Secondary data centres. Hence under normal operation for these systems, the DR storage which will always hold an exact copy of the storage at the Primary data centre.

2.2.4.1 Availability

The table below states the SLTs for overall availability. It shows both core hours, and this as a percentage (e.g. the 9s notation). The SLT applies to a five year rolling average.

Service Level	Max downtime core hours per year	Percentage availability
Outages in Core Hours where the Core Solution (Central & Branch Network, Core Infrastructure and Branch Database) is unavailable at > 10% of Branches per SLT year	3	99.89568%
Outages in Core Hours where the Banking Solution (CAPO, A&L, Link) is unavailable at > 10% of Branches per SLT year. This includes time when the Banking Solution is unavailable because the Core Solution is unavailable.	8	99.72181%
Outages in Core Hours where Other Services (ETU, DVLA, PAF, APOP, DCS) are unavailable at > 10% of Branches per SLT year. This includes time when the Other Services are unavailable because the Core Solution is unavailable.	14	99.51316%

Table 1 SLTs for overall availability

In the case of a service outage where the SLT is currently failing (not for unavoidable disasters) and Fujitsu Services recommends that the service is switched to DR operation whereby the entire service be moved to the secondary data centre, the measurement of the above SLTs stops until Post Office Ltd give official notification to move the service, after which the measurement continues in parallel with the DR SLTs

Availability for individual branches and counters is stated in the table below, again both shown as core hours and a percentage. Availability targets increase from March 2009.

Table 2 Branches and counter availability

Service Level	Max downtime core hours per year	Percentage availability
Branch availability during Core Hours until March 2009 per SLT year	18.1	> 99.37%
Counter availability during Core Hours until March 2009 per SLT year	25.8	> 99.10%
Branch availability during Core Hours from March 2009 per SLT year	15.0	> 99.48%
Counter availability during Core Hours from March 2009 per SLT year	23.3	> 99.19%

These values represent the maximum time per year that any counter or branch is unable to trade.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

2.2.4.2 Availability Principles

2.2.4.2.1 No data loss

Paragraph 5.1 of RM/CDE/0031 v 0.4 states that the current recovery models are assumed.

The recovery model for Horizon assumes that unrecoverable data loss is unacceptable for essential services. HNG-X will assume the same.

2.2.4.2.2 Fail-over for resilience within the data centre

Requirements SCD-41 and [ARC-445](#) state that automatic fail-over shall be used for defined services within the Data Centre, and for the network connections within the Data Centre.

Components provided for resilience (fail-over) can be incorporated into the normal service, to provide load balancing and improved performance. In this case, there should be sufficient spare capacity within a single data centre to continue to provide the service if one component fails.

2.2.4.2.3 Single point of failure

Requirements SCD-40 and [ARC-444](#) state that there shall be no single points of failure (SPOFs) that can cause the loss of any Business Capabilities or Support Facilities.

Requirement SER-2155 states that SRRCs will be updated to reflect new Architecture. SCD-39 states SRRCs will define the priority assigned to incident depending on the business impact and contingency for all components in the infrastructure.

There are redundant components within most of the data centre infrastructure which prevent any SPOFs in the services that support the main branch business. However Service Resilience and Recovery Catalogues (SRRCs) will be created for main components within the system, indicating risks, repairs, contingency and resolution in the event of potential documented failures to or within a component.

Data centre components that are not duplicated and represent single points of failure within the HNG-X solution are itemised as indicated in the table below:

Single Point of Failure	Reason	Mitigation
Primary Storage	Cost	Highly resilient internally
POL-FS Main Host	Cost	Highly resilient internally
Blade Frame Cabinet	Architecture	Highly resilient internally
Broadband Access Servers	No Alternative	Backup Network
Wide Area Network to Data Centres	Cost	Triangulation between data centres provide resilience
Radius Servers	Solution Complexity	Triangulation between data centres provide resilience
Branch	Cost	Depends on number of counters (see below)

Table single points of failure within the HNG-X

Although the Primary Storage (for the Horizon solution an EMC Symmetric disk array) is a SPOF, it is highly resilient internally and therefore assumed to meet the requirements. In the extremely unlikely event that the entirety of one of these components fails, then the service will be recreated at the secondary data centre. However as there are 2 Primary Storage Arrays at each data centre it would be



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

possible to provide a service using the Branch Standby Database which is located on different Primary Storage to the Branch Database, therefore providing extra resilience. This may prove to be sufficient until possible to repair the failure.

The main POL-FS host runs on a highly resilient server, which is not replicated due to cost reasons (additionally the less important IXOS server is also a SPOF). In the unlikely event that this fails and it is not possible to repair/replace within adequate timescales, then the entire POL-FS service will be recreated at the secondary data centre if deemed necessary.

It is not possible to completely remove SPOFs at each counter and branch. However under HNG-X, problems with an individual counter or branch are isolated and minimised as much as possible (e.g. using the branch router and for larger offices hubs).

2.2.4.2.4 Recovery times

Requirement [ARC-446](#) states that the impact on Branch Users shall be minimised if there is a failure and subsequent recovery. Principles for this are described in the document *Agreed Assumptions on HNG-X Branch Exception Handling* referenced from Schedule B6/1.

In most cases, HNG-X uses the same approach as the existing Horizon system. If any transaction times out, the service should be available again by the time the transaction can be retried by staff at a counter.

To achieve this, all services running on central systems at the data centre that are critical to branch operation should recover within two minutes wherever possible (this would not be possible for a serious corruption requiring the restore of a backup).

Services that are not critical to branch operation, such as anti-virus services, should be available within a reasonable period. These are described in section 3.2.3 within [ARC/PER/ARC/0001](#).

2.2.4.3 Component availability

The table below states the maximum target time in which each data centre service should recover from transient faults in order to meet SLTs

[Table in both the System Qualities Architecture paper and the HNG-X Resilience and Disaster Recovery HLD [DES/PER/HLD/0001](#)

	Maximum Target Recovery Time
Business Systems	
Branch Database	2 minutes
Client File Transfer (DCS, ETU, Banking)	2 hours
DCS & ETU online	2 minutes
FTMS TIP Local & Track and Trace	2 hours
NBX Banking Agents	2 minutes
DVLA online, PAF & APOP Agents	2 minutes
Branch Access Layer	2 minutes
TES (application servers)	2 hours
NPS	2 minutes
APOP	2 hours
Main Host	2 hours
Storage Systems	
Audit Centera Array	Next day
Audit Server	2 hours
Backup Servers	2 hours
ECC Server (or equivalent)	Next day
Main Backup System (Disk or Tape)	2 hours



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

Support Systems	
Secondary Authentication Service	15 minutes
Antivirus Server	Next day
Application Monitoring Server	15 minutes
Certification Server	15 minutes
DNS Server	15 minutes
Domain Controllers	15 minutes
NBX Network Observer Server	Next day
NBX Network Probe Server	Next day
Network Alarm Point Server	15 minutes
Network CISCO Works Server	15 minutes
Provisioning Server	Next day
Radius Servers	2 minutes
Accounting Radius Servers	2 hours
SAS Server	15 minutes
Signing Server	2 hours
SQL Server (ACDB, OCMS, Athene, MTAS)	2 hours
SSC Branch Database	2 hours
SYSMAN Enterprise Managing Server	2 hours
SYSMAN Enterprise Monitoring Server	2 hours
SYSMAN Enterprise Event Servers	2 hours
SYSMAN Availability Server	2 hours
SYSMAN Enterprise User Interface Server	2 hours
SYSMAN Enterprise Database Server	2 hours
SYSMAN Enterprise Provisioning Server	2 hours
SYSMAN Enterprise Fanout Server	2 hours
SYSMAN Enterprise Staging Servers	2 hours
SYSMAN Enterprise Legacy manager	2 hours
SYSMAN Enterprise Monitoring Display	2 hours
.Network switches	
.Network routers	
.Cable failures / mis-cabling	

Table 3 Maximum Target Recovery Time

2.2.4.4 Disaster recovery

In the event of a catastrophic failure within or at the entire primary data centre, that prevents full (or possibly partial) operation; Fujitsu Services will recommend that Post Office Ltd make an official request to switch over to DR operation at the Secondary data centre. When this official request has been made by Post Office Ltd, the services should be restored within the following timescales (currently subject to agreement with Post Office Ltd):

Service Description	DR availability target from official notification
Core Solution and Network Banking, including: <ul style="list-style-type: none"> Branch Database Servers Client File Transfer (DCS, ETU, Banking) DCS Servers FTMS TIP Local & Track and Trace NBX Banking Agents Branch Access Layer servers NPS Servers (Database) 	2 hours
All remaining Services excluding POL-FS:	5 hours



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

<ul style="list-style-type: none"> ETU & DVLA online Servers PAF & APOP Agent Servers TES Application Servers APOP Servers (Database) Main Host (Batch Database Server) <p>(Priority would be given to any services crucial at time of DR e.g. DVLA if at end/beginning of month. PAF if at Christmas mailing peak period)</p>	
POL-FS	48 hours

Table 4 Service availability at DR

Whilst Business Continuity testing might result in times lower than those defined in the table above for POL-FS and the Remaining Services group (as is sometimes the case on the existing Horizon solution), a single BC test (e.g. POL-FS) does not take into account that the operations staff will already be busy working on resuming other services. Also the times defined are the maximum acceptable times for DR service resumption and the aim, where possible, will be to provide service at the DR data centre in as short a time as possible.

The SLT times however reflect the reality of the situation, that unlike a scheduled business continuity test, a genuine DR request will be unexpected and staff will not be fully prepared for it. However processes will be in place to ensure that DR will occur as planned even if key individuals are on sick or on leave. This process will be rehearsed during the business continuity tests.

2.2.4.5 Solution Overview - Availability

System availability within HNG-X falls into 2 categories.

- Systems crucial to the operation of the branch. These will require fast automatic failover.
- Systems not crucial to the operation of the branch. These will typically allow time for a service to be restarted without affecting the ability of the branch to trade.

The table below shows the recovery mechanisms for essential Business Systems required to enable branch trading (where recovery times of 2 minutes or less are required).

System	Local Recovery Mechanism	Connecting Systems Failover Mechanism
Branch Database servers	Primary - Oracle RAC (cluster) Secondary – Oracle Data Guard (replication) Branch Standby Database	BAL is connected to all Servers; on failure detection an alternative connection is used
DCS, ETU online Servers	Platform (Active/Standby Configuration) Software Heartbeat mechanism between them written to NPS	Active Server broadcasts message on start up
NBX Banking Agents servers	Platform (Active/Standby Configuration) Software Heartbeat mechanism between them written to NPS	Active Server broadcasts message on start-up
DVLA online, PAF, APOP & Help Desk Agent servers	Platform (N+1 Configuration)	Network device checks service and no longer uses a failed system
Branch Access Layer servers	Platform (N+1 Configuration)	Network device checks service and no longer uses a failed system
NPS database server	Oracle RAC (Cluster)	Agents are connected to all Servers; on failure detection

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE

		an alternative connection is used
APOP database server	Oracle RAC	Agents are connected to all Servers; on failure detection an alternative connection is used
Main Host database server	Platform (Active/Standby Configuration)	For current active standby configuration, use identity of Active by VIP

UNCONTROLLED IF PRINTED

**Table 5 Recovery mechanisms for essential Business Systems**

The System Qualities Architecture paper ([ARC/PER/ARC/0001](#)) provides more details about the Resilience, Failover and Recovery for all the main elements of the HNG-X Architecture. Reference should be made to this document in conjunction with the High level Designs when designing the High level Test Plan and detailed tests.

2.1.5 Estate Management

The following details are taken from the HNG-X Estate Management Component Architecture ([ARCSYMARC0005](#))

The key roles for Estate Management include:

- Provide storage for permitted Branch Configuration Data values
- Maintain the target OBC state for the Branch. (Noting that this may differ from live in that other systems and staff may be responsible for achieving the desired state).
- Provide relevant configuration data for EM endpoints within the Data Centre
- Manage Operational Business Change
- Provide a delivery capability to external suppliers
- Support the provisioning of Branch Router and Counters for new and replacement units (a.k.a. AutoConfig)



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

3 Recovery and resilience

The following EM platforms to be considered:

- EMDB Server
- Boot Platform
- BCMS [DN: Out of scope currently.]

The table below summarises the recovery and resilience requirements for each platform:

Platform	Application	Recovery (non-DR) Time Objective	Recovery (DR) Time Objective	D/B Clustering required?	Data to be Backed Up	Comments
EMDB Server	SQL Server 2005	See Note * a.1.	See Note * a.1	No	Yes (database + logs)	Single blade with SAN (no VIP necessary). Real Time access by EM configuration data consumers (RADIUS etc) – but endpoint must provide resilience for EMDB being unavailable. EMDB must poll data providers (and be capable of handling their unavailability).
Boot Platform	Radiator, RCAP, Router Bootserver, Counter BootServer	24 hrs See Note * a.2	24 hrs See Note * a.2	No	Yes (database, transaction logs, RCFs and BSFs)	Real time access but for Counters and Routers that are being provisioned (small number and could cope with 2 hour outage). 2 Blades with SAN.

Notes:

- Each endpoint will be able to function with EMDB unavailable. The EMDB RTO values will be the minimum length of time determined by the minimum value of:
 - the minimum time required by the set of all EMDB endpoints can operate without EMDB,
 - the minimum time required for final changes to be entered into EMDB by OBC Team is currently 17:00 on day-2 (which would imply that for a change on day-0, gives approximately 36hrs).
 - the maximum time for Branch Router, HNG-X and XP OS Migration changes can be delayed without affecting installation of replacement units. This is likely to be the most critical.
- The Boot platform as a replicated system will have similar recovery time to that of the Horizon BootLoader (i.e. 24hrs based on current Business Continuity plans).

**COMMERCIAL IN CONFIDENCE**

3.1 Hardware Resilience

The EM platforms: BCMAS, EMDb and the Boot Platform are provisioned as individual Blades within a Blade Frame Server. The Blade Frame technology effectively gives each server platform a fully resilient virtual LAN connection to their respective Network Security Domains in the Campus Network.

The Blade's storage is supplied via a common SAN that is configured with distinct areas for Operating Systems and Paging, Binaries and Data.

In the event of a single platform Blade failure, there are four likely failover scenarios.

- Re-provision the failing platform on a spare (v)Blade in the Blade Frame.

- Re-provision the failing platform on a spare (v)Blade in another Blade Frame.

- Restore Image on new (v)Blade in the Blade Frame

- Restore Image on new (v)Blade in another Blade Frame

In both of the re-provision cases above, the entire platform's storage is immediately accessible via the SAN.

[DC: A roll back of the database and a replay of transaction logs may be necessary]

3.2 Database Backups

EMDb requires regular backups and the ability to recover back to a position 5 minutes prior to a failure. The 5 minute value may yet prove to be too long due to potential loss of input from external systems.

3.3 EMDb Endpoint Resilience Requirement

Each EMDb endpoint is responsible for its own resilience to the loss of its interface with EMDb should it ever fail or be unavailable.

3.3.1 Time Synchronisation

3.3.2 The following details is taken from the High Level Design for Time Synchronisation at HNG-X ([DES/NET/HLD/0013](#))

3.3.3 NTP Servers

In the event of individual failure the stratum 1 and Active Directory client devices will be configured to use:

- Either lan port on the device.

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

- Either NTP server at the production site
- Either NTP Server at the DR site.

In the event of a collective failure of the stratum 0 devices the stratum 0 devices will have their names aliased to DNS entries for the Estate Time Servers on all stratum 1 clients, the Estate time server will have its TTL entry in the primary DNS Server set to one hour to ensure that within around 2 hours or so the change in address will have been propagated throughout the system.

3.3.1.1 Reliability

The NTP protocol explicitly addresses resilience and reliability issues and can generally ensure that time rarely exceeds 0.0128 seconds from the stratum 0 source, this is regarded as an acceptable upper margin; there is redundancy in the system such that if any time servers should fail it will be replaced operationally by another. The Windows based time service is a proprietary solution with characteristics similar to the NTP standard and it is believed to be equally reliable.

Broadly speaking all strata will follow the following guidelines during normal operation where Δ is the observed time difference between the time source and local clock.

Deviation	$ \Delta > 5$ seconds	$5 > \Delta > 1$ seconds	$ \Delta < 1$ second
Action	Report exception	Step change to source time, log action	Drift towards expected time

When platforms start up there will be an introductory period where Windows machines will be manually synchronised to the required tolerance, a function that Unix Machines will automatically perform (N.B. It will be the platform owner's responsibility to ensure that any significant negative time changes are preceded by appropriate housekeeping measures).

3.3.1.1.1 Stratum 1 & 2

To make a step change if the time is out by more than 1 second and less than 5 seconds time.

3.3.1.1.2 Stratum 3+

Not to report any events but drift towards the provided time; the time discrepancy will be reported by the BAL.

3.3.1.2 Recovery

Support documentation and Kels will be provided to describe the necessary recovery steps should any of the above occur.

3.4 Resilience

The system will be an active-active configuration at stratum 0 with two time servers at each site.

**COMMERCIAL IN CONFIDENCE**

Where a step change of more than 5 seconds occurs the Windows Service will stop synchronising and the Unix daemons will terminate without synchronising, the consequential events will be recorded and reported through Tivoli.

3.5 Business Application Dependences

Business applications should not be allowed to start unless the time service is running. The time service will not run if it is more than 5 minutes away from UTC and it can see a time server.

3.5.1

UNCONTROLLED IF PRINTED

3.5.2 Active Directory

Directory services are provided by Windows Active Directory. AD will be deployed in accordance with best practices for security and resilience to provide a continuous service in the event of disaster. This is described in DES/PPS/HLD/0003. Interface modules are provided to allow Solaris and Linux systems to interact with AD.

3.5.3 Secondary authentication is integrated with AD and managed transparently for applications. This is described in [DES/SEC/HLD/0001](#)

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

The following details is taken from the Active Directory HLD for HNG-X ([DES/PPS/HLD/0003](#))

To provide resiliency there will be 2 domain controllers per data centre, at each site there will be one domain controller hosted within the Fujitsu Siemens Bladeframe and one hosted on a discrete server.

The Active Directory will be hosted on a Bladeframe pBlade which provides the highest server availability model, the Bladeframe infrastructure provides the following resiliency (see Table below)

Additional domain controllers will sit outside the Bladeframe to ensure service is maintained in case of connectivity failure with the Bladeframe

High level disaster recovery and resiliency information is covered in the HNG-X Resilience and Disaster Recovery High Level Design document [DES/PER/HLD/0001](#) level information derived from the TRIOLE Active Directory templates can be found in the Active Directory Low Level design.

Failure	Effect	Action
p-blade failure	None	Replace hardware
p-server components failure	None	Resource components are not single points of failure, additional resources can be allocated to maintain resiliency upon single failure.
c-blade failure	None	Configuration failed over to spare c-blade and c-blade replaced
SAN connection failure	None	SAN connections provide resiliency (see SAN design for details)
Network failure	None	Network design provide resiliency (see Network design for details)

Table 6 Bladeframe Failure Effect and Action

3.5.4 Domain Name Service

A primary and secondary DNS service is provided based on dedicated linux servers. DNS is designed for resilience across sites, and therefore the resilience model is also the DR model. The Windows Active Directory domain controller infrastructure is also a secondary DNS server, and provides DNS services to the Windows platforms in the estate.

The following details is taken from the Domain Name Service High level Design ([DES/NET/HLD/0006](#))

The DNS must be resilient to:

Server Failure

Software failure

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE

Network failure

Data fill errors

The internal DNS is authoritative for the HNGx domain "*horizonng.com*" and comprises:

- One logical server at each data centre to host Windows services and DNS information. This information is hosted within the Active Directory domain on a Windows 2003 server.
- One logical server at each data centre to act as the primary name server and host non-Windows information. This information is hosted on a UNIX server running the BIND DNS software.

One Virtual IP address (VIP) across both data centres to provide a resilient access method to the DNS service.

The primary DNS for the information above is at the active data centre; this is termed the primary master DNS server. A secondary server exists for each service at the DR data centre; this is termed the secondary master or *slave* DNS server. The primary and secondary servers are master (authoritative) for the information served. The secondary DNS servers at the DR data centre will transfer their zone information from the primary servers in the active data centre by zone transfer. The secondary servers are dependent on the primary servers and will cache the data transferred from the primary server in local files³.

The following types of failover are provided for:

Network Interface Card (rEth/appliance NIC) failover

IP failover

DNS server failure

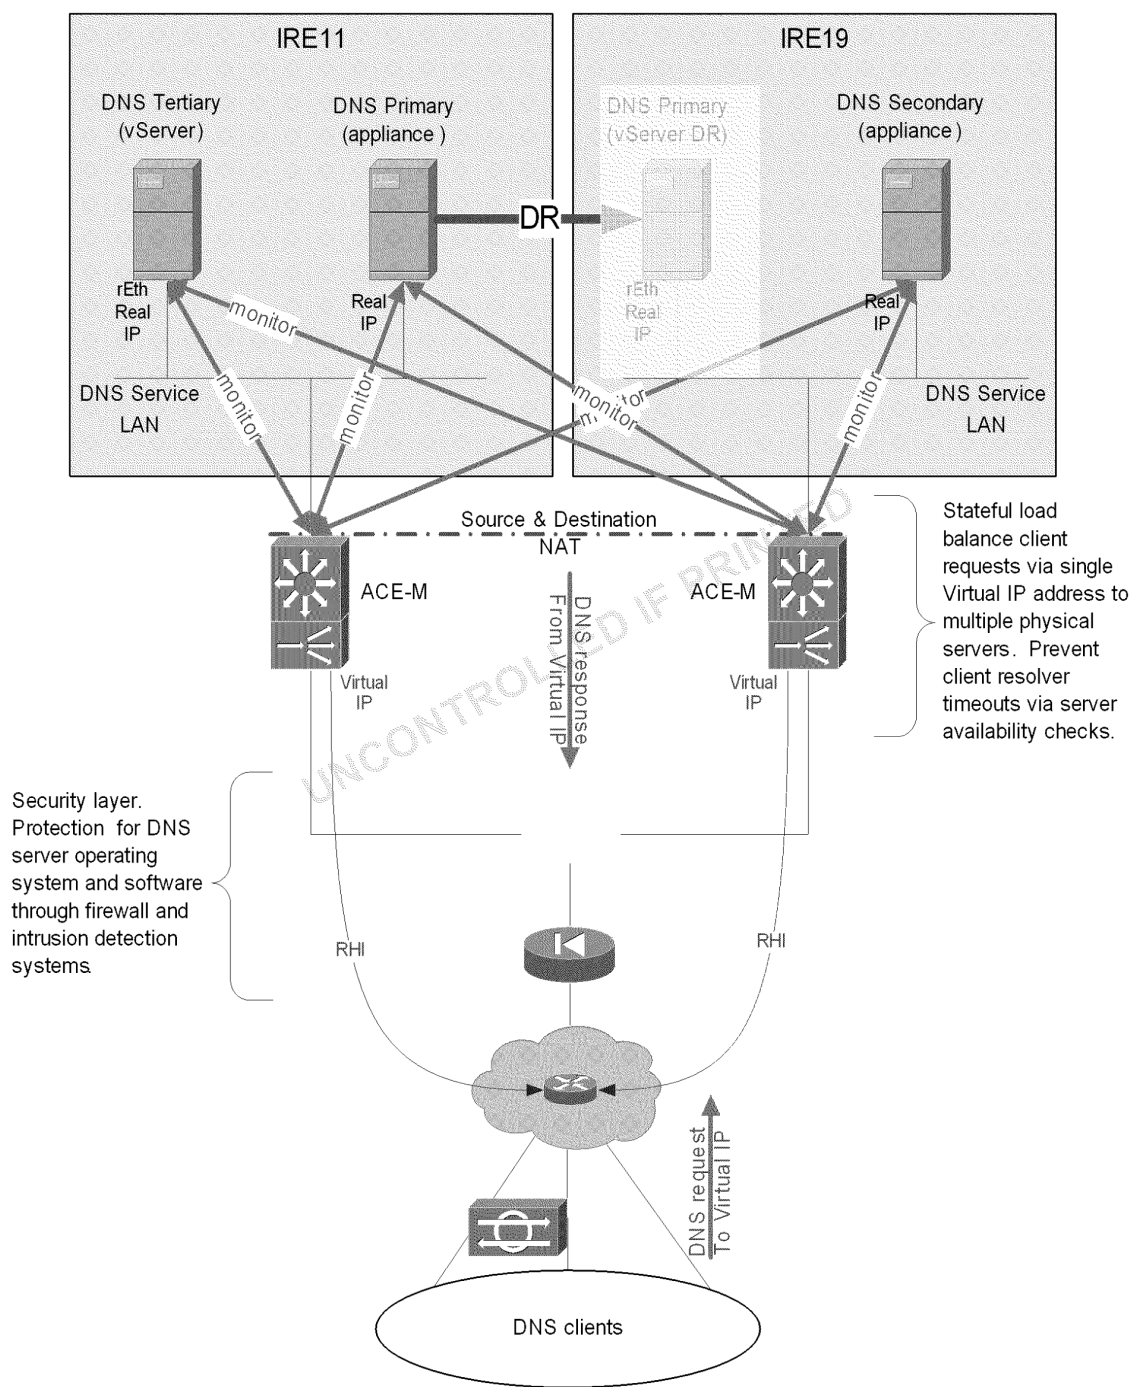
DNS service failover

DNS client failover

³ See the section on Resilience for why this is important.

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE



V0.2

Figure 4 DNS service resilience overview



3.5.5 Remote Support and Diagnostics

The following testing considerations are drawn from the Remote Support and Diagnostics Topic Architecture (ARCSYMARC0004)

The software facilities that together provide Enterprise Management are known as SYSMAN. SYSMAN2 is the generic name for all such software in the Horizon environment; this is updated to SYSMAN3 for the HNG-X solution

3.5.5.1 Remote Access Framework

The notional framework that will support remote access will require a number of discrete components. These will include:

- The Tivoli Management Framework.
- The Secure Access Servers (SAS) that will offer the point of access for all workstations, via Windows Terminal Server.
- An OOH facility that will allow remote access to the Campus via secured dedicated laptops connected to the SAS
- Secure 3rd party support access
- Secure Support workstations
- Secure Console access

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE

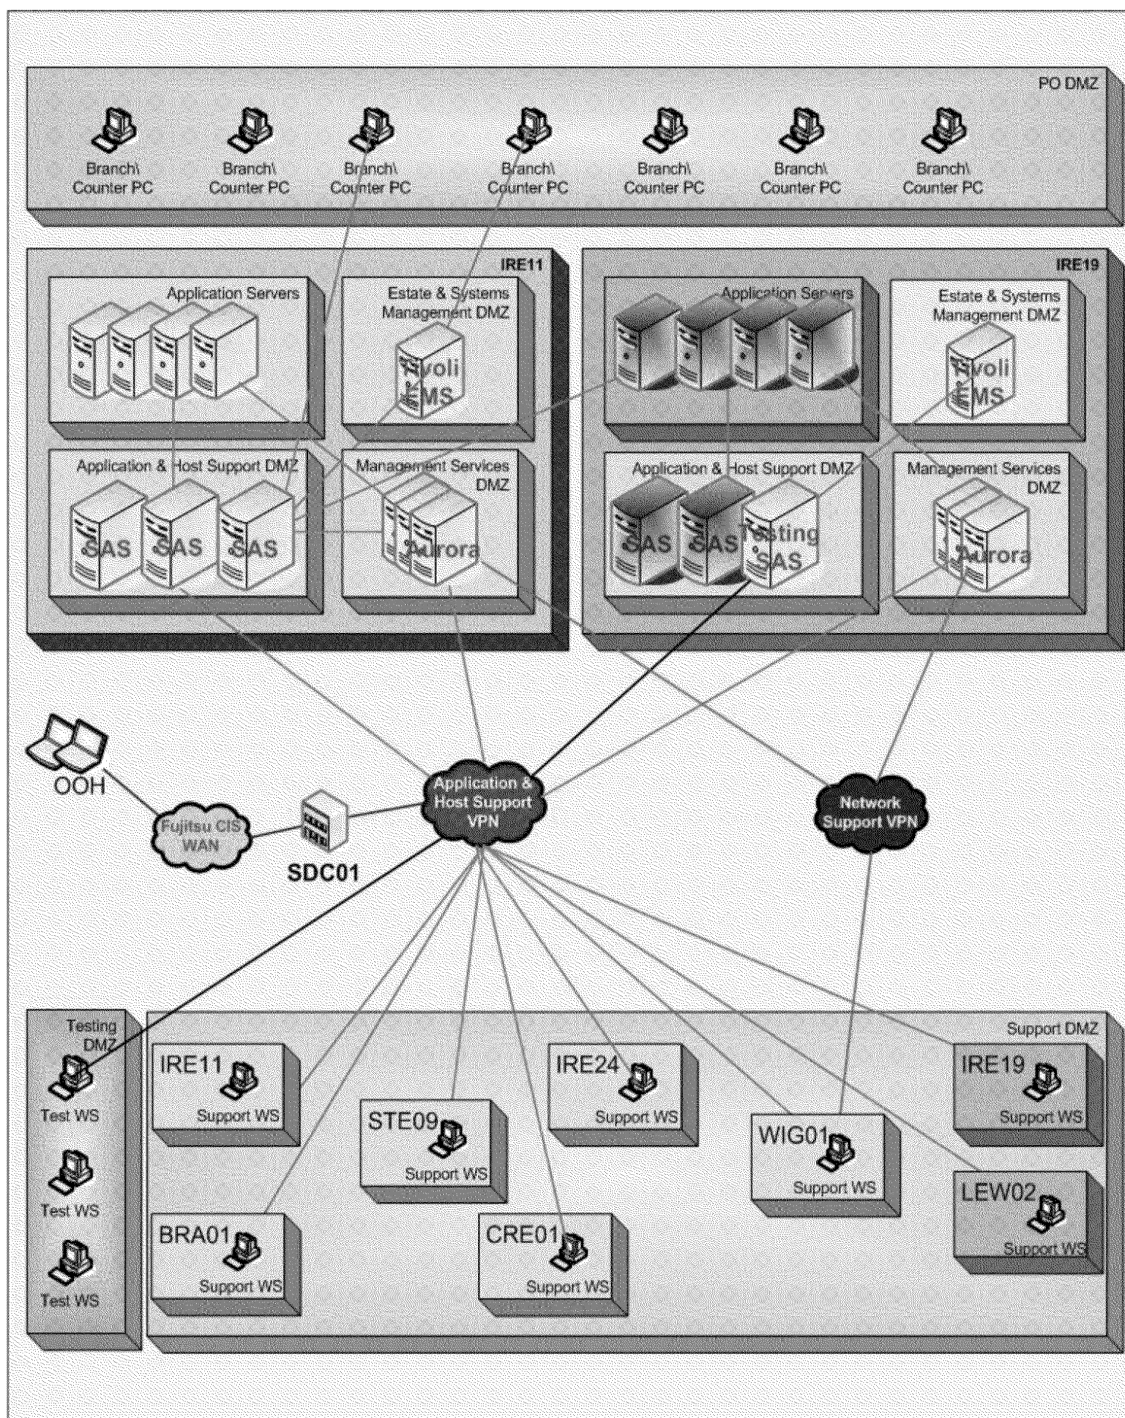


Figure 5 Remote Access Framework

**COMMERCIAL IN CONFIDENCE****3.5.5.2 Recovery and resilience**

In common with the RMGA application the SAS and Tivoli framework will be active in one campus, with failover to another server within the same campus. In the case of severe failure the complete environment will be reinstated the other campus.

3.5.5.2.1 SAS

The use of three SAS at each Data Centre will allow for failover to the remaining site in the event of a DR situation. In the event of a single server failure at one site one of the remaining SAS can be used for remote support access.

3.5.5.2.2 Test SAS

One Test SAS will be situated in the main campus. In the event of a failure the main SAS can be used until the Test SAS is brought back on line.

3.5.5.2.3 Aurora

There will be at least two Aurora servers located at each site providing a recovery mechanism in the event of Data Centre loss.

3.5.5.2.4 TMR

The TMR will be the subject of hot backups at regular intervals which will allow for recovery onto an alternative platform. Such backups must be available in both sites to allow consistent recovery

3.5.5.2.5 TPM

The TPM volatile data is held on the external database server so no special backup recovery is required; the exception to this is the local user database which must be backed up at regular intervals to ensure user/role changes are recoverable.

3.5.5.2.6 Tivoli Gateways

Since there will be multiple gateways and the Endpoints are agile in the case of gateway failure, resilience is automatic. There is little volatile data on a Gateway so on failure it is simply the case of loading the software onto a new server. This is similar in the case of Campus failure.

3.5.5.2.7 Workstations

There are many instances of workstations so there is no need for action in the case of single failure. Single failures would be handled by the normal incident management process. In the case of Campus failure the workstations will need to be reconfigured to point to the other Campus.

3.5.5.3 Migration

The main issue for remote support access is to maintain the access to newer systems and estate management tools as they are brought into HNG-X, whilst providing access to existing Horizon systems and tools.

The Volume and Integration testing environment will be accessed through the HNG-X SAS.

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

The Tivoli migration from Horizon to HNG-X will cause particular problems within the area of Remote Support and Diagnostics. The main reason for this is the current Tivoli TMR, Gateway and Endpoint versions are only compatible with Microsoft up to and including NT4 SP6, currently found at the Branch estate.

The target TMR, Gateway and Endpoint versions for HNG-X (SYSMAN3) are supported in Windows version later than, and including, W2000 but not NT 4. Some Campus servers are built around Windows NT and are not being rebuilt for HNG-X. Legacy Campus machines (NT SP6, Solaris 8); Horizon Branch and HNGX Branch (pre XP) will be managed by the Legacy TMR and Gateways. Any New Campus machines (i.e. RHEL, MS 2000+, Solaris 10) will be managed by SYSMAN3. During migration, therefore, there will have to be the ability to easily identify under which structure a machine is managed.

The above facts cause the need for a controlled migration with some Legacy support servers required in the short term.

3.5.5.4 Testing and Validation

To test the HNG-X Remote Support and Diagnostics architecture a test unit will need to access the servers in Belfast. Listed below are the high level testing requirements:

- Access to the SAS, Tivoli Management Framework and identity management solution for all phases of HNG-X testing.
- Access to provide the ability to connect to Databases in the Data Centre via the SAS. This will be provided using a specific Test SAS.
- Access to allow the take over of the GUI on servers in the Data Centre via the Test SAS.
- The access has to be secure, and not infringe Fujitsu Services or RMGA security policies.
- Large amounts of data may be transferred between workstations and the HNG-X System. This will require the selection of a tool to provide secure data transfer.

The following BCP/DR testing considerations are drawn from the Remote Support Secure Access Server High Level Design ([DES/SYM/HLD/0017](#))

Section	Design Text
4.1.1 (Access)	OOH Support access will be provided using OOH laptops which will provide access during disaster recovery periods
8.2 (Availability)	The platform will provide resilience and repair described in the Windows 2003 platform design. For the blade hosted SAS in IRE11 and IRE19. For HNG-X it is planned to have 3 SAS in each Data Centre.



3.5.6 Branch Support Database

The following testing considerations are drawn from the Branch Support Database High Level Design (DES/APP/HLD/0023)

The document describes the design of the Branch Support Database (BRSS), which will be built on a new database server using Oracle 10gR2 Enterprise Edition software.

The Branch Support Database will be populated with transactional, reporting and control data replicated from the Branch Database on a near real-time basis. Adequate security measures will be put in place to ensure that sensitive end-user information will not be replicated across from BRDB to BRSS.

Third line support (SSC) access to the Branch Database is limited and controlled to safeguard the performance of the Branch against the support actions of SSC. Hence a separate support database (single instance) is being made available for third line support.

The data will be retained in the BRSS for longer time duration as compared to the data retention in BRDB. The reason for this is to satisfy the requirement of support streams to be able to access such data over an extended period of time.

The Branch Support Database will provide a centralised point of access to live transactional, reporting and control data for the various business and technical support streams.

BRSS will also be used to generate intermediate reports in order to satisfy some of the SLA reporting requirements.

3.5.6.1 Availability

The Branch Support Database is a support system whose primary objective to provide a dedicated point of access to Live data for the support streams. If BRSS were unavailable, the data is accessible from alternate sources such as the Branch Database for more recent data and the audit storage system for historical data. Because of this, there are no hard availability targets set for BRSS.

However it is expected that in the event of a failure, the BRSS be available for use within 24 hours.

3.5.7 Restart/Recoverability

Process Control supports process restart and recoverability. For more details on the common process control, refer to HADDIS - Host Applications Design and Development Standards (DES/GEN/STD/0001).

The Restart/Recoverability at database/instance level is discussed in Section Failure & Restart/Recovery.

3.5.8 Fail-over

The Branch Support Database runs on the BladeFrame Linux environment. Fail-over is implemented using a remotely mirrored EMC file store while resilience to hardware faults can be handled by the BladeFrame technology's inherent fault management and resolution. This assumes that the BRSS environment operates under a N + 1 server strategy.

For a detailed discussion of Failure scenarios and failover, refer to Section Failure & Restart/Recovery.

3.6 Failure & Restart/Recovery

This section lists the various Oracle failure scenarios in detail and discusses the action to be taken in response to each different type of failure.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

3.6.1 Problem Scenarios & Recovery Options – Summary

Problem Type	Problem Detail	Recovery Action Details
Instance Failure	The Branch Support Database instance crashes	<p>The instance failure will be flagged by Grid Control and support will need to manually restart the instance.</p> <p>Data replication from Branch Database will pause without any data loss while the instance is restarted as long as the restart occurs within 24 hours as the storage area used to hold pending Streams packets is sized to hold 1 peak day's transactions.</p>
	Branch Support Database Node crashes	<p>The failure notification will occur via ITM Tivoli agent and via Grid Control (from Oracle perspective).</p> <p>BladeFrame technology will attempt to automatically restart the failed 'pserver' on a different physical blade in the LPAN. Once RHEL4 initialises in the new blade, BRSS database instance will automatically start via oratab.</p>
	The Branch Support Database instance crashes but fails to restart	<p>This scenario is a continuation from problem #1.</p> <p>If a manual restart of the failed database instance does not resolve the issue, the error must be diagnosed by support and the database instance restarted. As mentioned earlier, the Streams storage buffer in Branch Database can hold one day's transactions. If the fault resolution is in danger of exceeding that period, it is recommended that the size of the storage buffer⁴ be increased to avoid any data loss in the Branch Database.</p> <p>In the meantime the Branch Database can be used by support for running low-impact queries.</p>
	The Branch Support Database node crashes but fails to restart	<p>This scenario is a continuation from problems #2 & #3.</p> <p>If the BladeFrame PAN manager cannot automatically restart the blade, support will be flagged and will need to resolve the issue manually. The same Streams buffer storage constraints that were discussed in #3 apply here.</p> <p>In the meantime the Branch Database can be used by support for running low-impact queries.</p>
Media Failure	Alert-logs, RMAN backup (check logical), dbv, Grid Control detects a corrupt block	<p>If the block corruption results in a database instance failure, as may be expected if the corruption affects the data dictionary in the system tablespace or the online-redo logs, the resolution options discussed as a part of Instance failure (above) should be followed.</p> <p>If the block corruption does not result in a node / instance failure, block recovery using RMAN "blockrecover" command should be attempted. Block-recovery will make use of the backup sets to recover the last known image of the block and any changes made using the archived and online redo logs.</p> <p>There may be simpler ways to recover from corruption of certain objects by utilising application knowledge. Refer to Section 2.2.8.5.5 for further details.</p>
	The monitoring / verification tools detect a data file corruption	<p>As discussed for block corruption above, if the file corruption results in database instance failure, the resolution options discussed as a part of Instance failure (above) should be followed.</p> <p>If the file corruption does not result in an instance failure, the corrupt file and its contents can be recovered by standard RMAN recovery options.</p>
	The monitoring / verification tools detect a data file corruption along with corruption of online / archived redo log files	<p>The Branch Support Database writes archived redo log files to two different locations and each location resides in a separate EMC DMX. This makes the possibility of redo-log corruptions due to hardware issues unlikely. There is relatively a greater chance of redo logs being corrupt due to a logical corruption in the Oracle / EMC buffers.</p> <p>If such a corruption results in database instance/s failure, a full RMAN-based database point-in-time recovery should be performed. There is a potential for data loss here, which cannot be avoided.</p>

⁴ The size of the storage buffer can be increased by resizing the tablespace BRDB_STREAMS_DATA.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

	The monitoring / verification tools detect a corruption of the database control file	This will most likely result in a database failure, as the corruption will be replicated across all copies of the control file. Recovery is a straightforward process: Restore a previous backup of the control file and apply any file-level changes made since then and manually add any temporary files i.e. temporary tablespaces. Perform full RMAN database recovery to bring the restored control file in line with the change numbers in the data files.
Network Problems	Network slows down which results in the Branch Support Database to fall further and further behind the Branch Database	The network problem must be resolved manually. The Streams storage buffer in Branch Database may need to be increased if the Branch Support Database is in a danger of falling behind the Branch Database by more than 24 hours as a worst-case scenario.
User Errors	User inadvertently drops a table	The dropped table can be recovered using standard RMAN backup recovery procedures. RMAN will need access to the most recent Level-1 backup and all subsequent Level-2 backups along with copies of archived redo log files since the Level-0/1 backup. Refer to section Error! Reference source not found. for details on the location and retention period of the various backup components.

3.6.2 Failure Types

3.6.2.1 Instance Failure

An *Instance failure* occurs when software or hardware problems disable an instance. A common cause of instance failure is a Network / OS level glitch, which tends to disappear after a short period of time hence the automatic restart should be set up.

Grid Control will be used to monitor instances and alert in the event of a failure. There is an option of scripting in a restart of the failed instance into OEM Grid Control but it will not be used.

3.6.2.2 Media Failure

A *media failure* is a physical problem that arises when Oracle tries to write or read a file that is required to operate the database. An example is disk head crash causing loss of all data on a disk.

All Branch Support database components essential for the operation of the database will reside on SRDF replicated Rail-1 / Raid-5 disks. SRDF will ensure that every byte saved to the disk on the local (primary) site is synchronously replicated and saved to the secondary (DR) site before a positive response is returned back to the OS or disk management software such as ASM.

RAID-5 technology ensures that every byte of data written is mirrored on an erstwhile location in the disk array so even if a disk in the disk array were to fail, there will be no data loss.

The combination of SRDF and RAID-5 ensure that the possibility of media failure is greatly minimised.

The Oracle database will be running in archivelog mode and the archived redo logs are copied across to multiple locations so even if a media failure were to occur, provided the archived redo logs are accessible, a straightforward database point in time recovery will ensure that there is no data loss.

**COMMERCIAL IN CONFIDENCE****3.6.2.3 Block Corruptions**

A block corruption is said to occur when Oracle detects that the block wrapper (block header & footer) of one or more blocks is corrupt/invalid.

Once Oracle detects a corrupt block, it writes the following error message to the alert log for the instance that detects the corruption:

ORA-01578: ORACLE data block corrupted (file # %s, block # %s)

Where file# is the file ID of the Oracle datafile and block# is the block number, in Oracle blocks, within that file.

Any potential downtime is dependent on how quickly block corruptions are detected.

3.6.2.4 Detecting Block Corruptions

An Oracle error #1578 does not always mean that the block on disk is truly physically corrupt as the error message might just be indicating that there is a corruption in the cache. A number of means are traditionally available to verify that the block on disk is physically corrupted:

- Set DB_BLOCK_CHECKSUM parameter to TRUE (default).

DB_BLOCK_CHECKSUM determines whether DBWn and the direct loader will calculate a checksum (a number calculated from all the bytes stored in the block) and store it in the cache header of every data block when writing it to disk. This allows Oracle to detect block corruptions whenever the block is read. If the block is corrupted, error messages are written to the alert log but processes may not necessary fail.

This parameter will be set to TRUE for the Branch Support Database.

- Check messages in alert log on regular basis.

Oracle records any block corruption errors that it detects in the alert log. Automatic monitoring of corruption-type alerts in the alert log will be set up via Grid Control.

- Set DB_BLOCK_CHECKING parameter to TRUE (currently FALSE)

Setting this to TRUE causes Oracle to proactively check for corruptions and hence aid early detection.

- Detect logical corruptions using RMAN "CHECK LOGICAL..." command.

This check option will be used by the Branch Support Database as a part of the backup operation. Refer to Section **Error! Reference source not found.** for details.

3.6.2.5 Recovering from Block Corruptions

Once block corruption has been detected, the key task is to recover lost data. There are three options available for recovering data in the Branch Support Database:

- Use application-specific means of restoring data.
- Recover the corrupt blocks using RMAN
- Restore from Oracle export files

Oracle block corruption will be faithfully replicated by the mirroring system (SRDF) to the remote DR site hence failing over to DR is not an appropriate recovery method.



**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN**



COMMERCIAL IN CONFIDENCE

Index block corruptions are best resolved by rebuilding the index. There may be a performance impact of rebuilding and recommendations about time of rebuilding have been made in the later sections.

The appropriate recovery option to choose will depend on the type and criticality of the object being recovered and also whether the corrupted information has been consumed by all consumer systems. The Branch Support Database objects can be grouped together based on the recovery options available e.g. one group would be indexes where the recovery option would be to rebuild the index segment.

The discussion of object groups and recovery options available will be covered in the Support Guide.

UNCONTROLLED IF PRINTED

3.6.3 Aurora Console Access

The following is drawn from HNG-X Resilience and Disaster Recovery HLD ([DES/PER/HLD/0001](#)).

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

Aurora is used to manage all (serial) consoles for equipment such as cBlades, Solaris servers and Cisco switches in a controlled and logged manner without physical access.

There will be two Aurora systems at each site with interfaces on the management LAN, each managing complementary equipment, e.g.

Aurora1 manages bf001/cb1

Aurora2 manages bf001/cb2

Aurora1 manages core2 network switch (from mgmt LAN on core1)

Aurora2 manages core1 network switch (from mgmt LAN on core2)

Aurora1 manages Aurora2 and vice versa

Aurora connectivity is typically only used during "dead server" type recoveries, and is not required to be highly available, but is very useful for looking at the logs to see what went up the screen as the system died.

In the event of major disruption preventing access site access will be requested by the local Unix Support team who will gain emergency access via the Aurora physical console port until general connectivity is restored. This has never been required in Horizon.

There is no DR requirement for Aurora itself, but it is a critical component in Solaris DR to allow properly managed reboots

The following BCP/DR testing considerations are drawn from the Aurora Console Access High Level Design ([DES/SYM/HLD/0020](#))

Section	Design Text
6.1 Overview	A console server will be installed at each data centre.
8.1.3 Server Reliability	<p>An additional Console (standby) server, Aries multiport card and modem should be purchased, per installation, to provide resilience, protecting against Console servers failing. In the event the Console server fails manual actions will be required to disconnect the failed server from all remote devices replacing it with the standby server.</p> <p>The standby server should be connected to the network and kept up to date with software updates and configuration changes. In order to ensure configuration changes are applicable to all three Console servers (specifically with regard to device configuration files - see section 5.2.2.1 above), the same server at each site should be connected to the same port on the Aries card on all Console servers.</p>

3.6.4 Authorisation Services

The Authorisation Services are described in the Topic Architecture document [ARC/APP/ARC/0005](#)

**COMMERCIAL IN CONFIDENCE**

The following is drawn from Generic Authorisation Services High Level Design [DES/APP/HLD/0006](#), additional client-specific design's are captured in the individual High Level Design documents that are children of this document. Specifically,

DES/APP/HLD/0007	DCS Authorisation Agent High Level Design
DES/APP/HLD/0008	ETS Authorisation Agent High Level Design
DES/APP/HLD/0009	NBS Authorisation Agent High Level Design

3.7 General

Mechanisms are required to ensure that where an instance of an Authorisation Agent fails another will take its place. Resilience is achieved by having two instances of the Agent running at the same time on different platforms. Both agents start in standby mode. Heartbeats exchanged through the NPS are used to decide which agent becomes active and when and if the standby agent should take over.

When an Agent instance fails, all outstanding transactions being processed by that Agent will be abandoned. No attempt will be made by the replacement instance to recover such work (e.g. from the NPS). If possible, the failing agent informs the Branch Access Layer of its change of state by sending an ASTS message (see **Error! Reference source not found.**) to each Online Service Router connected to it. The replacement agent informs the Branch Access Layer of its change of state in the same way. The OSRs route subsequent messages to the active Agent. Requests that were waiting for responses from the failed Agent may be failed immediately or left to time-out.

Successful failover relies on the NPS to provide the necessary reliable communication mechanism between the active and standby instances. The NPS is implemented with two Oracle instances. If an instance fails the Agent process continues operating using the other NPS Oracle Instance while trying to re-establish the failed connection. Failure of both instances leads to termination of the Agent and a Tivoli controlled restart.

Controlled failover from an active instance to a standby instance occurs when the active agent "resigns" following loss of a critical resource and writes a heartbeat showing itself to be unavailable. (An active agent won't resign if the standby is unavailable.) Unilateral failover occurs when the standby agent fails to read a heartbeat from the supposedly active agent for some configurable period.

Note that an active agent will not fail over on loss of its connections to the external client. The resilience built into the network is such that the standby agent (which is running in the same Data Centre as the active agent) is unlikely to fare any better.

3.8 Heartbeats

Each Agent instance writes a Heartbeat to the NPS periodically (configurable in registry). Each of an instance's Heartbeats overwrites its previous Heartbeat, so only its most recent Heartbeat is available. A separate Heartbeat History table makes Heartbeat information available for systems management purposes. The History table is updated each time a heartbeat is written that differs from the previous heartbeat in a field other than the timestamps.

The Heartbeat includes the following information for failover purposes:

- *Agent Active* – Whether the Agent instance is actively processing NBX transactions.
- *Agent Available* – Whether the Agent instance is operational, i.e. whether it would be able offer an authorisation service at all. Note that a standby instance is normally operationally available.
- *Client Available* – The availability of the Agent instance's connections to the external client. The meaning of 'available' varies according to the client in question. This datum is only relevant for



COMMERCIAL IN CONFIDENCE

an Active instance - a Standby instance does not establish any connections. The value can indicate all, some or no connections available.

- *Priority* – A static priority value obtained from Agent's registry. This is used to determine which Agent instance has priority at Agent service start-up. It is also used to determine which Agent should 'resign' where some error situation has led to both being Active. The values 1 and 2 should be used, with 1 being the higher priority

As Heartbeats are also used for operational monitoring of the Authorisation Agent instances, an Agent instance will write a Heartbeat even if it is operationally unavailable.

Each Agent instance connects to both NPS Oracle instances, so that it can read Heartbeats from both. (Note that they are written to only one Oracle instance.) Failure of either connection will require the Agent to enter a retry loop until connection can be re-established. Failure of both connections will result in the Agent exiting, at which point it will be restarted by SYSMAN.

Each Agent instance will attempt to write a final Heartbeat when it closes down, in both normal and exception scenarios, to inform its Partner what it is doing. Agent_Active will be set to 'N' and Agent_Unavailable to 'Y'.

3.9 Agent Failover

There are two distinct failover scenarios to be considered:

- *Controlled failover* – both agent instances are running and can control the decision
- *Unilateral takeover* – an expected Heartbeat refresh by the Active Agent has not happened

The algorithms for both controlled failover and unilateral takeover take account of the exceptional running of three or more Agent instances.

3.9.1 Controlled Failover

Upon reading a Heartbeat, an Agent instance has to decide whether it is the one that should be the Active instance.

Controlled failover from an active instance to a standby instance occurs only when re-homing (see 2.2.10.7) or when the active instance resigns following loss of a critical resource. The resigning agent's final Heartbeat will set *Agent_Unavailable* to 'Y'. The active instance will not resign if its Partner is advertising itself as unavailable.

If more than one instance is claiming to be available, the decision-making between them is as follows:

- *Agent_Active* – an active instance takes precedence over a standby instance (i.e. a standby instance will never negotiate to wrest control from an active instance)
- *Priority* – everything else being equal, the instance with the lower priority number will take precedence⁵. This is relevant if both instances are in standby, as could happen if both are started at near enough the same time, or both are active (which should never happen)

Not all resource losses cause a failover. The networks between the Authorisation Agents and their clients are designed to be highly resilient. As such, an active Agent will not fail over to its standby in the event of network problems – the standby Agent is very unlikely to find the network in better shape. Failure of an external client connection is also not a reason for failover, since the external client is outside the HNG-X operational domain. Failover may not improve the situation and would unfairly count against HNG-X availability rather than the client in service level statistics.

⁵ A configuration error could leave both instances with the same priority. In that event the Agents will use the alphabetical order of their host names as the final discriminator.

**COMMERCIAL IN CONFIDENCE**

3.9.2 Unilateral Takeover

A Standby instance will take over when it detects the non-refresh of the Active instance's Heartbeat for some period (configurable via registry).

3.10 NPS Failure

The NPS is accessible via two different Oracle instances. Each Authorisation Agent Server is physically connected to both Oracle instances; the Agent is capable of connecting to one or other of them at all times.

When a connection to one Oracle instance fails, any Agents connected to it need to automatically fail over to use the other Oracle instance. Failover is expected to be achieved within a few seconds, therefore minimising the period that the authorisation service is affected. If the Agent loses its connection to both Oracle instances, it will fail.

During the connection phase when the Agent is first loaded, the Agent attempts to connect to both Oracle instances. Until it has succeeded connecting to at least one of them, it adopts the standard Agent approach of retries as appropriate; retries continue until the configured *total_connection_timeout* period has elapsed, after which the Agent fails. Each Agent is configured to have a preferred Oracle instance, for overall performance reasons. There is a delay between attempting to open the preferred instance and the non-preferred instance.

Once it has connected to one Oracle instance, the Agent can enter its main processing phase. Attempts to connect to the other Oracle instance continue indefinitely (and are not controlled by the *total_connection_timeout*).

Similarly, following a failover from one Oracle instance to another, attempts to re-establish resilience are made by trying indefinitely to reconnect to the original Oracle instance.

3.11 Re-homing

For efficiency, an Authorisation Agent should write its Transaction Journal records to the same Oracle instance as any process that is harvesting those records. This is material during the core day, merely desirable outside core hours. Such components are configured to prefer a particular Oracle instance, but failures may cause one or more participants to fail over to the non-preferred instance. In other words components accessing the NPS can get out of step with regards to their active Oracle instance, and it is desirable to resynchronise at some point.

As an independent requirement, it is helpful to operations staff if they can predict where active instances will be running, and resynchronisation following failover is also desirable.

These two (independent) re-homing tasks are configured via the registry to occur together at some time overnight. The active agent stands down (terminates) in favour of the standby agent if its priority is lower than that of the standby (see 2.2.10.4). Both active and standby agents revert to their preferred Oracle instance.



COMMERCIAL IN CONFIDENCE

3.11.1 File Transfer Managed Service (FTMS)

The following details are taken from the File Transfer Managed Service High Level Design (DES/APP/HLD/0051):

Section 4.6 Resilient and Recoverable	<ul style="list-style-type: none"> The solution must be able to recover and retransmit files from the last successfully completed transfer after a system crash or lockout (provided the system is NOT dead or corrupt) The solution must have a retry facility that overcomes any temporary aberration on the network/systems If recovery is not achieved automatically by the solution then simple operating procedures must be developed. The solution must detect and report errors.
Section 5.2.4 Multiple FTMS Channel Architecture	<ul style="list-style-type: none"> The failure of one channel must not stop other channels from working
Section 5.6.1 Networking Services	<p>There are currently three supported physical connections for an FTMS link: ATM, Frame Relay and ISDN.</p> <ul style="list-style-type: none"> Resilience and recovery may vary from one connection type to another
Sec 5.6.3 Domain Design Resilience and Recovery	<p>The loss of an individual platform should not prevent FTMS from being able to process files via an alternative route if there is one. When the system is fixed it should be easy to restore the system for use. Protection of data also needs to be considered and an appropriate backup strategy defined.</p> <p>The structure of domain design and placement of domain controllers directly affects the resilience and recovery of the NT infrastructure.</p> <p>There are also known issues with dual network cards.</p>
Sec 8 System Qualities 8.1 Availability	<p>FTMS application must be developed to be able to support 24-hour use with minimal recovery time in the event of software or hardware failure.</p> <p>The network link to the remote sites should be designed with built-in protection against single points of failure.</p>

The generic recovery procedures for an FTMS link are covered in [TD/STR/007]

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

3.12 Features Not to be Tested

The full Q4 2008 DR activity is outside the scope of this plan and is documented in HNG-X Business Continuity Test Plan ([SVM/SDM/PLA/0003](#)).

Security testing should be covered via other test streams and where particular tests, perhaps by virtue of the sensitivities surrounding secure materials (e.g. Live keys, encryption algorithms, etc.), require specialised or highly secure test environments, then it will not be practicable to combine such tests with the mainstream threads, and these will be dealt with in a separate stream of testing – Security Test. The resilience aspects of Security such as loss of a Firewall should be tested by the Integrity test team but under the guidance of the Security test team, utilising their expertise and knowledge to determine the impact on Security.

The Test Environment at IRE19 will not be tested separately for resilience. There should be no resilience issues as the HNG-X elements should be identical to IRE11. IRE19 is critical in the maintenance of the HNG-X solution and there is a known issue regarding DR where the availability of IRE19 for a POL project delivery may take precedence over its availability for DR.

The Help Desks though critical in support of the Post Office Business are outside the scope of this document. Currently there is no documentation showing how the HNG-X/Hydra help desk solution differs from that currently implemented.

Performance (Load, Volume and Stress) testing as defined in [TST/GEN/STG/0001](#) is outside the scope of this document and is covered specifically by the Volume test HLTP [TST/SOT/HTP/0003](#). However, there are elements to the solution, such as loss of 1 part of N+1, which will require that Performance testing be undertaken. This is documented where appropriate and tests devised. It is envisaged that these tests will be performed by the Integrity team with close cooperation of the Volume team.

NB Both teams are directly under the same manager.

Testing of the complete migration process is not in the scope of Integrity Test Team.

A backup EDG platform will be implemented at the PRISM data centre in Maidstone (Post Office CR PSOCR01178). This will be linked to the Live HNG-X Data Centre. The Live EDG at the Northern Data Centre will link to both the Live and DR HNG-X Data Centres.

Testing of the EDG back-up/failover process is outside the scope of HNG-X testing.

3.12.1 (Hydra) 'Migration States' Integrity Testing

The target end state of the HNG-X project is that all POL Counters will run HNG-X (rather than Horizon) software and those counters will be served by the data centre IRE11 with the data centre IRE19 as standby. Supporting Horizon and HNG-X counters out of IRE11/19 will be achieved through transitioning multiple stages. (i.e. Weekends A, B, C etc)

To assure the confidence in capability *for each and every* migration (Hydra) state one might consider the following testing tasks appropriate:

1. Validate and Verify the processes required to effect the Migration step
2. Verify and Validate the processes required to assure a completion of migration step
3. Test for any regression on the Functional Requirements of the Solution
4. Test for any regression on the Non-Functional requirements including
 - o Performance (Volume, Load, Stress)
 - o Integrity (Resilience, Recovery, Error Reporting)

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE

Tasks 1, 2 and 3 will be carried out by RV (Release Validation) test stream. These assertions however beg the question of how assured will POL be regarding the various Hydra configurations' capabilities to satisfy the Non-functional requirements.

Without the availability of a fully specified test rig to undertake the Non-Functional testing for each Migration (Hydra) State it will not be possible to give the highest level of assurance of System Capability before is goes live.

This document is concerned with the Integrity aspects of Non-functional testing in the project so the following comments should be read in that context.

Until Migration Weekends A, B, C and D are complete in the production/business domain, POL business workflow requires both the Bootle/Wigan pair and the IRE11/IRE19 pair to be running; with full resilience provisions. Whereas the current test plan provides for testing the new and replacement hardware/software at the (IRE11/19) data centres using the IRE11/19 kit configured as test rigs, the plan makes no provision for the use of test rigs (of similar specification to Bootle/Wigan) for non-functional testing of the dual (i.e. England/Ireland) data centre migrations. The (IRE11) VI test rig *will never be taken through the planned migration steps*. I.e. the VI test rig will only be available to the VI Team when it is in the state where all Horizon and HNG-X Counter business is transacted through IRE11(19) Data Centre Services {i.e. Shared Batch | Shared Online | HNG Branch | Horizon Branch}. The conclusion therefore is that the VI test team will have no opportunity to test the IRE11/19 migration configurations.

In summary, the current plan is that no physical, E2E integrity or performance testing of the various, migration states (other than the final state that will go into Pilot) will be possible on the VI test rig.

This document envisages Integrity Testing effort with respect to the Hydra States to comprise the following:

- Participate in the process of Risk Assessments of the Multiple, Transitory Migration States
- Identify which of the planned Integrity Tests (i.e. those applicable to the "Final" Hydra State [= post Weekends A+B+C+D]) are applicable to the Transitory States
- Produce a Gap Analysis between the outputs of the two tasks immediately above



4 Risks

The general project risks are to be covered in the overall HNG-X Projects Risk Register.

Assumptions, Risks & Constraints specific to testing are detailed in [TST/GEN/STG/0001](#)

4.1 Hardware and Software Risk Issues

- As the Resilience Test is testing a number of new technologies there is risk that the architecture will change. This may result in additional Test Cases being required or in the removal of or changes to existing Test Cases.
- The complete test environment is available in IRE11 i.e. Will the software & platforms be built and delivered on time to meet the start of testing. If there are significant delays it may be the case that the tests are prioritised and successful coverage of tests deemed of high importance is sufficient. NB The current intention is that the V&I environment in IRE11 will become the live environment.
- If the personnel involved in the collaborative working areas are not appropriately empowered, then the decision making processes will become protracted, significant levels of rework will result, and so costs and schedules will be adversely impacted.
- If the Acceptance Criteria are not produced at the outset, and couched in terms of the necessary test scenarios, it may not be possible to accommodate them collaboratively when planning and engineering the test materials as intended, and so may necessitate the re-separation of Post Office Ltd testing stages, with all the duplication of effort and increased timescales that will involve.

4.2 Planning Risks and Contingencies

- The documents on which this document is based are not all base lined. When they are re-issued it will be necessary to review this document and it's associated Test Cases to ensure complete coverage of design requirements. The resources required to achieve this review and implement changes may put the planned Test dates at risk.
- The SRRC is not available – this would prevent the proper design of tests as the expected outcome of a failure would not be apparent
- Integrity and Volume testing is taking place on the same rig. To attempt to achieve the challenging timescales will necessitate 24/5 testing. Shift working and the appropriate payments needs approval and agreement.



5 Quality

Quality Centre is the chosen product for test documentation and administration for the HNG-X delivery.

There is a link between Quality Centre and the Doors product which allows the direct linking of test cases to specific use cases within the Doors product.

The Test Cases and the design requirements from which they were derived will be recorded in Quality Centre. Within Quality Centre the requirement will include a reference to the document (Reference Number, Version and Section Number) from which it was derived. If a document, from which a design requirement was derived is changed the requirement, Test Case and Test Case reference to the requirement will need to be reviewed and possibly changed.

All test scripts will be generated from Quality centre.

During the Resilience Tests and DR Tests (if appropriate) the Test Cases in Quality Centre will be updated with the result of every attempted execution.

All Quality Centre data will be secured on a daily basis through normal processes that backup the server on which Quality Centre is hosted.

6 Approach

6.1 General

Everything except the Host and Agent layers are being entirely re-engineered, and with widespread revision of the supporting infrastructure across the whole solution, it is important that the system itself be used to confirm continuing integrity (DR test system). Resilience within the data centre is to allow the solution to survive the failure of a single part of the infrastructure.

The structure and contents of a HLTP are defined in [\[TST/GEN/PRO/0003\]](#)

The approach to non-functional testing is defined in [\[TST/GEN/PRO/0012\]](#).

Defect management is defined in [\[TST/GEN/PRO/0010\]](#).

The extraction of Test Case definitions from Quality Centre for input to a HLTP is defined in [\[TST/GEN/PRO/0006\]](#).

The testing involves two activities:

- (1) Resilience testing of the HNG-X Components
- (2) DR Testing of the HNG-X data centres

The testing execution will not be confined to the actual running of scripted tests against the components, but can be one of the following acceptance methods:

CT	Component Test
RV	Release Validation Test
DR	Document Review
DW	Design Walkthrough
SOO	Statement of Obligation
SOF	Statement of Fact
ST	System Test run by a Test team (not necessarily System Test Stream)

As an example:

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN****COMMERCIAL IN CONFIDENCE**

When resilience is provided by a backup/restore mechanism, then the testing would not be carried out by the test team but a support team and this could then be confirmed by either or all:

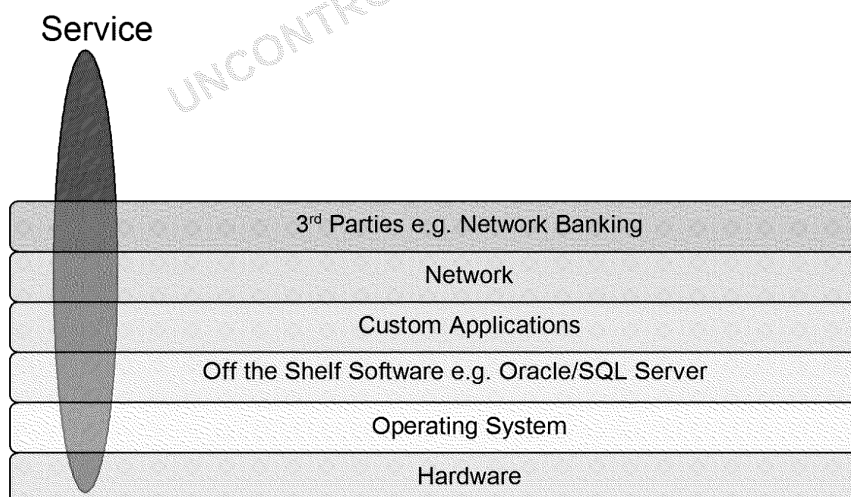
- 1 Reviewing the design (DW)
- 2 Ensuring the Work Instructions give details of backup/restore procedures (DR)
- 3 Observe backup/restore being performed by support team(RV)

There is a danger that the Resilience testing just tests either a component or the systems ability to deal with failure. However, equally important is how the system falls back. For instance if a server fails and the backup server takes over the workload what happens when the original server is repaired. No assumptions will be made on the fall back strategy, these are design considerations and the tests will be based on the design.

As important as the system having resilience is that the support team is made aware of any problems. Therefore, the testing must not only ensure that the components have the resilience as set out in the design but that the appropriate alerts are raised when an event, such as a component failure, occurs.

The testing is to ensure that the stated Requirements for Resilience and DR can be met. As Resilience is a system quality then all Architecture and High Level Design documents will need to state how they meet the appropriate Resilience requirements. Past experience would dictate that failure to do so will have a dramatic impact on the amount of time and resource needed to carry out the testing. This increase more in line with orders rather than a few percentage point increases.

Figure 6 Services Spear



There is a potential mismatch between the Business Continuity requirements of the Post Office and the testing carried out. The major concern of the Post Office is to carry out its business through the delivery of its services, spearing down each layer effectively looking end to end. However, the method by which testing will be performed is to concentrate on the resilience of the hardware and software components which support/supply the services, layer by layer. Attempts will be made to align the services with the underlying components to give confidence in the resilience of the service. This is similar to the approach taken with DR.

The main emphasis of the testing is the data centre as very little resilience actually exists in the Branch. The following is from HNG-X Systems Qualities Architecture (ARC/PER/ARC/0001):



6.2 DR Test Cycles

There are 4 DR failover test cycles planned with an initial walkthrough, these tests will be controlled by ITU and operated by the actual support staff. The number of test cycles allows all staff to be made familiar with the solution, and for F3 and F4 (as the final two tests are known) to be run in cooperation with the Business Continuity Manager and their counterpart from the Customer and could serve as acceptance tests.

V&I Integrity team will run a separate programme of Integrity testing designed to test the N+1 resilience features and recovery from backup. This will start at the earliest opportunity and run right through to the start of the migration weekends.

6.3 Test Case Analysis

The HNG-X testing strategy states that a full set of Business Requirements are to be formally stated for HNG-X, covering both functional and non-functional aspects. These will employ Use Case Modelling techniques, and will be elaborated to System Use Case level, and developed using UML, Sequence Diagrams, and non Activity Diagrams, etc. as and where appropriate and held in a central repository under formal change control these will in turn produce our low level test scripts.

6.4 Test Case Execution

V&I test team will complete the targeted Integrity testing within its own stream, collating the results where necessary from other test streams.

As a general rule tests are executed by a single test team within a test stream, in this instance for the DR testing it is likely that test execution/ completion will require the involvement of support personnel and other business-as-usual teams as well as the V&I and members of the joint test team.

6.5 Approach to Resilience

There are HNG-X Architecture documents covering:

- Counter
- Branch Access Layer
- Online Services
- Batch Applications
- Branch DB
- Counter Business Applications
- Network
- Platforms and Storage

Each of these areas will be examined to assess how resilience is catered for. This will be done by a document inspection of the Architecture, the High Level Design and the appropriate Low Level Designs. The owning Architects and Designers will also be contacted to discover what tests have been performed or are planned to ensure the resilience of that part of the solution. Resilience is not just how a piece of hardware copes with a failure but how the application deals with the failure. It is anticipated that a number of the tests will be performed by other test teams as part of the normal testing. However, it is likely that further testing will be needed particularly in the boundaries between applications and systems. Services have to be resilient to failure however if a service does fail then there is a target time by which



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

it needs to be available again. Testing is unlikely to proof conclusively that the time will be met just whether or not it is likely to be met. However, testing is likely to show if target recovery time will not be met.

6.6 Approach for DR

This is as document in HNG-X Business Continuity Test Plan ([SVM/SDM/PLA/0003](#))

7 Environmental Needs

V&I test rig will consist of the IRE11 HNG-X environment for Resilience testing and both IRE11 and IRE19 HNG-X environments for DR testing. This in turn will allow the resilience and recovery testing to be completed on what is essentially the live environment.

Both streams of the V&I team will work on the V&I test rig simultaneously, as well as the other outside demands or request that are expected as we draw closer.

8 Responsibilities

Resilience testing will be under the direct control of ITU Fujitsu Services staff.

DR testing responsibilities are outlined in detail in HNG-X Business Continuity Test Plan ([SVM/SDM/PLA/0003](#)) the 4 DR test slots will be as planned by the ITU with the inclusion of the POA BCP manager and IS Support.

IS Operation in IRE11 will be responsible for all Data centre DR and Resilience Process and Procedures.

All test cases and results will be reviewed by the Joint Test Team with exceptions, issues, etc referred to both the POL and FJS Business Continuity Managers.

9 Dependencies

- ✓ Post Office Ltd Business Requirements and Fujitsu Services Operational Requirements will be produced at the outset of the programme, to formally identify both the functional and non-functional requirements for Business Equivalence and the reduction in Total Cost of Ownership.
- ✓ Any Acceptance Criteria specified as to be satisfied by means of testing, will be couched in terms of the test scenarios necessary to demonstrate that they are met.
- ✓ The HNG-X development area will adopt exhaustive, generic, component level verification methods, in accordance with this strategic approach - Component Test (CT) and Component Integration Test (CIT) – and retain comprehensive Regression Test packs at this level.
- ✓ Recovery and Resilience testing requirements documented in Quality Center
- ✓ Quality Center kept up to date
- ✓ Guidance on Business Continuity areas of concern

Dependencies specific to testing are further detailed in [TST/GEN/STG/0001](#)

10 Schedule

TBA



11 Resources

The Resilience testing has been broken down into the Architectural areas. The intention is then to carry out the Resilience testing as far as possible when these areas are being tested. Therefore, assistance will be required from the appropriate teams to run the tests.

The Recovery testing will largely need the co-operation of the IS staff based in IRE11.

The DR testing will need the Data Centres in Ireland set up completely with links to 3rd parties and the standby system available as a Test Environment for Bracknell.

12 Entry Criteria

Due to the number of tests and the differing forms of the tests the Entry Criteria will be test dependent. As such each test within Quality Center must have the Entry Criteria specified. For the Resilience aspect of the testing examples of Entry Criteria are:

- ✓ Architecture and design documents approved
- ✓ Appropriate hardware available
- ✓ IS Resource available to demonstrate process
- ✓ Functional area Test Team available
- ✓ Low level test Plan for Functional Area testing
- ✓ Functional area testing results available in Quality Center
- ✓ V&I Rig available
- ✓ Injector Available from V&I Volume team

This will then lead to:

- ✓ Test Plan in Quality Center complete

For the DR testing the Entry Criteria will be documented in

HNG-X Business Continuity Test Plan [SVM/SDM/PLA/0003](#)

13 Exit Criteria

The following Exit Criteria are copied from a check list that is document in HNG-X Testing Process – Entry and Exit Criteria ([TST/GEN/PRO/0001](#)) and are applicable to this activity:

- ✓ Has all the testing been completed and recorded in Quality Center?
- ✓ Quality Center test reports produced?
- ✓ Have all critical defects been closed?
- ✓ Tested Component Code and Data
- ✓ Test Plan
- ✓ Test Outputs
- ✓ Test Report
- ✓ Product Breakdown Structure
- ✓ Agreed Deliverables

**HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN****COMMERCIAL IN CONFIDENCE**

- ✓ Have all other (non critical) defects been cleared? If not, have outstanding defect(s) been ratified for passing into next phase and a report produced?
- ✓ Have all the Objectives (Objective Driven Testing) designated for this stream, been covered within the testing cycles?
- ✓ Have all the System Acceptance Criteria been met in full or non conformances been agreed and signed off (and filed)?
- ✓ Are details of versions of all deliverables available for the handover to the next test stream? Particular attention to be paid to components and modules that have been amended due to fixes being applied.
- ✓ Test Results Package
- ✓ PMS Work Requests
- ✓ Test Report Review Output
- ✓ Is there a full set of testing reports available and filed?
- ✓ Is there any possibility that the system is unstable?

14 Test Pass / Fail criteria

Each Test Case defined in the Appendixes will have an "Expected Result" associated with it. The "Expected Result" will be recorded in Quality Centre which will provide an audit trail if it becomes necessary to change the "Expected Result". Changes may arise for example if the functionality is revised during the Testing, if the functionality defined in a HLD is revised or if the "Expected Result" is incorrectly set.

If the execution of a Test Case results in the "Expected Result" it will be "Passed".

If the execution of a Test Case does not result in the "Expected Result" it will be "Failed". The V&I Test Manager will be responsible for the progressing of all failed Test Cases to a status acceptable to HNG-X programme.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN

COMMERCIAL IN CONFIDENCE



A Appendix – Manual Testing

A.1 Explicit Requirements

The following listings constitute the planned deliveries into this < enter Test Level Name here >. Each delivery has a listing of expected tests to be carried out using manual testing methods

(Delete any of the 'Delivery' paragraphs (below) that are not required for this report. Rename the remaining paragraphs as required / appropriate)

The following is from TST/GEN/HTP/0002, then amended

DOORS Req ID	Acceptance Criteria ID	Description	Accept Method	Risk Impact	Acceptance Criteria
Key : CT = Component Test, RV = Release Validation Test, DR = Document Review, DW = Design Walkthrough, SOO = Statement of Obligation, SOF = Statement of Fact, ST = System Test					
ARC-484	ARC-437	The system shall switch automatically from a primary Branch telecoms network connection to a secondary connection where available in the event of connection failure, and back again at an appropriate point on restoration of the primary circuit and User may be informed.	DR Add : CT (demo failovers per component)	Interrupt	Documents provided with requirement cross reference, compliance with requirement in document. (TDN : N. Williams comment : A DR will demonstrate how the design meets this requirement) Add : Test to demonstrate failover of each component
ARC-484	ARC-466	The system shall switch automatically from a primary Branch telecoms network connection to a secondary connection where available in the event of connection failure, and back again at an appropriate point on restoration of the primary circuit and User may be informed.	ST	Interrupt	Tests to demonstrate network connectivity are defined and contained within a detailed test plan, the successful completion of which shall determine the acceptance of this requirement. To clarify 'successful completion' success criteria will be identified in advance to test execution (TDN: N. Williams comment: The functions should be testable. Clear success criteria for the tests (identified in advance to test execution) to be agreed and passed into PO DOORS)



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN

COMMERCIAL IN CONFIDENCE



ARC-443	ARC-490	The resilience capabilities for the Data Centre and the Data Centre network will be as specified in Data Centre Operations Service Description. (amended as per review 2/2/07)	SOO	Minor	Contractual Documents provided with requirement cross reference, compliance with requirement in document.
ARC-443	ARC-488	The resilience capabilities for the Data Centre and the Data Centre network will be as specified in Data Centre Operations Service Description. (amended as per review 2/2/07)	DW	Minor	Walkthrough (or agenda item in walkthrough) will be arranged which will demonstrate this requirement has been met.
ARC-492	ARC-444	Any single failure within the Data Centres shall not cause loss of any of the Business Capabilities & Support Facilities	DW	Interrupt	Walkthrough (or agenda item in walkthrough) will be arranged which will demonstrate this requirement has been met. (TDN : N. Williams comment : A DW will demonstrate how the design meets this requirement)
ARC-492	ARC-470	Any single failure within the Data Centres shall not cause loss of any of the Business Capabilities & Support Facilities	ST	Interrupt	Tests to demonstrate business continuity are defined and contained within a detailed test plan, the successful completion of which shall determine the acceptance of this requirement. To clarify 'successful completion' success criteria will be identified in advance to test execution (TDN: N. Williams comment: Is it practical to simulate the potential single-failure over modes, to demonstrate no-loss? Clear success criteria for the tests (identified in advance to test execution) to be agreed and passed into PO DOORS.
ARC-493	ARC-445	Switchover to backup systems within the Data Centres and for the network connections within the Data Centres shall be automatic where defined for that service.	DW Add : CT	Interrupt	? Walkthrough (or agenda item in walkthrough) will be arranged which will demonstrate this requirement has been met. (TDN: N. Williams comment: A DW will demonstrate how the design meets this requirement.) Add : Failover test



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN

COMMERCIAL IN CONFIDENCE



ARC-493	ARC-471	Switchover to backup systems within the Data Centres and for the network connections within the Data Centres shall be automatic where defined for that service.	ST	Interrupt	Tests to demonstrate business continuity are defined and contained within a detailed test plan, the successful completion of which shall determine the acceptance of this requirement. To clarify 'successful completion' success criteria will be identified in advance to test execution (TDN: N. Williams comment: The functions should be testable. Clear success criteria for the tests (identified in advance to test execution) to be agreed and passed into PO DOORS)
ARC-446	ARC-446	The impact on Branch Users due to data centre failure and recovery shall be minimised. The principles of exception handling and recovery are as described in the document 'Agreed Assumptions on HNG-X Branch Exception Handling' referenced from Schedule B6.1.	DW Add : CT	Interrupt	Walkthrough (or agenda item in walkthrough) will be arranged which will demonstrate this requirement has been met. (TDN: N Williams comment: The DW should demonstrate that impact will be minimised, with some indication of what the impact will be. Is achievement of this 'minimum' testable) Add : Failover test
ARC-446	NEW000	The impact on Branch Users due to data centre failure and recovery shall be minimised. The principles of exception handling and recovery are as described in the document 'Agreed Assumptions on HNG-X Branch Exception Handling' referenced from Schedule B6.1.	CT in ST or RV	Interrupt	(TDN: N. Williams comment: Is this 'minimum' testable?)
SER-2145	SER-2145	Fujitsu shall establish and maintain a Configuration Management Policy and processes with a single organisational accountability for the completeness, integrity and accuracy of the Configuration Management system	DR	Minor	An entry on the compliance matrix will be supplied which cross references this requirement with the appropriate document and the relevant clause within the document will be supplied. The document will describe how this requirement will be catered for in the design
SER-2146	SER-2146	Fujitsu shall develop and publish defined processes for the identification and registration of Configuration items into a Configuration Management system	DR	Minor	An entry on the compliance matrix will be supplied which cross references this requirement with the appropriate document and the relevant clause within the document will be supplied. The document will describe how this requirement will be catered for in the design
SER-2147	SER-2147	Fujitsu shall define the operation of the Configuration Management processes and Post Office Ltd shall approve		Minor	An entry on the compliance matrix will be supplied which cross references this requirement with the appropriate document and the relevant clause within the document will be supplied. The document will describe how this requirement will be catered for in the design



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN

COMMERCIAL IN CONFIDENCE



SER-2155	SER-2155	Fujitsu Services and Post Office shall update the Service Review and Resilience Catalogue (SRRC) to reflect the new System Architecture.	DR	Visible	An entry on the compliance matrix will be supplied which cross references this requirement with the appropriate document and the relevant clause within the document will be supplied. The document will describe how this requirement will be catered for in the design.
SER-2156	SER-2156	The Business Continuity Framework shall be modified to reflect the System Architecture changes, and the move to shared Fujitsu Services data centres and Live configuration and Standby configuration.	DR	Minor	An entry on the compliance matrix will be supplied which cross references this requirement with the appropriate document and the relevant clause within the document will be supplied. The document will describe how this requirement will be catered for in the design.
SER-2157	SER-2157	The Business Continuity Plan shall be tested with the frequency agreed within the Business Continuity Framework (BCF)	SOO	Interrupt	There will be an appropriate section in a CCD or a document referred to by a CCD which contains a clause confirming this obligation. A compliance matrix will be supplied which cross references this requirement with the appropriate clause will be supplied.
SER-2158	SER-2158	Fujitsu shall support DR / Service continuity testing of other suppliers to Post Office as a separately chargeable service as per Horizon baseline and subject to receipt of Change Requests. Fujitsu Services will continue to co-operate with Post Office to develop appropriate plans.	SOO	Minor	There will be an appropriate section in a CCD or a document referred to by a CCD which contains a clause confirming this obligation. A compliance matrix will be supplied which cross references this requirement with the appropriate clause will be supplied.
SER-2159	SER-2159	Failures in the System and service, leading to a Major Incident, shall be managed and reported to Post Office using defined communication channels and processes identified in Service Management Service document.	SOO	Interrupt	There will be an appropriate section in a CCD or a document referred to by a CCD which contains a clause confirming this obligation. A compliance matrix will be supplied which cross references this requirement with the appropriate clause will be supplied.
SER-2178	SER-2201	Fujitsu shall report major Business Continuity Incidents to the Post Office promptly in accordance with the timescales and through contact point described in "POA Customer Service Incident Management Process Details"	RV	Interrupt	Business Continuity Tests are defined and contained within a detailed test plan, the successful completion of which shall determine the acceptance of this requirement. To clarify 'successful completion', success criteria will be identified in advance of test execution.
SER-2178	SER-2200	Fujitsu shall report major Business Continuity Incidents to the Post Office promptly in accordance with the timescales and through contact point described in "POA Customer Service Incident Management Process Details"	DR	Interrupt	An entry on the compliance matrix will be supplied which cross references this requirement with the appropriate document and the relevant clause within the document will be supplied. The document will describe how this requirement will be catered for in the design.



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN

COMMERCIAL IN CONFIDENCE



SER-2179	SER-2202	Management Information needed to support joint management and resolution of incidents, which may include third parties, will be made available to Post Office in accordance with the Incident Management Process and provisions in Service Management Service document.	DR		An entry on the compliance matrix will be supplied which cross references this requirement with the appropriate document and the relevant clause within the document will be supplied. The document will describe how this requirement will be catered for in the design.
SER-2179	SER-2203	Management Information needed to support joint management and resolution of incidents, which may include third parties, will be made available to Post Office in accordance with the Incident Management Process and provisions in Service Management Service document.	RV		Incident management processes and procedures exercised satisfactorily during Release Validation testing (TDN: N. Williams comment: Unless the RV will simulate 'an incident' how will the process be exercised? If MO is required, it should be passive with a criterion of 'no evidence of failure of the process'.)
BUS-3189	BUS-3189	Operations Control shall undertake a Business Continuity review and assessment with Fujitsu and will update the Service Review & Resilience Catalogue (SRRC) accordingly	Email confirmation	Visible	Evidence of update activity in the BAU change plan. All necessary materials updated
BUS-3190	BUS-3190	Operations Control shall help develop and agree with Fujitsu the testing for Business Continuity	DR (The process) MO Confirmation received	Minor	Evidence of Operations Control inclusion in the Business Continuity Test Plan Evidence of Operations Control participation in test planning
BUS-3197	BUS-3197	Operations Control shall help develop and agree the Problem & Incident Management process for HNG-X with the programme	Compliance to process	Minor	Process agreed
SVC-392		Branch Outages shall be within the levels and frequency described by the Service Level Targets in the Branch Network Service document.			
SCD 41		There shall be automated fail-over for all Service components e.g. using Active 1 and Active 2 components performing optimisation of both load balancing and performance, but each component 'capacity rated' to support the total balanced work-load.			

Key: CT = Component Test, RV = Release Validation Test, DR = Document Review, DW = Design Walkthrough, SOO = Statement of Obligation, SOF = Statement of Fact, ST = System Test

©Copyright Fujitsu Services Ltd 2007

Commercial In Confidence

Ref: TST/SOT/HTP/0006
 Version: 0.2
 Date: 19-Oct-07
 Page No: 107 of 108

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN
COMMERCIAL IN CONFIDENCE

B.1 Counter Architecture

V&I ID	Description	Quality Center ID	Acceptance Method
CNTR-0001	A process exists by which Users can raise hardware faults. NB As there is no change from Horizon it is assumed that this already exists and no test is required	NFUN- CNTR-0001	
CNTR-0002	Intermittent failures affecting internal subcomponents. Ensure written to event logs and then picked up by SYSMAN3		
CNTR-0003	A process exists by which users can raise software errors on the counter application.		

C.1 Branch Access Layer

(Enter the Delivery title or ID as the title to this paragraph)

(Enter the Quality Center derived listings of the tests to be run below.)

(The listings that follow are purely for **manual** tests. **No** automated tests are included.)



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN
COMMERCIAL IN CONFIDENCE



D.1 Online Services

E.1 Batch Applications

F.1 Branch DB

G.1 Counter Business Applications

H.1 Technical Network

UNCONTROLLED IF PRINTED



B Appendix – Automated Testing

There is no Automated Testing

UNCONTROLLED IF PRINTED



C Appendix – Platforms/Components

I.1 Servers

Where it is deemed that there is no requirement for re-testing this will be indicated by N/A.

ID	Platform Name	Tested in Section	Network Zone
1	A&L Banking Agents		
2	Alarm Point		
3	Antivirus Server		
4	APOP Web Server		
5	Atalla HSM Appliance (HSM)		
6	Audit Server		
7	Audit Workstation		
8	BAL Server		
9	Boot Platform (HNG-X)		
10	Branch Change Management System Server		
11	Branch Configuration Database Server		
12	Branch Database Server - Main		
13	Branch Database Server - Standby		
14	Branch Support Server		
15	CAPO Banking Agents		
16	Certificate Server (CAN)		
17	Cisco Works		
18	Connect Direct Simulator		
19	ConnectDirect Gateway		
20	DCS & ETS Authorisation Server		
21	Debit Card Management Server		
22	Dimensions Signing Server (DSS)		
23	DNS Server (Primary)		
24	DNS Server (Secondary)		
25	Domain Controllers - Active Directory		
26	DVLA Web Server		
27	EMC ECC Server		
28	EMC Remote Support Gateway		
29	Enterprise Boot Server		
30	Firewall Security Manager		
31	FTMS EDG Local		



HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST PLAN



COMMERCIAL IN CONFIDENCE

32	FTMS EDG Remote		
33	FTMS TIP Local		
34	FTMS TIP Remote		
35	Generic Proxy		
36	Help Desk Web Server		
37	HNG-X NT4 Counter		
38	HNG-X XP Counter		
39	IN Pad Test Workstation (PPT)		
40	KMNG Server (KMN)		
41	KMNG Workstation (KSN)		
42	Link Authorisation Server		
43	Maestro Server		
44	MIS Client		
45	MIS Support Workstation		
46	Money gram Web Server		
47	Network Management Server		
48	Network Persistent Store		
49	NMS - Packet Capture		
50	Online Training Web Server		
51	PAF Web Server		
52	Performance Management Server (SPN)		
53	PIN Pad Key Generation Workstation (SKG)		
54	PIN Pad Proving Workstation (PPW)		
55	Radius Branch Router Management		
56	Radius Core Router Management		
57	RDMC Workstation		
58	RHEL Backup Server		
59	Router Operational Support Server		
60	Solaris Backup Server		
61	Solaris Host		
62	SSC Server		
63	SYSMAN's Enterprise Database Server		
64	SYSMAN's Enterprise Managing Server		
65	SYSMAN's Enterprise User Interface Server		
66	TES Web Server		
67	Tivoli Workload Scheduler		
68	Training PIN Pad Loading Workstation (TPP)		

HNG-X: ITU V&I BUSINESS CONTINUITY HIGH LEVEL TEST
PLAN

COMMERCIAL IN CONFIDENCE

69	Virtual Server Host		
70	Windows Backup Server		

J.1 Storage

Primary Storage Elements:

Storage Element	Number per data centre	Tested in Section
EMC Symmetrix DMX3	2	
EMC Clariion CX3-80	1	
EMC Celera NAS	1	
EMC Centera CAS System	2	
EMC EDL Virtual Tape Library	1	
Cisco 9509 Directors	2	
Fujitsu-Siemens Fibre Cat TX24 LTO3 autoloader	1	

K.1 Network

Network Element	Location	Tested in Section