POL00399072

[TITLE * MERGEFORMAT]

[DOCPROPERTY Subject * MERGEFORMAT]

Document Title:	[TITLE * MERGEFORMAT]
Document Type:	Architecture (ARC)
Release:	HNG-X Release 1
Abstract:	This document describes how HNG-X will meet requirements for performance and capacity, availability, and disaster recovery.
Document Status:	APPROVED This document contains sections that have been identified to POL as comprising evidence to support the assessment of named Acceptance Criteria by Document Review. These sections must not be changed without authority from the FS Acceptance Manager.
Author & Dept:	David Chapman
Internal Distribution:	
External Distribution:	POL Document Management
Security Risk Assessment Confirmed	YES. See section 0.9, Security Risk Assessment.

Approval Authorities:

Name	Role	Signature	Date
David Court	Programme Manager		
Amit Apte	CTO		

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

Documents are uncontrolled if printed or distributed electronically. Please refer to the Document Library or to Document Management for the current status of a document.

POL00399072

0 Document Control

0.1 Table of Contents

[TOC \O "1-3" \H \Z \T "POA APPENDIX HEADING 1,1,POA APPENDIX HEADING 2,2"]

0.2 Figures and Tables

[TOC \t "Caption,3"]

0.3 Document History

Version No.	Date	Summary of Changes and Reason for Issue	Associated Change - CP/PEAK/PPRR Reference
0.1	15-Nov-2006	Draft for review	
0.2	28-Nov-2006	Draft for review	
1.0	09-May-2007	For Approval	
1.1	27-June-2008	Draft for review	
1.2	1-Dec-2008	Draft for Acceptance by Document Review. Changes were made to the document for Acceptance by Document Review with the insertion of the Section containing the table of cross references for Acceptance by Document Review.	
1.3	14-May-2009	Draft for review; update to cover migration strategy	
1.4	12-Nov-2010	Revisions for HNG-X Release 1 baseline. Draft for review	
1.5	9-Mar-2011	Draft for review	
1.6	23-Mar-2011	Addition of further info to future changes section	
1.7	25-Mar-2011	Amendment of future changes section	
2.0	08-Apr-2011	Approval version	

0.4 Review Details

Review Comments by :	
Review Comments to :	David Chapman
Mandatory Review	
Role	**Name**
HNG-X Solution Design	Steve Evans / Adam Spurgeon
HNG-X Information Governance	Bill Membery
HNG-X Infrastructure Design	Alex Kemp
Capacity & Configuration Manager	Mark Brosnan
CISO	Ian Howard
Optional Review	
Role	**Name**

POL00399072

HNG-X Development	Graham Allen
HNG-X Test Design	
Service Network	Andrew Hemingway
Application Services	Peter Thompson
HNG-X System Test	Sheila Bamber
HNG-X SV&I Manager	Chris Maving
POL Test Manager	James Brett (POL)
HNG-X Testing	Debbie Richardson
Service Operations	Tony Atkinson
SSC	Steve Parker
Business Continuity	Adam Parker
Former Head of Service Change & Transition for Release 1 – retained for R1 Baseline review purposes only.	Graham Welsh
Integration Team Manager	Vijesh Pandya
HNG-X System Qualities	
HNG-X Design (Reference Data)	Duncan MacDonald
HNG-X Architect (Platforms and Storage)	Jason Clark
HNG-X Architect (Support Services)	Sarah Selwyn
HNG-X Architect (Customer Services)	Pete Jobson
HNG-X Architect (Branch Database)	Andy Beardmore
HNG-X Architect (Counter)	Andy Thomas
HNG-X Architect (Network)	Mark Jarosz/Dave Haywood
HNG-X Architect (Online)	Andy Williams
Operational Security	Donna Munro
Core Services & HNG-X System Qualities	Ed Ashford
Core Services	Andy Gibson
HNG-X Infrastructure System and Estate Management	Patrick Carroll
HNG-X Infrastructure System and Estate Management	Ian Bowen
Infrastructure Delivery Manager	Martin Brett
CTO	Amit Apte
POL Design Authority	Ian Trundell / Peter Stanley (as appropriate)
Issued for Information – Please restrict this distribution list to a minimum	
Position/Role	**Name**
Acceptance Manager	David Cooke

0.5 Acceptance by Document Review

The sections in this document that have been identified to POL as comprising evidence to support Acceptance by Document review (DR) are listed below for the relevant Requirements:

POL00399072

POL NFR DR Acceptance Ref	Internal FS POL NFR Reference	Document Section Number	Document Section Heading
ARC-475	ARC-450	3.1.2.4	Transaction Totals
ARC-475	ARC-450	11.3.1	Targets
ARC-442	ARC-491	4	HNG-X Capacity Management Service
MIG-3099	MIG-3240	4.2.1	Pro-active & Re-active Analysis
MIG-3099	MIG-3240	4.3.1	Migration
SER-2155	SER-2155	3.2.3.3	Single point of failure
SVC-840	SVC-794	4.4	Modelling
SVC-847	SVC-800	3.1.2	Throughput
SVC-849	SVC-802	3.1.2	Throughput
SVC-806	SVC-806	4.1.2	Capacity Operational Review Forum
SVC-859	SVC-812	4.1.1	Monthly Reporting
SVC-861	SVC-813	4.1.1	Monthly Reporting
SVC-864	SVC-814	4.1.1	Monthly Reporting
SVC-864	SVC-814	4.1.2	Capacity Operational Review Forum
SVC-851	SVC-803	3.1.3.3	HNG-X System Limits have not been set

0.6 Associated Documents (Internal & External)

Reference	Version	Date	Title	Source
PGM/DCM/TEM/0001 (DO NOT REMOVE)			Fujitsu Services Post Office Account HNG-X Document Template	Dimensions
			Branch Database Prototype Report on Linux	Dimensions
			Branch Database Prototype Summary Report on Solaris	Dimensions
SVM/SDM/SD/0003			Data Centre Operations Service	Dimensions
ARC/APP/ARC/0001			HNG-X Reference Data Architecture	Dimensions
ARC/APP/ARC/0002			HNG-X Integration Architecture	Dimensions
ARC/APP/ARC/0003			HNG-X Counter Architecture	Dimensions
ARC/APP/ARC/0004			HNG-X Branch Access Layer Architecture	Dimensions
ARC/APP/ARC/0005			HNG-X Online Services Architecture	Dimensions
ARC/APP/ARC/0007			HNG-X Batch Application Architecture	Dimensions
ARC/APP/ARC/0008			HNG-X Branch Database Architecture	Dimensions
ARC/APP/ARC/0009			HNG-X Architecture – Counter Business Applications	Dimensions
ARC/GEN/REP/0001			HNG-X Glossary	Dimensions

POL00399072

Reference	Version	Date	Title	Source
PA/PER/033			HNG-X Capacity Management and Business Volumes	PVCS
ARC/SEC/ARC/0003			HNG-X Security Architecture	Dimensions
DES/PER/HLD/0001			HNG-X Resilience and DR HLD	Dimensions
DES/PER/HLD/0002			HNG-X Capacity and Performance HLD	Dimensions
DES/PER/HLD/0003			HNG-X Branch Trading Resilience HLD	Dimensions
ARC/PER/REP/0001			HNG-X Performance Model	Dimensions
REQ/CUS/BRS/1049			HNG-X Post Office Non-Functional Requirements for Release 1	Dimensions
REQ/CUS/BRS/1050			HNG-X Customer Services Release 1 Requirements	Dimensions

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

0.7 Abbreviations

Abbreviation	Definition
A&L	A&L Santander
AVG	Average
ACDB	Access Control Database
AP	Automated Payment
AP-ADC	Automated Payment – Advanced Data Capture
APOP	Automated Payment Out-Pay
ADSL	Asymmetric Digital Subscriber Line
BAL	Branch Access Layer
CAPO	Card Account Post Office. An external client providing banking services.
DC	Data Centre
DCS	Debit Card Service
DNS	Domain Name Server
DR	Disaster Recovery
DVLA	Driver and Vehicle Licensing Authority
EDG	Electronic Data Gateway
EPOSS	Electronic Point of Sale Service; HNG-X service that supports Retail functions in Branches
ETS	Electronic [Phone card] Top-up Service
ETU	Electronic Top-Up

POL00399072

Abbreviation	Definition
FAD	Finance Accounts Division, part of Post Office Ltd
FTMS	File Transfer Management Service; HNG-X process that provides configurable file transfer services between Horizon and Post Office Ltd. Clients. Services available include data compression and encryption
HNG-X	Horizon Next Generation, Plan X
ISDN	Integrated Services Digital Network
JVM	Java virtual machine
LAN	Local Area Network
MGRM	MoneyGram
MSVS	Microsoft Virtual Server
MTAS	MID / TID Allocation System.
NBS	Network Banking Service – one of the A&L, CAPO or LINK Authorisation Services
NBX	Network Banking Engine
NPS	Network Banking Persistence Service
OCMS	Outlet Change Management System; Fujitsu System for scheduling estate management changes
PAF	Postal Address File
POA	Post Office Account
POL	Post Office Ltd
POL-FS	Post Office Ltd Finance System. SAP based system providing financial accounting for the branch based business.
POL-MIS	Post Office Ltd Management Information System
POLSAP	New single SAP system hosting both POL-FS and SAPADS
RAC	Real Application Cluster
REC	Reconciliation data files
SAN	Storage Area Network
SAPADS	SAP Automated Distribution System. POL's Advanced Distribution System (based on the SAP package) that interfaces to LFS
SAS	Secure Access Server
SLT	Service Level Target
SPARC	Sun Microsystems Processor (Scalable Processor Architecture)
SPOF	Single Point Of Failure
SQL	Structured Query Language; language commonly used to access Relational Database systems
SSL	Secure Socket Layer
TES	Transaction Enquiry Service
TIP	Post Office Ltd's Transaction Information Processing system

POL00399072

Abbreviation	Definition
TLCM	Telecom
TPS	Transaction Processing Service; Horizon service that formats data for transmission to TIP
UDP	User Datagram Protocol
VIP	Virtual IP
VOCALINK	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.
VPN	Virtual Private Network
VSAT	Very Small Aperture Terminal
WAN	Wide Area Network

0.8 Glossary

See also [REF ARCGENREP0001 \h].

Term	Definition

0.9 Changes Expected

Changes
Update chart in section 2.0 and where possible introduce new charts throughout to help illustrate document
Add an appendix which shows all data required for measurement and where it is obtained from
Update Network Section 7 (incl chart in section 7.1) once new Network Design Complete
Update Requirements with latest references and descriptions from SRS or consider replacing with a cross-reference to another source of this information.
Update Estate Management section.
Add section for Confidentiality and Integrity
Add section for Audit
Clarify section 7.3.1 in line with the Branch Exceptions CCD
Further coverage required of resilience from risks of corruption to platform image

POL00399072

0.10 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.

0.11 Security Risk Assessment

Security risks have been assessed and it is considered that there are no security risks relating specifically to this document.

POL00399072

1 Management summary

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

- Capacity/Performance. Adequate capacity is provided throughout all components of the HNG-X system, both for real time transactional performance and data persistence. This has been confirmed through a combination of analysis of live usage, knowledge of Horizon solutions, prototype testing, performance intensive simulations and modelling. Additionally the HNG-X Capacity Management Service will proactively and reactively investigate system performance and business volumes, where required in conjunction with Post Office Ltd, to monitor the solution and validate that it is running as expected from the perspective of business volumes and system component utilisation.
- 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 have been 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 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 is persisted on the storage both at the Primary and Secondary data centres. Hence under normal operation for these systems, the DR storage will always hold an exact copy of the storage at the Primary data centre.

POL00399072

2 Introduction

2.1 Purpose of the document

This document describes both how the HNG-X Architecture will meet requirements for performance and capacity, availability and disaster recovery, as well as the constraints on designers to meet those requirements. This document does not cover specific security issues as, for example, should a security issue arise that requires DR to be invoked, the standard DR process would be followed regardless of the cause. Additionally capacity and performance has to be provided based on typical utilisation and, for example, cannot be provided based on any extra load from a potential security breach (although the system design should minimise any impact). The security to provide a mechanism to prevent such issues occurring is covered elsewhere.

2.2 Structure of the document

Section 3 describes targets for performance, capacity, availability and disaster recovery. It explains how the targets are derived from the HNG-X requirements. It breaks down high-level targets to targets for different parts of HNG-X.

Section 4 describes the Capacity Management Service. It explains its roles and responsibilities, and how these will be achieved.

Section 5 gives an overview of how the targets from Section 3 will be met by the HNG-X solution. Where the solutions are going to be provided by the Systems Qualities team, the architectural details of these solutions will be explained.

Sections 6 to 14 provide more detail of how each primary part of the solution will meet the targets. It covers the counter, network, branch access layer, online services, branch database, host batch services, Hydra, the Post Office Financial Systems (POL-SAP) and Support Services.

Section 15 lists the HNG-X Release 1 requirements from Post Office and Fujitsu Customer Services that are relevant to performance, capacity, availability or disaster recovery.

2.3 Background to HNG-X

The Horizon systems supported Post Office branches, enabling them to trade and providing management facilities. Horizon included counter systems at the Post Office branches, and data centre systems.

The Horizon Next Generation Plan X (HNG-X) systems are a replacement for the Horizon systems. Many of the Horizon data centre business applications are taken forward into HNG-X (possibly enhanced), but the counter business applications are replaced. In HNG-X, some functions that are carried out on the Horizon counter systems are replaced with online functions that are carried out at the data centre. Some data that is stored on the Horizon counter is replaced with data stored at the data centre in a new Branch Database. The HNG-X systems (including the migrated Horizon systems) will be hosted at new data centres. HNG-X makes extensive use of 2 virtualisation technologies – vBlades (Xen based Linux and Windows virtual machines running on Egenera Blade frames) and to a lesser extent Microsoft Virtual Server (MSVS).

The Horizon systems retrieved and sent branch data using a messaging product called Riposte. HNG-X systems will communicate using the equivalent data to and from the Branch Database.

POL00399072

The Post Office Financial System (POL-FS) ran at the same data centres as the Horizon systems, and will be transitioned to the HNG-X data centres and integrated with other SAP based services as part of the POL-SAP solution.

HNG-X introduces a new Branch Access Layer. This provides access to the Branch Database and also routes requests on to other online systems. The Branch Access Layer runs on multiple application servers. Intelligent network devices will be used to distribute requests to the Branch Access Layer, and to distribute requests from the Branch Access Layer to certain other online systems.

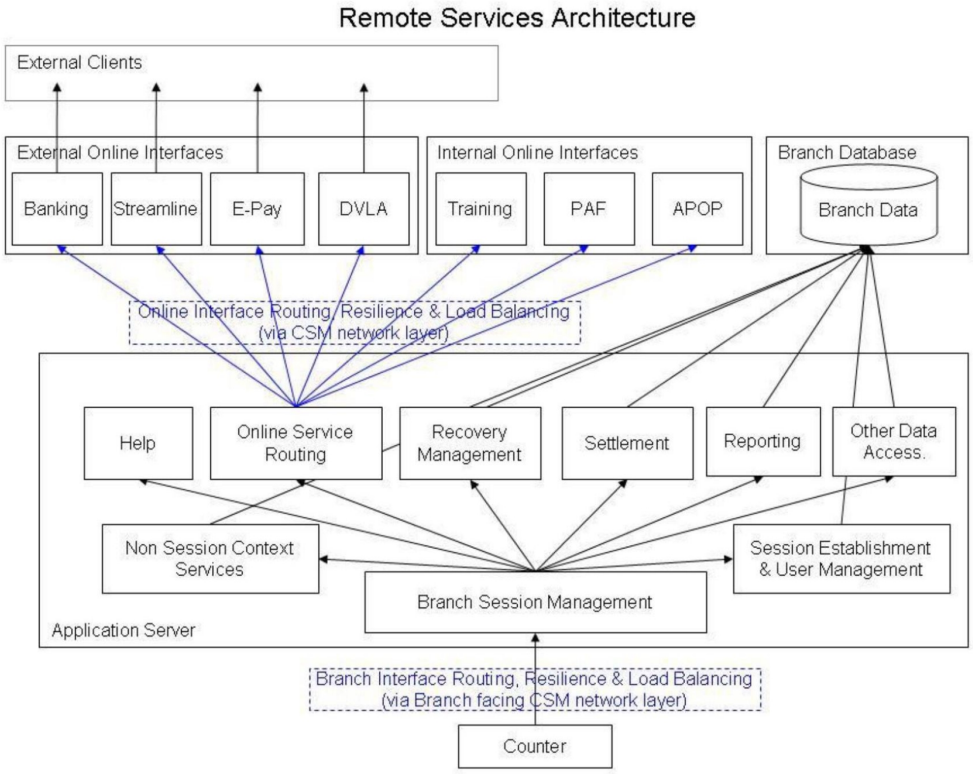

Figure [SEQ Figure * ARABIC] – Remote Services Architecture

A chain is only as strong as its weakest link. To meet overall requirements for performance, capacity, availability and disaster recovery, every part must meet requirements. This document shows how each of the following parts of HNG-X meets requirements:

- HNG-X Counter
- Network (including VPN & Radius)
- Branch Access Layer
- Online Services
- Branch Database
- Host Batch Services

POL00399072

- Hydra
- POLSAP (including Post Office Financial System and SAPADS)
- Support Services

POL00399072

3 Targets

3.1 Performance and capacity

3.1.1 General comments

3.1.1.1 Confidence in actual transaction rates

The performance and capacity figures are based on measurements of actual peak transaction rates seen over 5 years.

These figures are unlikely to rise for the following reasons:

- At peak times, the transaction rate is limited by the number of counters and the speed at which customers can be served.
- The number of branches and counters has declined since baselining. (Ongoing maximum numbers of branches and counters per year as defined in the HNG-X version of the Horizon Capacity Management and Business Volumes document [PA/PER/033] show a continual decrease. Therefore volumes will not increase.)
- Since initial baselining overall business volumes have declined with the removal of products (TV licence, internet DVLA, closure of CAPO to new customers), as well as the decline of certain services (e.g. APS).

However overall volumes are unlikely to decrease significantly due to a combination of new services (Telecoms/Kahala), more extensive use of existing services (PAF, APOP, MGRM, DCS), and changing business patterns (e.g.Ebay shipping has increased the number of mails transactions).

On balance, these figures can therefore be confidently used as a basis for peak volume calculations.

3.1.1.2 Exceeding contracted limits

If contracted transaction volumes are exceeded, the systems are not obliged to meet service level targets (SLTs).

However, the systems will be designed to meet transaction volumes that are higher than the contracted volumes. Even if the volumes exceed these design limits, the new HNG-X systems will be designed to handle excess volumes in an appropriate way where possible, to avoid overall failure. This would enable normal service to be resumed as soon as transactions volumes return to their contracted volumes.

3.1.1.3 Calculating data volumes

Data volumes are calculated based on the expected peak transaction volumes for the period that the data is retained.

For example, volumes for data that are retained for one day only are calculated using the peak daily transaction volumes. Volumes for data that are retained for one month is calculated using the peak monthly transactions volumes, not 30 times the peak day.

POL00399072

3.1.1.4 Dealing with peak loads

The online HNG-X systems have two different types of peak load: a sales peak (mostly comprising of banking transactions, typically occurring early in the morning), and a reporting peak (typically occurring late in the day); although other peaks also exist – evening mailing, lunchtime for end of month DVLA, as well as various seasonal peaks. The volume of reporting transactions is negligible during sales peaks when most counters and staff are serving customers, and the volume of sales transactions is very low during reporting peaks. All systems must be able to handle both peaks, but do not have to handle both full sales and reporting transaction volumes concurrently. A further smaller peak occurs during counter logon (typically occurring early in the morning as branches open), however this also needs to be considered due to the high volumes and particular characteristics of logon transactions.

3.1.1.5 Allowing for actual per second peaks

The contracted volumes and design limits are shown and measured as transactions per second averaged over 5 minutes of peak load. The system must be able to sustain this transaction rate.

However the actual number of transactions varies from second to second (including over this 5 minute average peak), and is of course at times higher than the average. At peak times where volumes approach the contracted volumes, some components may have to process higher actual per second peaks than their contracted volumes measured as the average over 5 minutes. In this scenario, the contingency provided by the design limit should provide adequate capacity to handle this per second variation. With the extensive knowledge and experience of post office business volumes, which are very consistent, it is possible to accurately predict expected volumes and therefore it is considered very unlikely that we will receive volumes that break the contracted volumes. However as it is necessary for the solution to be capable of running up to the design limits (with no SLTs), the figures for individual components show an additional contingency (found from analysis of actual volumes) to indicate a possible per second peak while running at the design limit. This value is not contractual and is for illustration purposes only (although system components should be designed to cope with these per second peaks varying from the sustained design limit).

3.1.2 Throughput

3.1.2.1 Business transactions

Requirement SVC-800 states that the system should support the volume of business transactions defined in the *HNG-X version of the Horizon Capacity Management and Business Volumes* document [PA/PER/033]. For existing services these volumes are the same as those supported by the Horizon solution.

Volumes are defined for each service.

For each service, the volumes are defined for the peak month, peak week, peak 2 days, peak hour, and peak 5 minutes. The peak 5 minute volume is presented as an average transactions per second (TPS) for that period (e.g. volume divided by 300 (the number of seconds in 5 minutes)).

Two sets of figures are presented in the tables below. The first table shows what HNG-X is contractually obliged to meet. The second table shows what HNG-X is designed to meet.

The first table additionally shows the peak 5 minute transaction rate actually recorded on the Horizon system, also presented as an average TPS over that period

POL00399072

Contracted Volumes

Service	Peak Month	Peak Week	Peak 2 Days	Peak Day	Peak Hour	Peak 5 Min avg TPS	Baseline 5 Min avg TPS
EPOSS	139,952,283	41,868,562	20,001,597	11,081,629	2,010,680	622	320
APS	20,320,847	5,755,944	2,813,111	1,515,787	280,421	90	139
NBS	34,200,639	9,234,172	4,408,320	2,553,600	618,240	204	176
DCS	6,315,000	1,897,152	848,924	432,638	119,052	33	22
ETU	2,293,203	561,415	213,783	121,200	17,254	5	5
DVLA	3,964,264	2,288,121	1,128,761	737,774	101,996	30	34
PAF	11,969,806	4,002,430	1,883,664	1,100,788	162,246	46	18
Bureau	1,351,863	439,901	203,351	115,163	16,601	5	5
Settlement	105,205,376	35,319,442	15,620,323	8,602,518	1,230,160	344	288
APOP (Generic online services)	8,200,000	2,050,000	1,045,000	660,000	120,000	39	9

Design Limits

Service	Peak Month	Peak Week	Peak 2 Days	Peak Day	Peak Hour	Peak 5 Min avg TPS
EPOSS	167,942,740	50,242,275	24,001,916	13,297,954	2,412,816	746
APS	24,385,017	6,907,133	3,375,733	1,818,944	336,505	108
NBS	41,040,766	11,081,007	5,289,984	3,064,320	741,888	245
DCS	7,578,000	2,276,582	1,018,709	519,165	142,863	39
ETU	2,751,844	673,699	256,540	145,440	20,705	6
DVLA	4,757,116	2,745,745	1,354,513	885,329	122,396	36
PAF	14,363,767	4,802,916	2,260,397	1,320,945	194,695	55
Bureau	1,622,235	527,880	244,020	138,195	19,920	6
Settlement	126,246,451	42,383,331	18,744,387	10,323,021	1,476,192	413

POL00399072

APOP (Generic online services)	8,200,000	2,050,000	1,045,000	660,000	120,000	39

3.1.2.2 Online business services

Some of the business transactions are routed to additional online services. This applies to NBS, DCS, ETU, DVLA, PAF and APOP (APOP counter volumes include MGRM, Telecoms, Kahala etc); it does not apply to EPOSS, APS and Settlement which just update the Branch Database.

These online business services are subject to additional overall contracted volumes and design limits for the combined total over any contractual period (this is less than the sum of all the services). These are shown below.

Online Business Services	Peak month	Peak week	Peak 2 days	Peak day	Peak hour	Peak 5 min avg TPS
Contracted Volumes	46,101,471	12,236,079	6,228,817	3,556,145	739,012	234
Design Limit	55,321,766	14,683,294	7,474,580	4,267,374	886,814	281

This can be broken down further into the individual services (and in the case of NBS, to individual banking services) and indicates the capacity each individual service has to be capable of processing.

Service	Design Limit (Peak 5 minute avg TPS)	Allowance for actual per Second Peak
Banking (A&L)	14	19
Banking (CAPO)	192	250
Banking (LINK)	29	38
Streamline (DCS)	39	51
E-Pay (ETU)	6	8
DVLA	36	47
Training	5	7
PAF	55	72
APOP (Generic online services)	42	55

POL00399072

3.1.2.3 Administrative and back office transactions

Requirement SVC-800 states that the system should support the volume of administrative and back office transactions defined in the *HNG-X version of the Horizon Capacity Management and Business Volumes* document [PA/PER/033].

This includes help text retrieval, log on/off (this also covers other user management operations such as edit user, new user), and reports.

The table below shows the contracted volume and design limits for each of these transactions, as TPS for the peak 5 minutes for each transaction type.

For each service the table also shows the peak 5 minute transaction rate (except for help text, which has no equivalent in the existing system), as well as the peak 5 minute transaction rate recorded during the peak transaction period, recorded on the existing live system from available sample data.

Only TPS figures are given. These transactions do not create additional persistent data of consequence, and so there is no need to know their volumes over longer periods.

The per second peak values for reporting and user management are significantly higher than for other services; this is because the reporting and user management peak periods are relatively short (administration tasks such as User management and SU management do not have separate volumes as they are not significant in terms of transaction volumes).

Service	Contracted peak 5 min avg TPS	Design limit peak 5 min avg TPS	Live 5 min avg TPS at overall peak period	Live 5 min avg TPS
Help Text	50	60	N/A	N/A
User Management (Log On/Off)	100	120	5	77
Reports	125	150	5	109

Reporting volumes here include back office administration tasks.

Help Text not used – allocation shown for illustrative purposes.

3.1.2.4 Transaction totals

The table below shows the overall contracted volumes and design limits for all Online Services (i.e. the total of all business transactions, administrative and back office transactions).

These figures are lower than the sum of the parts and have been jointly formulated by Post Office Ltd and Fujitsu Services (from investigation of actual live peaks and analysis of business volumes and trends). Not only do administrative operations not typically occur at the same time as the peak business transactions (as discussed in section 3.1.2.3), but the various business transactions peaks do not typically occur at the same time. Additionally routed online services are already subject to overall contracted volumes and design limits that are lower than their sum (see section 3.1.2.2).

For illustrative purposes the sum of the maximum possible contracted volumes is shown in addition to the actual overall Contracted Volume and Design limit for all Online Services.

POL00399072

Service	Peak month	Peak week	Peak 2 days	Peak day	Peak hour	Peak 5 min avg TPS
Total Online contracted volumes	328,532,165	101,713,363	47,781,332	26,645,231	4,567,828	1,676
All Online contracted volume	257,379,880	74,360,052	34,642,912	19,371,437	3,356,988	940
All Online design limit	308,855,856	89,232,063	41,571,494	23,245,724	4,028,386	1,128

In addition to these business, administrative and back office transactions, the counters perform reference data synchronisation transactions, which take place whether or not any session is active. Adding these to the design limit, and adjusting for actual per second peaks, gives the following overall transaction levels for the Branch Access Layer.

	Design limit (Peak 5 min avg TPS)	Allowance for actual per second peak
Branch Access Layer overall limit	1128	1466

This can be broken down per service running on the Branch Access Layer application servers.

Service on Application Server	Design limit (Peak 5 minute avg TPS)	Allowance for actual per second peak
Reference Data*	60	78
Help **(Unused)**	60	78
Online Service Routing	281	365
Recovery Management*	300	390
Basket Settlement*	393	511
Reporting*	150	300
Branch Data Access*	10	13
Session Establishment & User Management*	120	200

Not all transactions cause significant load (e.g. updates and inserts) on the Branch Database, online service routed transactions only authenticate session (those that do are marked with * in the chart above). Allow for 300 reports a second whilst other transaction levels are low – else 250. The figures for the branch database are:

POL00399072

	Design limit (Peak 5 minute TPS)	Allowance for actual per second peak
Branch Database overall limit	1068	1388

The branch database should perform with a maximum response time between 50 and 100ms for a typical transaction when running at its peak design limit of 1068 TPS.

3.1.3 Response times

3.1.3.1 General response times

Requirement SVC-815 states that banking and related transactions shall have an end-to-end response time as described in the *Data Centre Operations Service Document*. SVC-816 says the same for basket settlement. Requirement TR455 states that the new system shall have response times comparable with the existing systems.

These requirements translate to the following response time targets (SLTs):

- Network Banking transactions will take on average 2.5 seconds or less within the total of the HNG-X systems and infrastructure. This is the total time to and from the counter, excluding the time in the banks' infrastructure and systems.
- Settlements will take on average 2 seconds or less.
- No Settlement within the 95th Percentile will take 7 seconds or more.

These targets are the same as those for the Horizon systems where applicable.

The online services used by Horizon (Network Banking, DCS, ETU, DVLA, etc) are carried forward into HNG-X, the main changes only for PCI compliance (which would have been necessary for Horizon anyway). However any possible performance impact from PCI should be offset by improved hardware. It is therefore reasonable to assume that these services will at the very least continue to provide adequate response times. Furthermore performance based simulations have shown the overhead of vBlades to be minimal and therefore based on hardware comparisons (both through specInt ratings and System Qualities Team simulations) and the existing utilisation of these systems in Horizon, the HNG-X systems will be significantly faster. For NPS there will be a further improvement both from the faster and more efficient EMC Symmetrix DMX3 and the fact the SRDF latency time will be substantially reduced as the distance between the 2 data centres is considerably shorter than between Bootle and Wigan.

3.1.3.2 Target response times for new services have not been set

Requirement SVC-817 states that branch administration and back office transactions should meet any targets defined in the *HNG-X version of the Horizon Capacity Management and Business Volumes* document [PA/PER/033] from benchmarking. SVC-818 says the same for stock unit and branch reporting transactions. SVC-819 says the same for help pages.

No targets have been set for these new services in the current version of the *HNG-X version of the Horizon Capacity Management and Business Volumes* document [PA/PER/033], so these requirements are not currently measurable.

No measurements of the response times for the equivalent functions within the Horizon systems have yet been made. However it is assumed that response times will be similar for operationsin the Horizon solution.

POL00399072

3.1.3.3 HNG-X System Limits have not been set

[SVC-803][SVC-804] System will support parameter and system volumes as defined in the *HNG-X version of the Horizon Capacity Management and Business Volumes document* [PA/PER/033].

No System Limits have been set for HNG-X in the current HNG-X version of the *Horizon Capacity Management and Business Volumes* document [PA/PER/033], so these requirements are not currently meaningful. These system limits will however be expected to enable similar capabilities to the Horizon solution.

3.1.4 Data storage volumes

Data storage volumes are calculated from the maximum transaction rates and persistence periods. Maximum transactions rates for data for different databases are shown below.

Service	Design Limit Per Week	Design Limit Per Month
Network Banking (DRS)	11,081,087	41,040,766
DCS (DRS)	1,517,721	5,052,000
ETU (DRS)	673,699	2,751,844
Branch Database	77,986,779	269,932,941

3.1.5 Network traffic

Requirement SCD-139 states HNG-X networking capabilities will ensure that the bandwidth available to every branch and counter is sufficient to handle the delivery of any data without compromising other branches or counters.

The network has two peak requirements: peak transaction period and peak reporting period. The expected HNG-X requirements for each are shown below. Both upstream and downstream network traffic values are shown. It is assumed that extra contingency will be provided within the network infrastructure to allow for any fluctuation.

Peak transaction period

Peak transaction period	Total upstream bits per second	Total downstream bits per second
Total for all branches	24,915,660	9,796,693
Average per branch	1,747	687
Average per counter	713	280
Average per busy branch	3,433	1,350
Average per busy counter	891	350
Likely workload for largest site (20 busy counters)	17,826	7,009

Peak reporting period

POL00399072

Peak reporting period	Total upstream bits per second	Total downstream bits per second
Total for all branches	13,983,078	7,519,838
Average per branch	981	527
Average per counter	400	215
Average per busy branch	1,927	1,036
Average per busy counter	500	269
Likely workload for largest site (20 busy counters)	10,004	5,380

Furthermore we have added confidence based on the comparison from Horizon to HNG-X:

- 50% of traffic on the Branch to Data Centre Network was made up of Riposte Marker Interval; this will not exist in HNG-X, which will only handle the equivalent of the other 50% of data.
- There is no concept of EPOSS/APS messages in HNG-X (from counter to BAL/BranchDB) and as such EPOSS/APS are now properties of a basket settlement which is compressed before being sent to the data centre. This reduces the extra overhead that was required in Horizon for each message (headers, footers and generic information) and the amount of data. However extra basket items will increase the size of the settlement message in HNG-X.
- Horizon data showed there was contingency within the bandwidth used by the Horizon solution at peak periods; therefore with reduced volumes for HNG-X (due to the reasons above), we can be confident of sufficient bandwidth.

However under Horizon, reports were processed on the counter using local data; under HNG-X the request for this data will be made on the Branch Database at the datacentre, with the report data using the branch to data centre network to be passed to the counter. Whilst reporting peaks are quite high and the volume of data per report could be relatively high, this will be mitigated by compression of communications between the counter and data centre and the removal of the marker interval traffic.

3.2 Availability

3.2.1 General availability

General requirements for availability are summarised below:

- Where choices occur, the architecture shall take support for serving of customers as the priority, unless agreed otherwise with Post Office (ARC-411).
- Data Centre and Network resilience capabilities will be as specified in Data Centre Operations Service Description (ARC-443).
- Branch Outages shall be within the levels and frequency described by the Service Level Targets in the Branch Network Service document (SVC-792).
- Central Systems shall be available to provide the Core Solution during Core Hours according to Service Level Targets set out in the Data Centre Operations Service document (SVC-779).

POL00399072

- Central Systems shall be available to provide the Core & Banking Solution during Core Hours according to Service Level Targets set out in the Data Centre Operations Service document (SVC-780).
- Central Systems shall be available to provide the Core & all Other Services during Core Hours according to Service Level Targets set out in the Data Centre Operations Service document (SVC-781).
- The availability and recovery times of POL-FS are defined in the Data Centre Operations Service document (SVC-782 and SVC-783).

The availability targets assume the contracted transaction volumes given in section 3.1.

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%

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, as illustrated in the chart below:

POL00399072

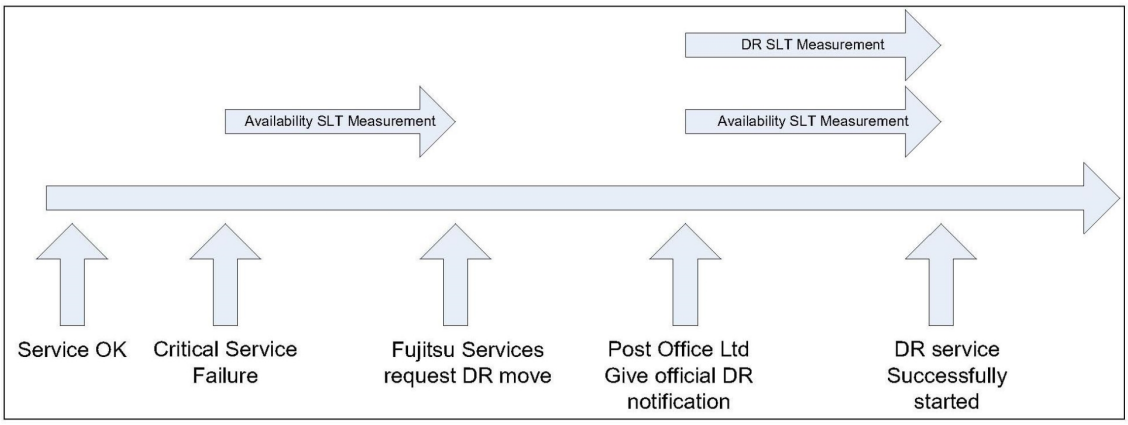

Figure [SEQ Figure * ARABIC] – SLT measurement during DR operation

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.

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.

3.2.2 Availability principles

3.2.2.1 No data loss

Paragraph 5.1 of RM/CDE/003 v 0.4 states that the current recovery models are assumed.

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

3.2.2.2 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.

POL00399072

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.

3.2.2.3 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 the main physical components within the system (BladeFrame, Discrete Systems, Core Switch & Storage), indicating risks, repairs, contingency and resolution in the event of potential documented failures to or within a component. These are complemented by the Branch Trading Resilience HLD (DES/PER/HLD/0003) and in further detail within the Resilience and DR HLD (DES/PER/HLD/0001).

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
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 and the actual fault within the branch (see below)
Blade Frame PIM (Power Input Module)	Solution Design	No resilient service is dependent on the same PIM within a single Blade Frame Cabinet

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 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.

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).

POL00399072

3.2.2.4 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 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 below.

3.2.3 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.

POL00399072

	Maximum Target Recovery Time
Core 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 & MGRM Agents	2 minutes
Branch Access Layer	2 minutes
TES (application servers)	2 hours
NPS	2 minutes
APOP	2 hours
Main Host	2 hours
VPN and Radius Servers	2 minutes
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
Support Systems	
ACE Server	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
Accounting Radius Servers	2 hours
SAS Server	15 minutes
Signing Server	2 hours
SQL Server (ACDB, OCMS, Athene, MTAS)	2 hours
Branch History Database (SSC)	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

POL00399072

3.3 Disaster recovery

A DR requirement from (RM/CDE/003, v0.4) states that HNG-X will have one primary data centre and one disaster recovery (secondary) data centre. Requirement GEN1 in the document states that the DR (secondary) data centre will be used for testing. Section 5.1 states that the current recovery models (from the perspective of SLTs) are assumed.

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: • VPN & Radius • Core Switch & Data Centre Network • Branch Access Layer • Branch Database • Banking Agents • Debit Card Agents • NPS (Database)	2 hours
All remaining Services excluding POL-FS: • ETU, DVLA & MGRM online Servers • PAF, APOP & Service Hub Agent Servers • TES Application Servers • APOP (Database) • Main Host (Batch Database Server) (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)	5 hours
POLSAP	48 hours

Whilst Business Continuity testing might result in times lower than those defined in the table above for POLSAP and the Remaining Services group (as was sometimes the case on the Horizon solution), a single BC test (e.g. POLSAP) 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.

Possible triggers for DR to be recommended are:

- 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

POL00399072

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, flood).
- Security breach

It is envisaged that unless there are extreme circumstances (e.g. the complete destruction of a data centre) the decision to invoke DR will depend on the exact circumstances and will be made based on negotiations between Post Office Ltd and Fujitsu Services. For example it could be considered desirable to run with a lower capacity or reduced level of service outside a peak time, to give more time for an attempted repair and prevent an outage necessary for DR failover.

POL00399072

4 HNG-X Capacity Management Service

Requirement ARC-442 states Fujitsu Services will document potential performance and capacity bottlenecks in data centres as per the current capacity service via information, reports and models.

Requirement SVC-801 states Any changes to supported volumes will be handled by a change control via the Capacity Management service.

Requirements SVC-767, 768 and 769 state Within 6 months of contract confirmation it is necessary to confirm the methods of Service level measurement for SLTs and Performance Metrics. This will be performed in test and on a regular basis in live.

It is necessary for the Capacity Management Service to operate as described in the *HNG-X version of the Horizon Capacity Management and Business Volumes document* [PA/PER/033] and provide data on request internally and externally. This is both to understand what has happened in the past and also to predict what will happen in the future.

4.1 Business Volumes

Requirement SVC-812 states 3rd party transaction (response) times will be measured

Requirement SVC-813 states End to End transaction (response) times will be measured

As per the Horizon Solution, Business Volumes data will be available to show the service volumes for each service as shown in section 3.1.2.1 (i.e. EPOSS, Network Banking, DVLA etc) in the contracted periods of monthly, 2-weekly, weekly, 2 daily, daily, peak hour and peak 5 Minute average TPS. In addition the volumes for the new online services (Reporting, User management etc) will be available. For both the existing and new services, failures and response times will be provided (again as per the Horizon solution).

It is a Systems Qualities requirement that this data is provided to the Capacity Management Service through the data warehouse and other services.

4.1.1 Monthly Reporting

A monthly capacity report showing data including business volumes, failures and response times (as per the current Horizon solution) will be produced for internal and external distribution. This will include sections for new services introduced..

Additionally as per the current Capacity Management Service, other standard reports and data provided to POL on a regular basis (e.g. a monthly report showing details of the current Web Service based transactions (DVLA, PAF, APOP, MGRM, Telcoms & Kahala)) will also be produced.

4.1.2 Capacity Operational Review Forum

Requirement SER-2137 states Management reporting from Capacity Management Service will include performance reporting and service volumetric data (plus relative SLT adherence).

As per the Horizon Solution, the Capacity Operational Review Forum will take place every month, during which Fujitsu Services and POL Ltd can review the Business Volumes, response times and failures. Additionally the forum can examine adherence to contracted volumes and SLTs, predicting any potential future issues. Capacity/Performance issues from either Fujitsu Services or Post Office Ltd can be addressed within the forum and any further concerns or issues propagated appropriately.

POL00399072

The Capacity Operational Review Forum will provide the opportunity to adjust the flexible capacity for any changes required in service levels, and also raise any opportunities to POL for reduction in capacity for component or services.

4.1.3 HNG-X Capacity Management & Business Volumes Document

Requirement SVC-793 states the HNG-X version of the Horizon Capacity Management & Business Volumes document [PA/PER/033] will be maintained and kept up to date as described.

The capacity management service will keep the HNG-X Capacity Management & Business Volumes document [PA/PER/033] updated. This document holds all the contractual volumes and design limits (as described in section 3 of this document).

4.1.4 Capacity Management Service Database

The Capacity Management service database stores an assortment of live data samples for trend analysis and investigative purposes. Current contents include details of all transactions in the entire estate for annual sample weeks, samples of reports and user management operations. For HNG-X this database will continue to exist and remain a key part of the Capacity Management Service, with new processes created to populate HNG-X data for samples. The database will be used to verify usage expectations, patterns and correlate with other technical or business volume statistics. Additionally this database will store data from the following sources: Network, Xen/MSVS Virtualisation, Branch Access Layer, Oracle RAC and EMC disk statistics; the server will run the processes to import them. More details are provided in the Capacity and Performance HLD (DES/PER/HLD/0002).

4.2 Component Utilisation

Requirement SER-2180 states FS shall perform LIVE MONITORING of central systems and networks such that they are able to state at any time which system and networks are operating to specification and within defined thresholds – as comparable to current Horizon.

As per the Horizon solution System Component utilisation data will be collected on a regular basis, and mainly be stored in the HNG-X Performance Database, with additional statistics stored in the Capacity Management Service Database. It will be necessary for the Capacity Management Service to manage and maintain the Performance Database, as well as the processes both for its population and the supply of information from it.

Collected statistics will cover Network, Platforms, Storage, Processes, Databases and will include extensive metrics for CPU, Memory, Disk, Network, System Objects etc. More details are provided in the Capacity and Performance HLD (DES/PER/HLD/0002).

4.2.1 Pro-active & Re-active Analysis

Requirement ARC-435 states Capacity Management Service network utilisation statistics will enable any potential network cost savings or improvements that could be made.

It is envisaged that the Performance Database will be used both for pro-active and re-active analysis. A series of standard reports will be configured for a typical selection of basic statistics on key platforms; these will then be analysed regularly on an ongoing pro-active basis to confirm the system is working as expected and there are no capacity, performance or resource issues (i.e. memory leaks, higher than expected component utilisation, unexpected results from correlation analysis with business volumes etc). Additionally however the Performance Database is invaluable for reactive analysis as the wealth of

POL00399072

statistics it provides makes it possible to understand exactly what happens in detail within the system at every point monitored within the HNG-X system.

In addition it is intended that a new monthly report showing key utilisation statistics in the Network shall be produced as part of the HNG-X Capacity Management Service. This will allow analysis that will indicate opportunities for network cost savings or improvements.

4.2.2 Automatic Alerting

As per the Horizon solution, appropriate thresholds will be set for all key system components (e.g. CPU Utilisation exceeding a set level), such that should these thresholds be exceeded a mechanism will be provided to propagate these alerts through the System and Estate Management Monitoring team. These breaches can then be immediately escalated to the appropriate channels and investigated appropriately. See the System and Estate Management Monitoring Architecture for more details. However extra care must be taken with virtualised platforms where the guest utilisation might not be representative of the actual physical component utilisation.

4.3 Special Requests

The Capacity Management Service will also have to undertake various special requests (deemed as not being regular activities such as the Monthly Capacity Report).

4.3.1 Migration

Requirement MIG-2958 states It may be necessary to use Capacity Management Service data to enable authorisation by POL Ltd that the migration is progressing correctly.

Requirement MIG-3099 states Capacity management Service will monitor Capacity and Performance of agreed system components during migration

Requirement MIG-3118 states Capacity management Service may have to provide data to show that the pilot adheres to SLAs and non contractual performance SLTs in the HNG-X version of the Horizon Capacity Management & Business Volumes document [PA/PER/033].

During the initial HNG-X Migration phase; it will be necessary to carry out extensive analysis on all key areas of the system to confirm that the HNG-X system and its components are behaving as expected based on the business volumes. Such analysis can help discover system problems such as memory leaks or inefficient operations, and potentially prevent a major incident from occurring. For the initial period of HNG-X running this will be carried out more regularly and in more detail than the planned regular standard pro-active analysis process (described in section 4.2.1).

4.3.2 POL

Requirement DEV-358 states Currently the Capacity Management Service will often help in their discussions/negotiations with their clients

The Capacity Management Service is regularly asked by POL to provide extra reports or data to help with their analysis or investigations. For HNG-X the Capacity Management Service will continue to provide POL these when requested.

POL00399072

4.3.3 Incident Analysis

In the event of a major incident with the HNG-X system, it is likely that in order to investigate the causes (either during or after) the Capacity Management Service Performance Database will be able to provide crucial information to explain system behaviour and potentially the causes, at all levels of the solution (the Performance Database captures detailed process information that could be useful for understanding non capacity related issues).

4.3.4 Internal

Requirement DEV-348 states Will show functional size and complexity of current and future components. Component stats provided from Capacity Management service will help enable this.

In order to understand system behaviour for design/development/customer services teams, the Capacity Management Service may be asked to investigate relevant issues. It has often been possible to provide answers to important questions from the data collected for a previous analysis.

4.4 Modelling

Requirement SVC-794 states HNG-X Capacity and Performance models will be shared with POL Ltd as per existing Capacity Management Service.

Requirement SVC-770 states A combination of Benchmarking and modelling is deemed adequate for analysis purposes to reduce the costs of replicating a full live network environment.

Requirement SVC-797 states Volume testing and Modelling to show adherence for MHTR

Requirements SVC-798 & SVC-799 state Benchmark Period analysis to validate assumptions and modelling parameters for transaction volumes, administrative and back office functions.

The Capacity Management Service is responsible for the creation and maintenance of HNG-X Capacity Model as described in the *HNG-X version of the Horizon Capacity Management and Business Volumes* document [PA/PER/033]. This model will initially be used to confirm the provided capacity is adequate for the design limits, contracted volumes and expected workloads. Once the service has started, live data will be used to validate the accuracy of the model and fine tune it for accuracy. As the solution matures and develops, the model will be used to check potential scenarios based on expected business volumes and also to show potential capacity available for new services.

POL00399072

5 Solution overview

This section provides an overview of the methods used to achieve the targets set out in section 3 and where applicable indicates the architecture for any solutions provided by the Systems Qualities team. Sections 6 to 11 provide more details of how each of the other areas of the solution meet the targets.

5.1 Performance and capacity

HNG-X is sized to meet processing peaks. The sizing covers all the types of infrastructure: CPU, memory, I/O, network and disk.

The main components within HNG-X are scalable, and so can be easily sized to meet the required business volumes. Here are examples:

- Servers that can scale high. This means making the individual servers more powerful, by adding more processors, upgrading to more powerful processors, using multi core processors, or adding memory.
- Servers that can scale wide. This means adding more servers and sharing the load between them, typically using an intelligent network device.
- Different storage types. Where disk I/O is a bottleneck on secondary storage, faster primary storage can be used (however if already on primary storage any further improvements would be very complex and subject to availability and cost).
- Software/Application tuning. Where cost alternatives for hardware based solutions (such as those above) are prohibitive, application tuning may be viable (e.g. using a more efficient algorithm, generating less network traffic or optimising a Database or Operating System). It may also be possible to amend a software based solution to use the existing hardware more effectively.

Where Horizon services will be used (even if minor changes are being made – such as agent interfaces and software version upgrades), we can be confident that there is adequate performance and capacity based on known volumes and system component utilisation. Additionally newer higher specification hardware being used for HNG-X will actually provide further improved Performance.

Testing and modelling has been carried out to analyse each migrating platform; additionally simulation applications were ran on Horizon platforms (for comparative purposes) as well as the proposed hardware. These have given confidence in the overhead from using virtualisation technologies – vBlades and MSVS, and also shown the actual processor capacity improvements from the new hardware. In conjunction with performance statistics from Horizon these indicate there are no performance issues in any of the migrating or retiring platforms, Furthermore the EMC DMX-3 is significantly faster and more efficient than the Symmetrix used in Horizon, along with the reduced latency for SRDF (given the 2 HNG-X data centres are much closer than the Horizon ones); this will improve performance for migrating platforms with storage dependencies (e.g. NPS. POL-FS & DAT). Further details are covered within each section for that platform.

Where new services are being provided for HNG-X (e.g. Branch Access Layer and Branch Database), volume testing and prototyping, in conjunction with modelling has been carried out to confidently size these systems. Additionally these solutions provide excellent extensibility, so should it ever be necessary to provide additional capacity for the Branch Access Layer an additional server could seamlessly be added (scale wide). Similarly with the Blade Frame solution, it would be easy to use higher power blades for the Branch Database or other components (scale high). However this is always subject to availability

POL00399072

and cost (e.g. if already using the highest specification blade available it would not be possible to scale high – moving from single core to dual core blades may have licence cost implications).

5.1.1 Architecture for Performance and Capacity Solution

In order to provide the Performance and Capacity solution, it will be necessary for the Systems Qualities team to provide 3 components.

5.1.1.1 Performance Database

The Performance Database solution will be based on version 8 of the Metron Athene product (the Horizon solution used version 7 of Athene). Athene 8 uses a SQL Server database instead of the DBase based solution used in Athene 7. The architecture for the HNG-X Performance Database solution will be significantly different from the Horizon Performance Database solution which had both a long term and short term Performance Database.

The solution will be provisioned on a single platform running SQL server. The Performance Database will be made up of two databases, one of which will be the active database whilst the other will be used for backup purposes. No concept of a long term/short term database will therefore exist, with resilience now provided by replication of data to the opposite data centre (which additionally will be regularly backed up). Metron Athene acquires will installed on each new platform at the Data Centres and will be configured to collect all types of data (for example platforms which run a database will also collect database statistics). By default these will be set to a 5 minute sample time to ensure that adequate information is captured during the Horizon to HNG-X migration. The Metron Athene Control Centre on the SQL Server platform will collect the statistics using the standard Athene collection mechanism on a configurable basis (however if required it would be possible to collect and load any available data at any time).

The HNG-X Performance Database solution provides 5 significant benefits when compared to the Horizon Solution:

- Ease of configuration. Under Horizon changes to capture intervals had to be manually applied by raising an OCP, a costly process taking significant time t implement. With the HNG-X solution this is just a setting and gets applied instantaneously. Therefore should it be necessary to increase data capture granularity for a system of concern, this can be done in minutes.
- Quality of data. Athene 8 gives significant improvements in the quality of captured data. For example Network and Process data were previously unavailable for many systems.
- Availability of data. Under Horizon data was copied off overnight via a Maestro schedule, taking a significant time to load. At best data would be available 24 hours later. With the new solution, data is typically available for analysis within 30 minutes and data for key systems often available within 5 minutes.
- Default frequency of data capture. Under Horizon the default period of data capture was 20 minutes; the new solution has a default period of 5 minutes – this gives a far more accurate and reliable indication of what was happening on the systems; in particular this allows an accurate isolated break down of the various peak periods.
- Scope of data capture. Under Horizon not all systems had data capture enabled and therefore it was only possible to analyse those systems. With the HNG-X solution Athene is a part of the core platform build on all platforms and therefore we capture data for the vast majority of the systems.

POL00399072

Requirement SVC-811 states System Diagnostics from counter will be available when necessary, therefore Metron Athene Acquires will also be present on all the counters in the estate. However like the Horizon solution, this data will only be collected when required.

5.1.1.2 Alerting Thresholds

The alerting solution will be developed in conjunction with the Systems and Estate Management Monitoring Team. The Metron Athene product already installed for the Performance Database can be used in addition to Tivoli (as per the Horizon solution for Windows systems) to provide alerts on threshold breaches for key components (e.g. CPU utilisation, available memory), however will not be configured to do so unless deemed necessary by System Qualities. These alerts can be propagated by a process designed and implemented in conjunction with the Systems and Estate Management Monitoring Team if required.

5.1.1.3 Capacity Management Service Database

As per the Horizon solution a database will store network statistics provided from the HP-Open view solution, as well as various statistics from the Branch Access Layer Servers (exposed through Java), Xen Virtualisation statistics, EMC disk statistics and other data provided for the Capacity Management Service (as described in section 4.1.3). Additionally storage or database data captured for the capacity management service will be stored on this system.

As this database can be very large, it is necessary to have a system with a very large storage capacity. The system will run on the SQL Server database on the same platform as the Performance Database, with the DR system populated using EMC disk replication. Existing databases and data from the Horizon system will be transferred and ported to a local Capacity Management Service system, to provide the ability to compare previous volumes and patterns.

The HNG-X database will hold all non Athene generated performance data and also the applications necessary to load it (e.g. automatic loading of network statistics from HP-Openview and reports to provide access of data to the Networks team).

5.2 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).

POL00399072

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, MGRM & Service Hub 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	Agents are connected to all Servers; on failure detection an alternative connection is used
APOP database server	Bladeframe	Once blade failed over, service reintroduced.
Main Host database server	Platform (Active/Standby Configuration)	For current active standby configuration, use identity of Active by VIP
VPN servers	Platform (N+1 Configuration)	Counter is allocated a VPN cluster which has 4 different VPN Servers.
Radius servers	Platform (N+1 Configuration)	Network device checks service and no longer uses a failed system

5.2.1 Automatic fail-over

To achieve contracted service levels, some systems must typically be restored by automatic fail-over within two minutes of a failure. The main components that provide this are listed below:

- The Oracle Real Application Cluster (RAC) is used for the Branch and NPS databases. Oracle RAC is an established widely used system that provides a high level of resilience. Connecting applications use a software heartbeat mechanism to determine availability of the clustered Servers and have connections to each server. If the heartbeat mechanism determines that a server has definitely failed, then the connecting agent switches to use an alternative connection to another server. This implementation was used successfully in the Horizon NPS solution.
- The network devices used to handle online requests in to and out of the Branch Access Layer, DVLA, PAF, APOP, MGRM and Service Hub web servers, all of which are provided in an N+1 configuration (so the entire workload can be handled by N servers). The network devices are intelligent switches (e.g. Cisco ACE) which provide many services, including load balancing to available servers and resilience; the switches are configured to send heartbeats that determine the availability of their respective services, and should they determine a failed system, all requests are sent to the remaining available server(s). This implementation was used

POL00399072

successfully for the Horizon Web Services (DVLA, PAF. APOP. MGRM, TLCM and KAHALA) solution.

- Application software based, for example within Network Banking, Debit card and ETU agents. Servers for these services are provided in an active standby pair configuration, where heartbeat monitors between each service pair written to the NPS, determine any failure within the current active server. On the standby determining the current active server has failed (or the current active server informing the standby of its failure), the standby server becomes active and informs any connecting systems that it is now the active system. Connecting systems will subsequently connect to this new active server instead of persisting with the failed previous active system. This implementation was used successfully in the Horizon Network Banking, Debit Card and ETU solutions.

For all cases HNG-X is sized so that if automatic fail-over takes place, the remaining servers and components always have sufficient capacity to deal with the additional load.

After automatic fail-over has taken place, the failed server or component will be recovered and then reintroduced to the live service (at the next opportunity) to bring the system up to its full working complement. However until it is possible to reintroduce the server or component, the system will be running with reduced resilience and a SPOF will potentially exist in the solution.

5.2.2 Blade Frame

Most HNG-X systems run on Blade Frame servers.

The Blade Frames and their server configurations are highly resilient. The Blade Frame servers additionally use resilient external storage devices (servers outside the blade will use RAID 1 mirroring for local disks).

As well as this inherent resilience, the Blade Frame also allows new servers to be configured and booted very rapidly, to take over from failed servers as each blade server only gains its identity once booted.

Where services do not need to be recovered within two minutes, resilience will be achieved by configuring and booting replacement servers very rapidly (or return to the full working complement for those that do). This means that every service can in effect have a standby server, but that the total number of standby servers required is minimised. However also, where services do need to be recovered within 2 minutes, Blade Frame offers a further layer of resilience when compared to the Horizon solution as a failover blade can be available within minutes to bring the solution back to its full working level (in terms of physical platforms) and doesn't require a lengthy repair or reprovisioning cycle like per Horizon. Furthermore there is no dependency of physical components as a layer of abstraction is introduced (e.g. where exact physical components are required as per Horizon systems, as these get superseded and become old – procurement can become increasingly complex).

However blade frame technology runs the same platform image on a standby blade in the production data centre as well as at the DR centre. There is therefore an inherent risk that a corruption of the platform image will render the server unusable until support downgrade it to an earlier backup and a procedure to minimise this scenario is necessary.

5.2.3 Data centre

Adequate provisions are made at both data centres, therefore allowing for a full service in the event of DR. Full facilities are provided at each data centre, with adequate space, suitable flooring and reliable air conditioning.

The power supply at the primary data centre will have two separate power providers, each with an uninterruptible power supply (UPS) device.

POL00399072

While there is only a single power provider at the secondary data centre, there will be duplicate UPS and power distribution units (PDU) to provide extra resilience (this is deemed adequate).

Many systems (including the Blade Frame, Primary Storage, Batch Host) will have redundant power supplies, with a main and backup connected to different PDUs.

This will ensure that failure of a power provider (at the primary data centre), power cable or power socket will not interrupt the service to these systems.

5.2.4 Network

The main networks are entirely resilient, with no single points of failure across switches, routers, cables or sockets (except those documented 3.2.2.3).

Most branches have an alternative network connection to use if their primary network connection fails. In most cases, the primary ADSL connection has a back up network in the event of an ADSL failure.

Within larger branches, hubs and routers are used so that failure of one counter PC does not impact the others.

5.2.5 Monitoring

System components such as processors, disks, fans and switches are monitored to detect actual failure and intermittent or potential future failure of the component (or the entire system) (see Monitoring sub Architecture from the System and Estate Management Architectures). This allows failing components to be replaced before any serious impact may occur.

When a component has failed, a remaining functioning component may become a single point of failure. This proactive monitoring minimises the time for which the solution is running with reduced resilience or a single point of failure.

As far as possible, standard components will be used so that it is easy to hold components in stock and replace failed components quickly (until 2015).

However while the solution provided by the System and Estate Management Monitoring team will determine a failed system, which can then be replaced, it will also be necessary for the System and Estate Management team to provide a software based solution to determine the availability of software based services for mechanisms that do not provide this functionality (network based resilience solutions do). The solution will be similar to the Horizon EACRR (Enhanced Agent and Correspondence Server Recovery and Resilience) solution; it will regularly check service availability, and if failed it will attempt to automatically restore the service. If a failure occurs, the details will be propagated through a mechanism provided by the System and Estate Management Monitoring team. This will allow for appropriate action to be taken in event of a failure where it is not possible to restart the service, but the system remains available. Even if it is possible to restart the service, it will help indicate any potential problems occurring in the system that will need further investigation.

Other new or monitoring services in the Horizon System such as the Outlet Monitoring Agent, will be updated and designed in conjunction with the System and Estate Management Monitoring team.

5.2.6 Data recovery

Although the HNG-X infrastructure is highly resilient, data may be corrupted due to hardware or software error, or become invalid from an intentional or unintentional action (a file might be accidentally or maliciously altered/deleted).

POL00399072

Data corruption will typically be written to locally mirrored disks (including in the storage systems) and replicated to the secondary data centre.

In the event of system tools being unable to repair a corruption without data loss, to allow recovery from this situation, backups are taken at both the primary and secondary data centre. This allows data to be restored back to a known point. In many cases, database logs are also taken which can be replayed to correct the corruption without any data loss.

Hot backups are taken for databases that support an on-line service, such as NPS and the Branch Database. Hot backups allow the database to remain active while the backup it taken.

For backup solutions crucial to branch trading, RPO (Recovery Point Objective) is 0, such that a requirement of no data loss is required.

RTO (Recover Time Objective) is to adhere to the availability SLTs described in section 3.2.1.

In order to achieve high availability for the Branch Database, which together with the Branch Access Layer is a crucial component in allowing branches to trade, it has been decided to provide an extra standby branch database through Oracle Data Guard replication, that will provide potential resumption of service from a corruption on the main branch database without having to revert to restoring backups which can be a lengthy and complex process.

5.2.6.1 HNG-X Backup Strategy and Architecture

Both persistent data and certain system images (e.g. SAP) need to be backed up. Persistent data can be system registries, DNS data or databases. Oracle databases that need to be backed up include the Branch, NPS, APOP and other batch databases (TPS, TES etc) on the main host (DAT).

The Branch Database, NPS and APOP databases will use Oracle RMAN hot backups, which will enable service to continue during the backup process, but the batch databases will continue to shut down and use cold back ups – their batch nature and availability requirements do not require a hot backup implementation.

Other smaller systems requiring backup will be backed up across the network to their media servers.

For the non RMAN based backups Symantec (previously Veritas) Netbackup is the chosen backup software package being used. This product is widely used in successful implementations and has an excellent reputation; Netbackup was also successfully used for the Horizon SMDB backup solution and as such there is experience of its use. The media servers will be connected directly to the media servers and backup storage through fibre channel. There are multiple media servers for each OS. There are 2 Solaris media servers and 1 each for Windows and Linux, as well as a backup server for cataloguing.

All backup is disk based, with the RMAN backups writing directly to EMC Clarion Disk Backup, whilst EMC CDL (Virtual Tape Library) is the choice of backup media for the Netbackup backups; disk based backup was chosen due to cost, simplicity, reliability (physical tapes and tape drives proved unreliable in the existing Horizon solution). Additionally recovery times are faster for disk based backup than tape. However some tape based back up facilities have been maintained in addition to the disk to provide support in situations where tape facilities are required (e.g. might need the capability to send a tape off site for support or development).

For business critical data it will be necessary to take back ups on a regular basis (along with archive redo logs) to allow for fast restore times. Furthermore consistency checks will be made to help verify the backups as being correct (and check for any corruptions) and assist in setting the retention times.

POL00399072

As well as application data, backups of system images are taken. This makes it easy to repair problems with the installation, such as restoring standard files.

Due to the critical nature of the Branch Database for trading, a further standby database will exist using Oracle Data Guard. This will provide an extra level of resilience and potentially allow for service in the event a corruption or other issue that causes the primary Branch Database to fail. This solution is deemed necessary due to the availability targets of the core system (which depends on the Branch Database) and high recovery times for backups.

5.2.7 Retiring Horizon systems

Horizon systems that remain in Bootle/Wigan will only be required for the period when Horizon and HNG-X coexist. This includes the Riposte messaging system that runs on the Correspondence Servers and the Generic Agent servers.

In Horizon, these achieved resilience by distributing workload across servers at both data centres. If a server failed at either data centre, the other servers would take over the work.

For the intermediate solution when both Horizon and HNG-X coexist, these systems will continue to use the Horizon resilience methods.

5.3 Disaster recovery

A secondary data centre is available for DR in case of catastrophic loss of crucial components at the primary HNG-X data centre (or the entire primary HNG-X data centre itself).

The secondary data centre solution provides the same functionality, configuration, capacity and resilience as the primary data centre. (see data centre facilities for power providers)

In normal operations, the HNG-X secondary data centre is not used as part of the live configuration (network access is triangulated and as such certain components run active/active). This is known as an active-passive configuration. This is different from Horizon for which the load was shared between the primary and the secondary data centre, which is known as an active-active configuration.

To make sure that no data is lost if DR is invoked, all business and system critical data is replicated to the secondary data centre. Synchronous replication is used, which commits changes to storage at both data centres for any transaction – therefore guaranteeing storage at the secondary data centre contains an exact duplicate of the storage at the primary data centre at any time. (The alternative, asynchronous replication, is more efficient but does not guarantee that data is never lost. However this could be used for other HNG-X storage where some data loss would be acceptable).

Some Horizon systems are required for the migration period when Horizon and HNG-X coexist. This includes the Riposte messaging system that runs on the Correspondence Servers, and the Generic Agent servers. These will use the active-active approach used for Horizon.

The secondary data centre will be used for testing. On the Vol rig, this will allow HNG-X to be tested on the exact same configuration as the live system, providing much more accurate test facilities than have previously been available.

If DR is invoked (official notification has been given by Post Office Ltd):

- All testing ceases. (testers might be forced off and the next step proceed immediately due to time constraints – testers should not expect a clean shut down!).
- Network is switched from Test to Live configuration (secure from test etc)

POL00399072

- DR systems are booted from centralised storage, and assume the identities of the systems that they replace.
- If this is any quicker than the permitted time (SLTs as defined in section 3.3), the remaining time can be used to investigate any problems due to the DR switch (if necessary).

5.4 Testing

HNG-X is tested to make sure it meets capacity, performance, availability, DR and recovery requirements.

- Volume testing is used validate that each HNG-X service can handle the design limits, with no significant reduction in response times.
- Volume tests are used to confirm the assumptions in the performance models that have been used to predict capacity requirements. .
- Performance testing for online transactions is carried out by running a process that injects transactions into the system. This is used to assess different services' ability to cope with different volumes of transactions, and to understand maximum achievable throughput.
- Resilience testing is carried out by deliberately failing components to simulate a failure and force recovery.
- DR testing is carried out by a regular DR test exercise and process walkthroughs.
- Backup recovery tests.

5.5 Measurement

Measurement is used to demonstrate that Service Levels are met. Measurement is also required as an input to the Capacity Management Services (see section 4.1).

The HNG-X measurement regime is very similar to that in Horizon.

The following are measured:

- Business volumes for each service for each contractual period. Peak hourly values and peak 5 minute TPS are measured.
- End to End Response times for each service.
- Response times for external services (and therefore also HNG-X times)
- Failures
- System availability and downtime for the entire solution and for individual components.
- Network utilisation and operational statistics.
- System metrics, such as CPU, Memory, I/O, Process Information, System Objects and Database statistics. These are typically regularly collected from Athene into the performance database. Further statistics will also be collected, including Network statistics, Xen/MSVS Virtualisation Statistics, EMC disk statistics, BAL operation statistics and Oracle database statistics provided by utilities such as Oracle AWR.
- Storage volumes and usage (including archiving, databases, POL-FS as well as standard storage).

POL00399072

5.5.1 Extra Measurement required for Virtualisation

VPN Servers will be running on virtualised Windows servers under Microsoft Virtual Server. Extensive use will be made of vBlades running virtualised Windows and Linux Servers. Once virtualised these systems will only provide a view of system utilisation with respect to what it has been allocated by the Host/Hypervisor. This is therefore not an accurate representation, nor does it provide a definitive indication of what the physical server is doing. Therefore for MSVS platforms additional performance statistics will be collected from the Host operating system. For VPN servers which will continue to operate in HNG-X these will be collected by the HNG-X PDB. For the vBlades, a process will run against the Hypervisor collecting various statistics which will then be imported into the CMS-DB.

5.5.2 TWS requirements for Hydra

Once the Main Host (DAT) has operationally relocated to the HNG-X data centres at Migration Weekend B, it will be necessary to maintain the collection of the Athene data on the remaining servers at the Horizon data centres at Bootle and Wigan. The process to populate this data into the Horizon STPDB and the LTPDB will continue.

POL00399072

6 Counter

6.1 Overview

Each counter within the Post Office branch network runs the HNG-X Counter system.

The HNG-X Counter is a Windows desktop PC running a Java application. It uses a number of specialised peripherals, such as a touch screen, printers and scales.

There is no separate back office system within the branches (except for the back office printer, which is LAN connected for offices with more than 9 counters). Every counter contains all the required software to carry out administrative functions.

The counters connect to a wide area network (WAN) using a router. The WAN connects the branch with the data centre.

The HNG-X Counter reuses hardware previously used by the Horizon Counter.

For more information on the counter, please refer to the HNG-X Counter Architecture and the HNG-X Counter Business Application Architecture documents.

6.2 Performance and capacity

The counter hardware has sufficient capacity to run a Java virtual machine (JVM) and the HNG-X Counter application.

There is enough network bandwidth to support simultaneous trading from every counter.

6.3 Availability

With the branch router, for a branch with more than one counter if one counter fails the others can continue trading.

Most branches connect to the WAN using an ADSL link. If this fails, the branch router provides alternative connections so the branch can continue to trade (although this would not be transparent). Network connections do vary from branch to branch, and some branches have different connections.

The router is a single point of failure within a branch. While the branch router becomes a single point of failure within the branch, the branch router is a much simpler device than a counter acting as Gateway and made up of fewer components to go wrong. Additionally replacement is much faster – both minimising the MTBF (Mean Time Between Failure) and the MTTR (Mean Time To Repair) for the branch.

Although it is not possible to provide complete resilience, any failure at a counter or branch is isolated to that counter or branch.

6.4 Disaster recovery

If the secondary data centre is used, the WAN will connect the branches to the secondary data centre. This would be transparent to the branches.

There are no DR plans for the branches within the scope of this document (POL have plans to enable a quick roll out of an entire new outlet in case of disaster).

POL00399072

6.5 Testing

Testing will take place on the counters for all aspects of System Qualities as described in the HNG-X Counter Architecture document. This will include counter performance testing, testing of basket mixes as specified and outlined within [PA/PER/033], video benchmarks and their target replacement using built in performance metering within the counter software.

6.6 Conclusions

The HNG-X Counter is simpler than the Horizon Counter, but uses the same hardware. Use of the counter is limited by the speed at which customers can be physically served. There is therefore a high degree of confidence that the HNG-X Counter system will perform adequately.

Although there are single points of failure within individual branches, any failure is isolated to the counter or branch. Post Office as a whole can continue to trade.

POL00399072

7 Network

The *HNG-X Technical Network Architecture* document (ARC/NET/ARC/0001) describes the network used for HNG-X. It includes detailed sections on resilience and performance, including targets and risks. This section provides a summary of the information.

7.1 Overview

The following diagram illustrates the HNG-X network.

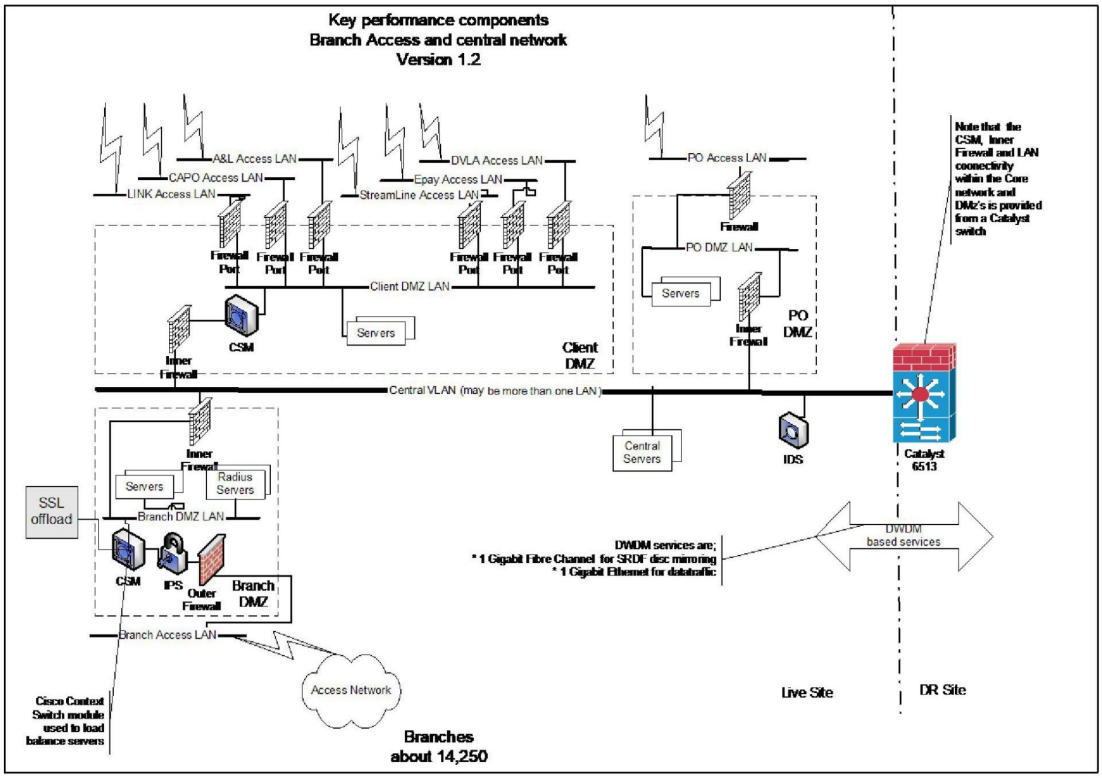

Figure [SEQ Figure * ARABIC] – HNG-X Network

7.2 Performance and capacity

7.2.1 Targets

Requirement ARC-438 states that the network capacity should meet the requirements of the *HNG-X Capacity Management and Business Volumes* document. Requirement SCD-139 states that the network to the branches and counters should have sufficient capacity so that delivery of data to one branch or counter does not impact others. (In Horizon there were problems with this, such as limits to the number of branches simultaneously connecting to the V-SAT network, and contention problems between Network Banking and Reference Data).

POL00399072

The *HNG-X Technical Network Architecture* document breaks these requirements down in detail. Some of the main design limits are summarised below.

Measure	Design limit
Branch Access Network, TCP connections per second	828
Branch Access Network, concurrent connections.	34,954
Branch Access Network, total throughput	35 Mbits/second (4.3 MBytes/second)

7.2.2 Approach

The Branch Access Network will use the same approach and component types as were in place for Horizon. This currently supports throughputs of more than 70 Mbits/second, and can achieve 90 Mbits/second. It will therefore meet the design target for HNG-X.

The LAN, Storage Area Network (SAN) and inter-campus links will use the same approach as Horizon. There are only a few new components, such as those that perform secure socket layer (SSL) encryption, and these all significantly exceed design targets.

Additionally where necessary QOS control will be used for traffic shaping. This will be implemented to prevent a single service (e.g. reference data) from monopolising the network to the detriment of other services (details are provided within the Network Architecture document)

Performance and throughput can be increased by adding alternative or further equipment (where available and cost effective).

7.2.3 Risks

Most of the network approach and equipment types were successfully used on Horizon, and so carries very few risks. Horizon suffered complication and risk of network congestion because of the proprietary UDP-based transport which is used by the Riposte messaging product. This is being removed for HNG-X, which reduces risk further.

There are some new parts of the network. These have been investigated, but will be tested further.

- The branches will connect with ADSL and with other network types such as GPRS. Some further testing of the Access Routers to which the branches connect will determine how many branches each can support, given the mix of network types.
- The components that support SSL encryption are new for HNG-X, and some further testing will be carried out to gain experience of these.

7.3 Availability

7.3.1 Targets

The network is critical to the solution. If any component fails, service must be restored to alternative components very quickly.

Detailed targets for each component are given in the *HNG-X Technical Network Architecture* document. In summary:

- For central network equipment, target recovery times are between 10 seconds and 60 seconds.

POL00399072

- Within a branch, maximum target local network recovery time is 210 seconds (this is the worst case scenario and will typically be much lower).
- If a branch ADSL connection fails, the switch to a usable backup service takes no more than 6 minutes (again this is a worst case scenario value and will typically be much lower).

7.3.2 Approach

High availability will be achieved using the same approaches as Horizon. In every case this involves automatic fail-over to standby components (e.g. Firewalls, Switches).

7.3.3 Risks

As with performance and capacity, availability risks are minimised by using the same approach and equipment types successfully used on Horizon.

Further testing will be carried out on new equipment types to verify that it can be recovered automatically within the target timescales.

7.4 Disaster recovery

7.4.1 Targets

Requirement ARC-443 states that the Data Centre and Network resilience capabilities will be as specified in Data Centre Operations Service Description.

Requirement SEC-3174 states that test systems shall only share logical network connections with live systems under carefully controlled conditions (they do share the same physical network).

It is therefore required that the network at the secondary Data Centre will provide the same capacity and resilience as the network at the primary Data Centre, but will be secure from any testing (or other) activity at the secondary data centre and be made available within the DR Targets as described in Section 3.3 once testing activity has ceased for data centre failover.

7.4.2 Approach

There is a single active network that covers both the primary and the secondary data centre.

The secondary data centre is usually used for testing. The live and test network traffic are separated using a variety of mechanisms such as virtual local area network (VLAN), virtual interfaces, and virtual route forwarding (VRF). Where components are shared for both live and test traffic, the resources available for test traffic are limited to avoid any impact on the live service.

Branches access the data centre using a virtual IP (VIP) address. If DR is invoked, these addresses will simply map to the secondary data centre.

7.4.3 Risks

If DR is invoked, there is very little change to the network. The risk of the network not providing service is very low. However for any subsequent releases that require network changes, it is essential that these integrate into the base solution in a way that minimises risk to DR capabilities. Additionally once

POL00399072

operating in DR the resilience provided by triangulation is no longer provided (as documented in the SPOFs table in section 3.2.2.3).

7.5 Testing

As described in the sections above, some further testing is required for new network equipment types, to validate that they can meet the required performance and availability targets. This will include testing that the Access Routers to which the branches connect can handle the potential mix of network types.

7.6 VPN Layer

The HNG-X VPN layer provides the equivalent functionality of the Horizon VPN layer, taking encrypted traffic from the counters and propagating it (to the BAL) unencrypted (and the reverse in the other direction). As per Horizon there are 12 VPN Servers in each data centre running in an active/active configuration, however these are hosted on 4 physical RX300 systems running MSVS, each hosting 3 virtual VPN Servers (additionally a 4th platform such as a domain controller can be hosted on the system, but all of which have minimum overhead). Besides that the layer is pretty much unchanged from Horizon.

7.6.1 Performance and Capacity

Testing and simulation on guests running under MSVS has found its overhead to be minimal. Each physical platform has 4 cores available and therefore each VPN server will have access to an available core (a limitation of MSVS is that it only allocates a single core for each guest system, however in this case it is probably advantageous). Each virtual VPN platform will therefore have significantly higher CPU/Memory specification than the Horizon VPN servers. Furthermore given the branch to data centre network has a limit of 100Mbps and each virtual VPN server has exclusive use of 2 Gigabit Ethernet ports there are no bandwidth limitations.

7.6.2 Availability

VPN availability is as per Horizon, with 12 VPN servers in each data centre. Each counter is allocated to a VPN cluster, which contains 4 VPN Servers, 2 in each data centre all on different physical servers (therefore avoiding any dependence on a particular physical server). In the event of a VPN server failing or being unavailable, an alternative VPN server within the cluster is used.

7.6.3 DR

If DR is invoked, there is very little change to the VPN layer as it runs in an active/active mode and the same resilience is used as if a local failure occurred. The risk of the network not providing service is very low. However for any subsequent releases that require network changes, it is essential that these integrate into the base solution in a way that minimises risk to DR capabilities. Additionally once operating in DR the resilience provided by triangulation is no longer provided (as documented in the SPOFs table in section 3.2.2.3).

7.7 Radius Layer

The HNG-X Radius layer provides the equivalent authentication and accounting functionality of the Horizon Radius servers. Each server provides separate Radius services giving authentication and accounting for counter traffic of each service type (ADSL, ISDN, VSAT and Wireless Wan), as well as

POL00399072

TACACS which gives authentication, authorisation and accounting for admin/support access. The HNG-X Radius layer is located on BladeFrame running as Linux vBlades with 2 Radius Servers in each data centre running in an active/active configuration.

7.7.1 Performance and Capacity

Each virtual Radius Server has significantly higher CPU/Memory specification than the Horizon equivalent and is considered to be more than adequate for the potential workload.

7.7.2 Availability

Radius server availability is provided by an N+1 configuration with the workload load balanced by a Network Device (Cisco ACE). Each service is allocated a VIP (Virtual IP address) for which the service is routed via the Network Device to an available server. Whilst there are 2 servers in each data centre, in the event of a server failing, the service will be introduced on a replacement blade with full level of service being restored seamlessly upon introduction.

7.7.3 DR

If DR is invoked, there is very little change to the Radius layer as it runs in an active/active mode and the same resilience is used as if a local failure occurred, with the Network Device routing services to the 2 Radius Servers in the Secondary Data Centre. The risk of the network not providing service is very low. However for any subsequent releases that require network changes, it is essential that these integrate into the base solution in a way that minimises risk to DR capabilities. Additionally once operating in DR the resilience provided by triangulation is no longer provided (as documented in the SPOFs table in section 3.2.2.3)

7.8 Conclusions

The network (including the VPN and Radius layers) has adequate capacity to handle peak volumes, and will meet availability targets.

Some controls will be required to make sure that Reference Data downloads and software updates do not occur during peak transaction or reporting periods, as this could increase the capacity requirement beyond the expected volumes.

POL00399072

8 Branch Access Layer

8.1 Overview

The Branch Access Layer (BAL) provides data persistence and service routing services to the HNG-X Counter system.

The HNG-X Counter sends requests to the Branch Access Layer. The BAL usually responds by accessing the Branch Database or in some cases passing service requests on to various online services.

Each of the 10 BALs runs 2 OSRs (Online Service Routers) in a Java-based server environment on 32 bit Linux virtual servers (although these are currently allocated to a single physical server – testing has shown that using a virtual server has minimal overhead, whilst giving flexibility for test by not binding a BAL server to a physical server were not necessary)

8.2 Performance and capacity

8.2.1 Targets

The overall Branch access layer should be able to handle an overall 1128 TPS over a 5 minute period (allowing for an estimated actual peak of 1466 TPS) for all transactions. This is further broken down into the volumes required for each individual service in section 3.1.2.4.

8.2.2 Approach

The BAL is run over multiple high-specification servers. Network devices (CISCO ACE switch) distribute the load evenly between the servers. Capacity can be scaled up and down either by adding/reducing the number of virtual servers over which it runs or by using more/less powerful virtual servers. The servers are sized to cope with peak rates.

The database does not run on the same servers as the BAL. All database access is passed to the Branch Access Database, which runs on separate, specialised servers, which are separated from the BAL by a firewall.

The BAL only carries out relatively simple database access. More complicated processing is passed on to other online services.

The BAL handles routing to other online services efficiently. The BAL manages multiple requests to other online services on a single set of threads. The BAL does not use more threads if other online services are running slowly (each OSR has different queues for services, therefore for example reports cannot effect basket settlements and an external online service could not effect basket settlements or reports).

The BAL software runs under Java 1.6 allowing for the maximum Java performance (Horizon web servers ran under Java 1.4 which is significantly less efficient) and is written to be efficient such that the software is focused on carrying out the required tasks, with no additional generic components that would use valuable resources.

POL00399072

8.2.3 Risks

Some online Horizon services (DVLA, PAF, APOP. MGRM, TLCM & Kahala) were java based web servers switched (load balanced) via a CISCO CSM and demonstrated that the general approach is sound.

An initial prototype to understand feasibility of the approach was created using a simple initial BAL of two simple services using Interstage version 8 on a 32bit version of Windows. Performance test/simulation showed that a single server with 4 dual-core processors could run at a sustained rate of 944TPS (669 TPS settlement plus 275 TPS reporting). The actual system will be more efficient, running on 32bit Linux and replacing Interstage and related components with efficient Java specifically designed to concentrate on the necessary workload. While this initial simulation was very limited the results also gave an initial indication of the number of BAL application servers that would be required.

If business volumes increase, or the processing load is higher or more complex than anticipated, the BAL can be scaled by either adding more virtual servers or using more powerful virtual servers. Configurations of hypothetically up to 250 servers could be used with this model, which is deemed more than sufficient to satisfy the required volumes and any future requirements (subject to availability and cost – e.g. the BDB would also need to be scaled according).

Further testing has shown that any risk in combining the several services that make up the overall BAL onto a single platform is negligible. This was originally perceived to be a risk as no Horizon systems used a similar configuration and the original feasibility simulation carried out only had 2 simulated services.

A further specific risk was found during early testing whereby should 1 or more external online service (over which there is no control) fail or run slowly, the number of outstanding connections would rapidly reach the maximum permitted number of connections, therefore causing all further transactions to fail. This was however addressed by the bespoke non blocking I/O development for online service routing and event queuing development for online service routing.

8.3 Availability

8.3.1 Targets

If a BAL server fails, normal service should be resumed within 2 minutes.

8.3.2 Approach

The BAL applications servers are provided in an N+1 configuration, such that if a single application server fails there is sufficient capacity to handle the contracted volumes and design limits on the remaining servers. If an application server or one or more of its services fails (OSRs), the network device (CISCO ACE) will automatically detect that through its monitoring mechanism. Once it has established this failure, it will route requests to the remaining active servers (or if at OSR level to the alternative OSRs). It will then be possible to introduce another blade as an application server which will bring the configuration back to its full N+1 level. As transactions to the BAL are stateless, for the period while the switch has not detected and acted on a failed server, when retying a failed transaction it only has a 1/N+1 probability of going to the same failed server (or 1/((N+1)*2) if an OSR failure), therefore further reducing the impact of a failure on an individual counter.

POL00399072

8.3.3 Risks

The online Horizon web services successfully used the same fail-over approach using an older, less powerful network device (Cisco CSM) than that proposed for this solution (Cisco ACE). The approach is therefore low risk.

8.4 Disaster Recovery

8.4.1 Targets

If DR is invoked, the BAL service should be restored within 2 hours.

8.4.2 Approach

The general approach to DR described in section 5.3 applies to the Branch Access Layer.

8.4.3 Risks

No specific DR risks have been identified.

8.5 Testing

Performance modelling and testing of the BAL Layer systems has been carried out and indicated adequate capability to satisfy Capacity, Performance and Resilience requirements (as discussed in 8.2.3).

8.6 Conclusions

Tests and simulations show that the proposed solution will support the required transaction volumes. The underlying hardware and software are highly scalable. The Horizon systems show that the load balancing and resilience approach is sound.

A decision has been made to implement bespoke User Management system within the BAL and the Branch Database for Branch users. Part of the justification for this is the fact that these are the only component within the system with the necessary performance and resilience characteristics (additionally it avoids introducing another dependency into the Counter-BAL-BranchDB).

In addition to the Performance Testing and modelling undertaken during the HNG-X development phase, the load on the BAL will rise gradually in line with branch migration. Hence pro-active trend analysis during the roll out phase will be able to spot any performance or capacity issues that may arise prior to the full workload hitting the BAL. (it may even be possible to limit the number of BAL application servers to stress the servers to a higher level that will closer represent their final workload by switch configuration).

POL00399072

9 Online Services

9.1 Overview

HNG-X will include all the online systems used within Horizon.

A few changes will be made to these systems:

- Online access using Riposte messaging will be replaced with HTTP. Online access will be controlled by the Branch Access Layer.
- Solaris based systems (NPS and APOP) will be ported to Linux.
- All Oracle systems will be upgraded to version 10gR2.
- All Web Servers except for PAF (DVLA, APOP, MGRM, Service Hub) will migrate to Linux.
- All Web Servers will migrate from Interstage 6 to Interstage 8.
- All Windows servers will be upgraded to Windows 2003.
- The underlying hardware will be upgraded. Windows and Linux-based systems will use Blade Frame servers (all servers except for NPS & APOP databases will use virtual vBlades).
- Full resilience will be provided at each data centre, with active/standby configurations for agents, N+1 for Web Servers and Oracle RAC for NPS database (APOP database will rely on blade failover) as per Horizon solution.
- HNG-X systems will use an active-passive DR approach, not an active-active approach.
- Banking and DCS changes will be made for PCI compliance.

To generalise, these changes incrementally improve performance and capacity as new higher specification hardware will be used and no inter data centre communications are necessary (apart from synchronously replicated storage). Because HNG-X uses an active-passive DR approach, there are some changes to resilience and DR, and a higher level of resilience and DR is provided as both data centres provide full resilience for all services.

9.2 Performance and capacity

The component targets for the migrated Horizon systems are the same as they were under Horizon. The systems are largely unchanged. The specification of new servers is based on well understood usage. The updated systems will be fully volume tested.

Performance simulation testing has shown the overhead of vBlades to be minimal, and these tests were also run on Horizon hardware; these, in conjunction with Horizon system utilisation statistics obtained from the Long Term Performance Database, indicated that the platforms all have adequate capacity and should provide a significant performance improvement. Furthermore there should be additional improvements (not quantified at this stage) from using Windows 2003, Linux and Interstage 8 (where appropriate). The NPS and APOP databases should also see an improvement in their I/O as the EMC DMX3 is significantly more efficient and faster than the Horizon EMC Symetrix, as well as much faster SRDF latency due to the closer proximity of the HNG-X data centres compared to the Horizon data centres at Bootle and Wigan.

POL00399072

9.2.1 Targets

The individual services should be able to cope with their design limits and contracted volumes as described in section 3.1.2.2.

9.2.2 Approach

The servers are sized to handle contracted volumes and design limits (actual levels of use are also well known).

9.2.3 Risks

The Atalla devices for PCI are network attached and will require 3 crypto operations per transaction. However initial analysis indicates this not to be a problem and improved hardware should mitigate this.

9.3 Availability

9.3.1 Targets

Availability targets for each service remain the same as for the Horizon solution. These are shown in sections 3.2.1 and 3.2.3.

9.3.2 Approach

The Online services will switch over automatically within the 30 – 60 seconds fail-over window. The resilience models for online services are the same as for the Horizon system:

- NPS uses Oracle Real Application Cluster technology. The workload is split across two nodes (servers). If one node fails, the other is capable of handling the whole workload by itself.
- In a similar way, some online services such as DVLA, PAF, APOP, MGRM, TLCM and Kahala ensure high availability by using duplicate servers (N+1) which share the workload. If one server fails, the remaining servers can handle the whole workload. Load balancing and fail-over is provided by the network layer, which uses a heartbeat mechanism to detect available servers/services.
- Other online services such as Banking, Debit Card and Electronic Top-Up, provide a standby server. During normal processing, this does not share the workload, but can rapidly take over if there is a failure. The fail-over will be provided by the application layer. (As per the Horizon solution, the fail-over is provided by local heartbeats to the NPS.)

The NPS and APOP databases will continue to be backed up as per the Horizon system.

9.3.3 Risks

As the resilience models are largely unchanged for online services in the Horizon system, which work effectively, the approach for HNG-X is considered very low risk. Furthermore as improved resilience is provided by having full resilience at each data centre (also reducing inter campus traffic) the HNG-X approach is potentially a major improvement.

POL00399072

9.4 Disaster Recovery

9.4.1 Targets

The DR targets for online systems are documented in section 3.2.3.

9.4.2 Approach

The general approach to DR described in section 5.3 applies to the online services. An exact duplicate of the online systems is provided at the secondary data centre so a full level of service (capacity, performance and resilience) is provided.

Some of the online systems have external links (for example to A&L Santander and VocaLink for banking). These links are available at the secondary data centre. System state is held in the database (in the case of banking, in NPS), so the recovered systems can continue where they left off.

9.4.3 Risks

No specific DR risks have been identified.

This has already been tested with the Horizon NPS which used oracle RAC technology.

The Online interfaces are already resilient across data centres and, where applicable, maintain state within the NPS.

9.5 Conclusions

Tests and simulations show that the proposed solution will support the required transaction volumes. The underlying hardware and software are highly scalable. The Horizon systems show that the load balancing and resilience approach is sound.

POL00399072

10 Branch Database

10.1 Overview

The Branch Database stores transaction and other information captured by the HNG-X Counter systems. It is also used as a staging area to pass data to and from the counters.

The Branch Database is implemented as an Oracle 10gR2 Real Application Cluster (RAC) running over four nodes (servers).

10.2 Performance and Capacity

10.2.1 Targets

HNG-X Capacity Management and Business Volumes states important targets for the Branch Database.

- Transaction capture from **33,951** counter positions (see PA/PER/033 section 3.9 for the contractual number of counters it is necessary to support over each year of contract – however the actual current number of counters at time of writing is 30,307)
- An average transaction capture rate of **1068** transactions per second during the peak 5 minutes (1388 per second peak). Each individual item sale and each payment count as a transaction.
- An average session (basket) commit rate of 393 sessions per second during the peak 5 minutes (511 per second peak)
- A requirement to store recovery records at an average peak rate of **300** transactions per second during the peak 5 minutes (390 per second peak), to allow certain transaction types to be reversed if there are failures.
- Average Daily, Weekly and Monthly reporting loads peaking of **150** report requests per second during the peak 5 minutes (250 per second peak).
- An average peak logon/logoff rate of 120 per second during the peak 5 minutes (each will have 2 separate transactions into the BAL – potentially distributed across different BAL application servers each creating/updating session data held in the branch database) – (200 per second peak).

Since the branch database consists of four balanced Oracle RAC nodes in an N+1 configuration, each node should be capable of handling one third of the these targets, a quarter under typical circumstances.

10.2.2 Approach

To achieve the required performance and capacity, an Oracle Real Application Cluster (RAC) is used. This allows for the database work to be distributed over multiple nodes (servers).

An Oracle instance only has a single log writer (LGWR) process, which can become a bottleneck. Each node in an Oracle RAC runs a separate instance. Using a RAC therefore increases the number of LGWR processes, and reduces contention.

Horizon used a two-node RAC for the Network Banking Persistent Store (NPS) database.

The Branch Database will use a four node RAC. In comparison to a two-node RAC, this further allows the load to be distributed and reduces the LGWR bottleneck.

The diagram below shows the proposed 4-node cluster, each with its own Oracle instance.

POL00399072

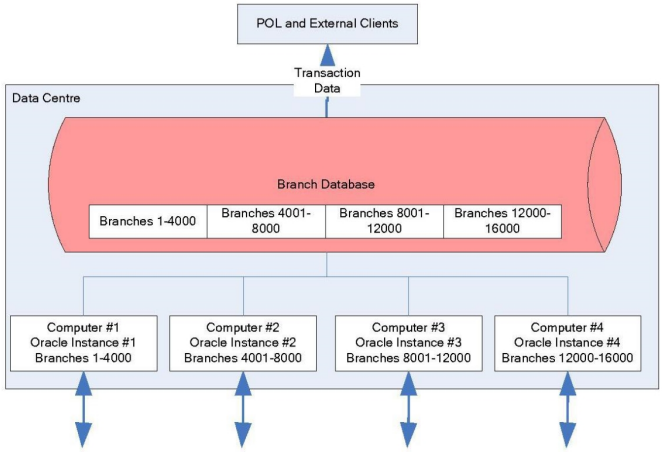

Figure [SEQ Figure * ARABIC] – Leveraging the Oracle Real Application Cluster

The Branch Access Layer manages connections to the appropriate nodes.

In addition to the RAC, the data is written to partitioned tables. The partitioning uses a pre-allocated hash code for each branch (the partitioning will be based on actual trading volumes for branches based on data provided by the capacity management services; this will create a very even split, improving performance and minimising the effect of a failure on a larger partition). This ensures that data for each branch is only accessed from one node, which reduces contention within the database and maximises caching.

Each Branch Database server has 2 very high spec quad core processors, offering significant processing capacity and in conjunction with the EMC DMX3 offers a very powerful processing capability.

10.2.3 Risks

The Branch Database is a core component of HNG-X. If it has any performance problems, these will have an immediate and significant impact on branch trading. This is different from Horizon which does not rely on online connections to a central database.

To mitigate this risk, a prototype Branch Database was built early in the HNG-X project, together with load generators. The results show that the proposed architecture can comfortably achieve the performance targets. The [Branch Database Prototype Report on Linux] and the [Branch Database Prototype Summary Report on Solaris] documents provide details of the tests and results.

In addition to the Performance Testing and modelling undertaken during the HNG-X development phase, the load on the branch database will rise gradually in line with branch migration. Hence pro-active trend analysis during the roll out phase will be able to spot any performance or capacity issues that may arise prior to the full workload hitting the branch database.

10.3 Availability

10.3.1 Targets

The Branch Database should support the Service Level Agreements set out in Schedule C1. The key availability targets are:

POL00399072

- The database must be available 24 x 7.
- There should be no single point of failure that causes the business service to fail.
- Failure of one node should not cause capacity problems.

10.3.2 Approach

The database is designed to run continuously. It need not be brought down for routine operations such as backup, for batch processes, or even for maintenance activities such as adding or reconfiguring storage. It will however be necessary to bring the whole branch database down (during agreed maintenance slots) so that oracle software upgrades can be applied.

The servers on which the Branch Database runs are inherently reliable. If one node fails, the other nodes automatically take over the load. The nodes and partitioning are designed so that the data for each branch continues to be accessed from only one node, even if one node fails. The responsibility for targeting a specific node for and individual branch using the pre-allocated hash code lies with the application code (BAL or batch Harvester process) rather than using any built in Oracle features. Following a node failure, requests are transferred to a pre-defined secondary node based on the hash code. Following recovery of the failing node, the workload does not automatically revert to the original node, but will revert at scheduled times (when volumes are low) to ensure an orderly failback process.

Hardware or software errors could cause data corruptions. It can take a long time to recover from these. To avoid this, on top of the standard backup process, Oracle Data Guard is used to maintain an additional standby copy of the database. If data corruptions are detected, this can be used to restore the database service within SLT.

10.3.3 Risks

Oracle 10gR2 RAC is an industry proven solution for high availability databases. Oracle RAC was used within Horizon for the NPS database. The procedures for maintaining 24x7 availability are well established.

10.4 Disaster Recovery

10.4.1 Targets

The service level target for recovery of the core data centre capabilities, which includes the Branch Database, is two hours.

10.4.2 Approach

The Branch Database uses the approach to DR outlined in section 4.3.

The secondary data centre will have a 4-node RAC with the same configuration as the primary data centre.

Transactional data will be replicated between the primary and DR data centres using a storage based synchronous replication system (EMC SRDF). This commits data changes to both data centres at the same time, and guarantees that no data is lost if there is a catastrophic failure at the primary data centre.

POL00399072

10.4.3 Risks

SRDF was used within Horizon, and there is a high degree of confidence in the solution. Historically there were some performance issues with SRDF, but these are well understood and have been modelled tested and simulated. Therefore there is not considered to be any problem.

Experience with DR tests of the Oracle RAC within Horizon showed that the database could be recovered within the two hour target.

10.5 Testing

Both performance and fail-over are tested as part of the HNG-X project process.

10.6 Conclusions

Testing and simulation have shown that the proposed database solution can cope with the design targets. Further testing showed that the database could cope with simulated transaction volumes up to four times higher than the contracted design limits (however as this was a limited functionality simulation, allowing for 'real world' scenarios, the system is estimated to be only 50% higher, which along with the necessity for resilience is seemed to be a reasonable solution).

POL00399072

11 Host Batch Services

11.1 Overview

HNG-X will include the main host batch systems used within Horizon.

A few changes will be made to these systems:

- Extracts and loads from the Riposte messaging product will be replaced with extracts and loads from the Branch Database.
- Solaris will be upgraded to version 10.
- Oracle will be upgraded to version 10gR2.
- Resilience will be provided within the data centre, with 2 systems running in an active/standby configuration.
- The secondary data centre will also have 2 systems running in an active/standby configuration to give full service resilience under DR.

The extra platforms have been sized to meet contracted volumes and design limits (actual levels of use are well known).

11.2 Performance and Capacity

The component targets for the migrated Horizon systems are the same as for Horizon. The systems are largely unchanged (beyond the interactions changing from Riposte via the Generic Agents to the Branch Database). Remaining scheduling, delivery and processing dependencies and SLTs must be adhered to. The specification of new servers is based on well understood usage. The updated systems will be fully volume tested.

Each server will have 4 Sparc 64 2.16GHz processors, as compared to the existing 8 Sparc 64 810MHz processors. Modelling and benchmarking shows that in raw processing power alone, this will give a significant performance improvement (25%) and in conjunction with Horizon volumes this will be adequate for the workload (and potentially provide faster processing). Furthermore the faster and more efficient EMC DMX3 (with reduced SRDF latency due to the proximity of the HNG-X data centres), will allow for much faster I/O which will further reduce the amount of CPU idle wait for I/O time, which on the Horizon solution amounted to a significant proportion.

11.3 Availability

11.3.1 Targets

A number of requirements deal with availability of the migrated Horizon systems.

- The Transaction Enquiry Service (TES) and the associated Query Application shall be available to receive transactions and to provide a query interface according to the Service Level Targets set out in the Data Centre Operations Service document. (SVC-784)
- Automated Payment (AP, AP-ADC) data files shall be available to the Electronic Data Exchange (EDG) according to the Schedule and Service Level Target defined in the Data Centre Operation Service document. (SVC-785)

POL00399072

- Reconciliation data files (REC) shall be available to Clients according to the Schedule and Service Level Target defined in the Data Centre Operation Service document. (SVC-786)
- Data from the Logistics Support System (SAPADS), such as Planned Orders, shall be available in Branches according to the delivery Schedules and Service Level Targets defined in the Data Centre Operation Service document. (SVC-787)
- Transaction Correction Records shall be available in Branches according to the delivery Schedule and Service Level Target defined in the Data Centre Operation Service document. (SVC-788)
- Data to the Logistics Support System (SAPADS), such as Pouch Receipts, shall be available in SAPADS according to the Schedules and Service Level Targets defined in the Data Centre Operation Service document. (SVC-789)
- Data to Post Office Systems (e.g. POLFS and POLMIS) shall be available according to the Schedules and Service Level Target defined in the Data Centre Operation Service document. (SVC-790)
- Global Reference Data shall be available at Branches according to the Schedules and Service Level Targets defined in the Reference Data Management Service Description. (SVC-791)

There are no changes to availability targets for the migrated Host Batch Systems. , These will be re-established on the standby platform within the standard fail-over times for manually initiated switch over. (Up to 1 hour during core time, and up to 2 hours for out-of-hours failures).

11.3.2 Approach

The host batch systems will run on a single server with a single Oracle instance. The Oracle transactions logging capabilities are used together with mirrored discs to protect against transient failures. Overnight the databases are backed up using a technique based on a third mirror break off solution. Batch schedules are re-runable.

In HNG-X, the resilience approach is modified to reflect the active-passive DR approach. Instead of the standby server being at the secondary data centre, an additional server is added at both data centres, and the fail-over will take place onto the standby server at the primary data centre. The switching technology will not change however. The location of the replicated or standby servers is transparent to the application.

The secondary data centre will provide the same configuration as the primary data centre. The secondary data centre will also have a standby server, to provide resilience if DR is invoked.

11.3.3 Risks

The host batch services use the same approach to resilience that was used in Horizon. The only change, housing the standby server in the same data centre, simplifies the solution and reduces risk.

The updated and ported components will be fully tested for resilience as part of the HNG-X test.

POL00399072

11.4 Disaster Recovery

11.4.1 Targets

The DR targets for host batch services are documented in section 3.2.3.

11.4.2 Approach

The host batch services are part of the general active-passive DR approach described in section 5.3. (In Horizon, they used an active-active approach.)

11.4.3 Risks

The change from an Active/Active to Active/DR configuration is little different to the Horizon failover mechanism from the perspective of the host batch systems. The only difference is that the service failover from the primary data centre to the secondary data centre becomes the DR process only. Under HNG-X service failover at either the primary or secondary data centre will be achieved by providing extra servers at the same data centre.

As such there is little additional risk specific to the host batch systems over the general risk in the technology required to switch the secondary data centre from a testing configuration to an Active configuration.

11.5 Testing

The HNG-X tests will include DR tests, showing that the secondary data centre can be switched from test and can recover the host batch systems.

11.6 Conclusions

The total risk to the migrated host batch systems is similar to that of a major Horizon release (such as Banking BI3, S50, S60, S70/S75, S80 or S90) where significant parts of the systems have been ported, upgraded to new database version or had significant functional change. The standard development and test cycle proposed for HNG-X will address these risks.

The main change to the systems is the data load mechanism, which changes from using Riposte messaging (via the Generic Agents) to the Branch Database. The HNG-X approach is more efficient, and therefore likely to improve performance.

POL00399072

12 Hydra

12.1 Overview

There are 2 main System Quality issues within Hydra:

- Providing solution system qualities for Horizon systems during and after the stages of data centre migration.
- Providing solution system qualities for the concurrent running phase.

In order to understand these, it is necessary to show the migration phases. The data centre migration will take place during weekends (weekend A is currently scheduled to be the final weekend – but otherwise they will be in the order shown with the principle systems migrating during that phase):

- Weekend A. POL-FS (to become POLSAP)
- Weekend B/C: Host Batch Services (TPS, APS, LFS, DRS, TES, RDMC, RDDS, KMA, DWH, Audit) and Online Services (Banking, DVLA, PAF, APOP, Track & Trace, MGRM, TLCM, Kahala)
- Weekend D: Audit

All Horizon platforms remaining at Bootle/Wigan (Correspondence Servers, Generic Agents, Routing Agents, Horizon PAF servers) will be retired once the entire estate has been moved onto HNG-X. In the intermediate period, when Horizon and HNG-X coexist, these systems will continue to run on their existing hardware and continue to use their existing Horizon availability and DR approach.

12.2 Performance and Capacity

Whilst SLTs will be relaxed over the weekends when the data centre moves actually take place, outside these weekends full performance and capacity capabilities are expected with full adherence to both Horizon and HNG-X STLs (as described in section 3 of this document and within the HNG-X Capacity Management and Business volumes document). While it is considered that there will be adequate capacity for these periods and their modes of operation, the potential risks are outlined below for each state:

- During moves: Where data centre moves involve the restoration of data (e.g. across the network or from tapes), there must be adequate performance and capacity to complete the work in the time available. Test moves will help indicate any potential problems or incorrect assumptions.
- Weekend A Complete: As all services within the POLSAP solution are largely self contained there is little risk.
- Weekend B/C Complete: The Host Batch Services will be communicating with the generic agents at a different site. There must be enough bandwidth, capacity and performance (especially by minimising network latency) to adhere to SLTs. However, due to the nature of the batch systems there is considered minimal risk in this state. This is because the current solution is designed to use Generic Agents at the other data centre, and as such the impact of latency on operation is minimal. Furthermore the Host (Dat) will be running on faster hardware (both CPU and I/O) and this will allow for more efficient processing. The VPN servers and correspondence servers will be communicating with online services at a different data centre. Also the NPS and Banking Authorisation Agents will be at a different data centre from the routing agents. There must be enough bandwidth, capacity and performance (e.g. minimising network latency) to adhere to the SLTs. As this state has the biggest impact on branch trading, this state is considered the biggest risk during the migration period. However the riskiest components for this mode of operation are

POL00399072

the DCS and ETU servers, which will not be migrating and therefore this risk is reduced. Furthermore, the Network Banking routing agents communicate with the authorisation agents via pipes minimising the effects of extra latency. Stateless Web Services using the HTTP protocol should not be significantly affected by running remotely. Further modelling, simulation, testing and analysis has helped confirm the viability of this. Volume testing will further prove it. In order to minimise the volume of remote traffic, the Horizon PAF servers at Bootle and Wigan will remain active and be used for all Horizon PAF requests.

- Weekend D Complete: The main feed from the Correspondence Servers to the migrated audit server will be now be remote. Whilst there is not perceived to be any problem from the effect of latency or bandwidth restrictions under normal utilisation, this service is not crucial to branch trading and would not cause any outage in the event of a problem requiring audit running in a "catch up" mode.
- Horizon PCI rollout Complete: As of Horizon PCI rollout DCS & ETU will now run remotely, however along with Banking transactions these will now use a different remote route – from the counter via the VPN Servers and CSMs at Bootle and Wigan and the HNG-X BAL servers into the authorisation agents (for banking as per the previous state). Whilst the state has more remote traffic with the addition of DCS and ETU requests, it is believed this state is lower risk with the removal of the Routing Agent dependency such that the remote communication is an HTTP request like the web servers.

12.3 Availability

Although availability SLTs will be relaxed over the weekends when the data centre moves take place, outside these weekends full contractual availability and resilience capabilities are expected, so that full adherence to both Horizon and HNG-X STLs (as described in section 3 of this document) will be provided.

Whilst systems and their services remain in the Horizon data centres, existing resilience will be provided with no risk (given the active/active Horizon resilience model HNG-X systems will already communicate with systems at both Horizon data centres).

Once systems and their services have moved to the HNG-X data centres, full resilience will be provided within the primary data centre therefore appearing seamless to the Horizon data centres. Any HNG-X platforms on BladeFrame will gain an extra level of resilience, as whilst on Horizon it currently would be necessary to repair or reprovision new hardware to replace failed hardware before being up to full resilience levels, on HNG-X any system will fail over to a spare blade and be at full resilience levels as soon as this has occurred.

Therefore as long as communications between the data centres operate correctly there is low risk.

Where systems require communications from both Horizon and HNG-X systems during the migration period (for example Batch Systems will be updated from both Riposte via the Generic Agents and the HNG-X Branch Database), in the event of a corruption or a failure it is assumed the Batch Services team will provide the recovery requirements.

12.4 Disaster Recovery

DR SLTs will only be relaxed over the weekends when data centre moves take place. Therefore should it become necessary to recommend DR operation for either of the primary data centres, it must be possible to adhere to the full DR targets for each state,

Should Horizon Primary (Bootle) fail, operation will be from HNG-X primary and Horizon secondary (Wigan).

POL00399072

Should HNG-X Primary fail, operation will be from HNG-X secondary and both Horizon primary and secondary (in its normal active/active configuration)

12.5 Conclusions

The Hydra period introduces risks both for migration activities (making sure they don't over run) and the dual running phases. Outside the move weekends full contractual obligations are required for all systems qualities and it is essential to make sure it is possible to adhere to these for each combination. Testing and analysis has confirmed the viability.

POL00399072

13 POLSAP - Post Office financial system

13.1 Overview

POLSAP is a new service that will comprise of existing services including POL-FS and SAPADS as a SAP based service for providing financial information for the Post Office business.

13.2 Performance and Capacity

POLSAP will include the Horizon POL-FS system, as well as the SAPADs system. The physical servers on which POLSAP will run will be on a Bladeframe located at the HNG-X data centres. In addition HNG-X will provide a new storage subsystem, as well as new Capacity Management and Backup solutions.

The pBlades are considerably faster than the Horizon systems used for POL-FS. In addition the new EMC DMX-3 storage is considerably more efficient and has a higher level of performance than the storage used for Horizon POL-FS, as well as the reduced SRDF latency between the HNG-X data centres; therefore the performance and capacity will be at least as good as for Horizon (and possibly significantly better).

13.3 Availability

Requirement SVC-782 states that POL-FS shall provide at least the availability set out in the *Data Centre Operations Service* document.

For most systems the POLSAP availability approach will be similar to the approach on Horizon POL-FS, but overall there will be significant improvements. The main POLSAP host will be on a pBlade, which if it fails, a cold standby server on a pBlade will reintroduce the system (previously if this system failed the entire service had to be recreated at the secondary data centre). Similarly with the IXOS archiving server which previously was a single point of failure, a cold standby server will be available. Within the data centre the EMC Centera is a single point of failure, but not for the overall solution as operation of these is switched to the secondary data centre without having to relocated the entire POLSAP system. All the other servers are provided in an N+1 configuration so that full service is available in the event of a failure to any server. Daily backups are taken of the SAP Oracle database with consistency checks ran at the weekend to confirm the backups do not contain any detectable corruptions.

The POLSAP availability approach is described in *SAP Resilience and Recovery Design* (EA/DES/001).

13.4 Disaster Recovery

The HNG-X POLSAP DR approach is exactly the same approach as for other HNG-X systems. Storage is synchronously replicated to the secondary data centre and the service introduced on the same infrastructure provided at the secondary data centre. The secondary data centre will continue to be used for testing purposes as in the Horizon solution (however as POLSAP QA is a hosted service the secondary site can also be considered as live).

The POLSAP DR approach is described in *SAP Resilience and Recovery Design* (EA/DES/001).

POL00399072

13.5 Testing

As the HNG-X POLSAP system retains the same active development/test paradigm as the Horizon POL-FS solution, relative Performance and Capacity remains unchanged, therefore any performance test results remain valid and testing will continue at the secondary data centre for this. Under the Horizon system, Business Continuity tests regularly took place, which successfully tested both the resilience and DR capabilities of the POL-FS system. These will continue under HNG-X for POLSAP.

13.6 Conclusions

The POLSAP solution provides a significant improvement over Horizon POL-FS for all system quality elements, with improved performance via faster hardware (processors, memory, storage), improved resilience through the availability of cold standby servers for servers that were previously single points of failure, improved backup solution,,capacity management and full capacity for DR systems.

POL00399072

14 Support Services

14.1 Overview

The support services can be divided into those for the technical support community and those for the business support community.

The technical support community include:-

- SMC in Stevenage/Bangalore
- MSS in Stevenage/Blackpool
- SMG in Blackpool
- ISD NT and UNIX in Belfast
- SSC in Bracknell
- Capacity Management Service in Bracknell

Providing services that include:-

- Software distribution
- Monitoring
- Remote access and diagnostic support
- DBA support

The business support community include:-

- Reference data team in Bracknell
- OBC team in Crewe
- Audit
- File Transfer

14.2 Performance and Capacity

In general the reduced complexity of the overall HNG-X solution as compared to Horizon is anticipated to reduce the operational overhead and thus the associated headcount dedicated to HNG-X. However the underlying support architectural components are sized to handle Horizon volumes so that migration and early HNGX roll out can be safely managed.

14.3 Availability

Each support community has diverse network routes to the HNGX data centres and a nominated DR site. Access is also available under controlled conditions from non FJ locations. The main point of support ingress to the data centre – the Secure Access Server – are multi provisioned.

POL00399072

Within the active data centre there is no single point of failure in the management components and the majority are provisioned in the BladeFrame and therefore partake in its resilience properties.

14.4 Disaster Recovery

The management components follow the main resilience and DR strategy in that the platforms datastores are replicated via the SAN network to the DR site. A full set of management platforms exist in the DR site and can thus be activated with the live datastore in a DR situation. Further design work will identify the exact point of recovery of the live datastore and whether post DR operational procedures will be required to bring the datastore up to the last management operation.

POL00399072

15 Requirements Reference

The following requirements are relevant to the System Qualities strategy.

15.1 Post Office Requirements

The following are the Post Office Release 1 baselined requirements associated with this System Qualities Architecture. The full set of all Post Office Release 1 'non-functional' requirements are held in the Fujitsu DOORS system and also in REQ/CUS/BRS/1049.

DOORS reference	Requirement Text
ARC-411	Where choices occur, the architecture shall take support for serving of customers as the priority, unless agreed otherwise with PO Ltd
ARC-438	Size of network shall be flexible up or down within the limits set in 'HNG-X Capacity Management and Business Volumes' document
ARC-442	For the data centres, Fujitsu Services shall document potential performance and capacity bottlenecks, and demonstrate cost-effective scalability at those critical points. Fujitsu Services will continue to share with Post Office capacity models, information and reports as per the current Capacity Management Service. Scalability of the various layers of the architecture are expected to be addressed as part design assurance process for HNG-X.
ARC-443	The resilience capabilities for the Data Centre and for the Data Centre Network, will be as specified in the Data Centre Operations Service Description.
ARC-444	Any single failure within the Data Centres shall not cause loss of any of the Business Capabilities & Support Facilities
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.
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.
SVC-779	Central Systems shall be available to provide the Core Solution during Core Hours according to Service Level Targets set out in the Data Centre Operations Service document.
SVC-780	Central Systems shall be available to provide the Core & Banking Solution during Core Hours according to Service Level Targets set out in the Data Centre Operations Service document.
SVC-781	Central Systems shall be available to provide the Core & all Other Services during Core Hours according to Service Level Targets set out in the Data Centre Operations Service document.
SVC-782	The Post Office Financial System (POLFS) shall be available during Supported Hours to an extent not less than that set out in the Data Centre Operations Service document

POL00399072

DOORS reference	Requirement Text
	within any Service Management Period.
SVC-783	The duration of a single period of any Outage of POLFS shall not exceed that set out the Data Centre Operations Service document.
SVC-784	The Transaction Enquiry Service (TES) and the associated Query Application shall be available to receive transactions and to provide a query interface according to the Service Level Targets set out in the Data Centre Operations Service document.
SVC-785	Automated Payment (AP, AP-ADC) data files shall be available to the Electronic Data Exchange (EDG) according to the Schedule and Service Level Target defined in the Data Centre Operation Service document.
SVC-786	Reconciliation data files (REC) shall be available to Clients according to the Schedule and Service Level Target defined in the Data Centre Operation Service document.
SVC-787	Data from the Logistics Support System (SAPADS), such as Planned Orders, shall be available in Branches according to the delivery Schedules and Service Level Targets defined in the Data Centre Operation Service document.
SVC-788	Transaction Correction Records shall be available in Branches according to the delivery Schedule and Service Level Target defined in the Data Centre Operation Service document.
SVC-789	Data to the Logistics Support System (SAPADS), such as Pouch Receipts, shall be available in SAPADS according to the Schedules and Service Level Targets defined in the Data Centre Operation Service document.
SVC-790	Data to Post Office Systems (e.g. POLFS and POLMIS) shall be available according to the Schedules and Service Level Target defined in the Data Centre Operation Service document.
SVC-791	Global Reference Data shall be available at Branches according to the Schedules and Service Level Targets defined in the Reference Data Management Service Description.
SVC-792	Branch Outages shall be within the levels and frequency described by the Service Level Targets in the Branch Network Service document.
SVC-800	The System shall support the volume of Business Transactions defined in the "HNG-X Capacity Management and Business Volumes" document at the Peak Rates identified.
SVC-802	The System shall support the volume of Administrative and Back Office Transactions defined in the "HNG-X Capacity Management and Business Volumes" document at the peak rates identified.
SVC-805	Transaction Time Benchmark measurement shall be carried out in accordance with the Service Management Service and the method described in the "Transaction Time Benchmarking, Joint Working Document" (CS).
SVC-806	Transaction Time Benchmark evaluation shall be conducted in accordance with the process "Counter Transaction Time Performance - measurement and results(CS/PER/046)." document
SVC-807	New Video Benchmarks shall be defined for a limited number (to be agreed) of representative Single and multi-product Baskets that will at least include an example of

POL00399072

DOORS reference	Requirement Text
	each Transaction Type.
SVC-808	During testing New Video Benchmarks shall be conducted on Baseline Horizon and equivalent Video Benchmarks shall be conducted on the System, in accordance with the process described in "HNG-X Test Strategy - HX/STR/001 and assessment made according to the Video Benchmarking Tolerances agreed therein.
SVC-809	System Measurement of Counter Performance will be introduced in the System as a future alternative to Video Benchmarking.
SVC-810	System Measurement of Counter Performance may be calibrated with the Video Benchmark results obtained for the System for equivalent sequences of operation.
SVC-815	The end-to-end performance of Banking and related Transactions shall be no worse than that set out in the Service Level Target described in the Data Centre Operations Service document.
SVC-815	The end-to-end performance of Banking and related Transactions shall be no worse than that set out in the Service Level Target described in the Data Centre Operations Service document.
SVC-816	The end-to-end performance of the Basket Settlement transaction shall be no worse than that set out in the Service Level Target described in the Data Centre Operations Service document.
SVC-817	The Time taken to complete Branch Administrative or Back Office Transactions shall meet Design Targets (if specified) in "HNG-X Capacity Management and Business Volumes"
SVC-818	The time taken for Stock Unit and Branch Reports shall meet Design Targets (if specified) in "HNG-X Capacity Management and Business Volumes"
SVC-819	The time taken to retrieve Help pages shall meet Design Targets (if specified) in "HNG-X Capacity Management and Business Volumes"
TR455	The speed of response for transactions in the HNGx training system shall be broadly equivalent to the response time of the HNGx live system.
SEC-3174	{CISP 8.5.1j} Test systems shall only share logical network connection with operational systems in carefully controlled circumstances. Test systems shall be configured to connect in this manner for the minimum duration necessary to support testing. The logical connection shall only be permitted after an assessment has confirmed that live operation will not be adversely impacted or as otherwise agreed by Post Office Limited.

15.2 Customer Service requirements

The following are the Customer Service Release 1 baselined requirements associated with this System Qualities Architecture. The full set of all Customer Service Release 1 requirements are held in the Fujitsu DOORS system and also in REQ/CUS/BRS/1050.

POL00399072

CS-BR-148 (derived from GEN1 in RM/CDE/003 v0.4)	The data centre configuration when running the HNG-X application will have one active data centre and one DR data centre There will be Test Environment(s) including testing of network. Assumption: The DR data centre will be used for testing (See Comments re network testing being not all possible within the DR data centre). It is possible within the DR centre to switch from test to live within the service levels agreed for each component.
CS-BR-150 (derived from NUN4 in RM/CDE/003 v0.4)	The Data Centre disaster recovery model should ensure compliance with contractual SLTs & OLTs giving a switch-over to the Disaster Recovery site within 2 hours for core systems and Network Banking, 48 hours for POLFS & 5 hours for other systems. Additionally, server build standards should be reviewed to ensure that a server rebuild is done efficiently and is automated.
CS-BR-22 (derived from SCD 139 in RM/CDE/003 v0.4)	Application Development when designing the Reference Data Distribution System shall take account of the network bandwidth and ensure that the overnight distribution of Reference Data and Code etc does not exceed the practical network bandwidth and impact the service provided by POL. Note: Problems in past include Counters not being able to work due to other Counters downloading large amounts of Reference Data at the same time.
CS-BR-179 (derived from SCD 40 in RM/CDE/003 v0.4)	The infrastructure should not have any single points of failure (e.g. with the network, Branch connection to the telephone exchange) unless specifically agreed with POL (e.g. covering specific Branches or where there is an acceptable level of built-in resilience within the system). For Branch Infrastructure, the Branch Counters themselves are excluded.
CS-BR-180 (derived from SCD 041 in RM/CDE/003 v0.4)	Service components should have resilience and failover solutions sufficient to support the HNG-X SLTs and OLTs and performance metrics. The resilience model for the HNG-X Infrastructure should be provided to, and agreed with the POA CS Business Continuity Manager.

.