Document Title: [TITLE * MERGEFORMAT]

Document Type: [DOCPROPERTY "Document Type" * MERGEFORMAT]

Release: N/A

Abstract: [COMMENTS * MERGEFORMAT]

Document Status: DRAFT

Author & Dept: [AUTHOR * MERGEFORMAT]

Internal Distribution:

External Distribution:

Approval Authorities:

Name	Role	Signature	Date
David Chapman	Systems Qualities~~y~~ Architect		
Geof Slocombe	Infrastructure Design		

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

0 Document Control

0.1 Table of Contents

0.2 Document History

Version No.	Date	Summary of Changes and Reason for Issue	Associated Change - CP/PEAK/PPRR Reference
0.1		Initial Release	CP4143
0.2	23-Nov-2007	Incorporated initial comments and updates from meeting with Will Russell and John Halfacre and reviewers.	
0.3	01-May~~Feb~~-2008	Incorporate comments Added ref to DEV/GEN/SPE/0007 for PIM failure	
0.4	23-Jul-2008	Incorporate comments and reformat last section to use as a summary table in the overview document. Incorporate changes to EST platform services after informal comments from Pat Carroll. Changes to VPN platforms Update to reflect latest network design. Reference made to RAD changes KMS moved into 'frame with software random number generator	 CP0219 CP0097 CP0184 CP4506

0.3 Review Details

*** NB: NOT SUBJECT TO APPROVERS & REVIEWERS ROLE MATRIX**

Review Comments by :	Fri~~Mon~~day, 29~~16~~0th August~~December May~~ 2008~~7~~
Review Comments to :	[HYPERLINK "mailto:Edward.ashford GRO] & [HYPERLINK "mailto:PostOfficeAccountDocumentManagemen GRO]

Mandatory Review	
Role	Name
Development~~Business Continuity Manager~~	Paul Stewart, Joseph Diffin~~Tony Wicks~~
Architecture~~Systems Quality Architect~~	David Chapman~~David Chapman~~
Development~~SSC Manager~~	Adrian West~~Mik Peach*~~
SSC	Mik Peach
Business Continuity	Tony Wicks
Migration Architect~~Data centre Service Manager~~	Brian Ridley~~Peter Thompson*~~
System Test	John Rogers
Optional Review	
Role	Name
~~Apologies if roles are not quite right; new announcement wasn't too explicit~~	
Programme Manager~~Network Architect~~	Phil Day~~Mark Jarosz~~

Formatted Table

Formatted Table

Applications Architecture	Dave Johns
Architect	Jason Clark
Security Architect	Jim Sweeting
Test Design	Peter Robinson
Test Design	George Zolkiewka
Head of Service Management	Steve Denham
Head of Service Change & Transition	Graham Welsh
Service Support	Peter Thompson
Service Network	Alex Kemp
Data Centre Migration	Caroline Montgomery
Infrastructure Design	Geof Slocombe
Testing	Peter Dreweatt
SV&I Manager	Sheila Bamber
Tester	Hamish Munro
RV Manager	James Brett (POL)
VI & TE Manager	Peter Rickson
HNG-X Acceptance & Risk	Wayne Roberts (POL)
Integrity Testing	Alan Child
Integrity Testing	Michael Welch
Core Services	Mark Walsh
Core Services	Andrew Gibson
Business Architect	Gareth Jenkins
~~Security Architect~~	~~Jim Sweeting~~
~~Systems & Estate Management Architect~~	~~Ian Bowen*~~
~~Head of Engineering~~	~~Barbara Perek~~
~~Quality~~	~~Jan Holmes*~~
~~Security~~	~~William Membery*~~
~~CTO~~	~~Matt AdbyGiacomo Piccinelli~~
~~Requirements & Architecture Manager~~	~~Martin CairJohn Lake~~
~~Implementation & Transition PM~~	~~Martin BrettCaroline Montgomery~~
~~Requirements & Architecture (Architecture)~~	~~Dave Johns~~
~~Requirements & Architecture (Requirements)~~	~~Dave Cooke~~
~~Head of Infrastructure~~	~~Dave SackmanGeoff Slocombe~~
~~Head of Applications & Integration~~	~~Graham AllenAdrian West~~
~~Head of Test~~	~~Pete Dreweatt~~

~~Programme Manager~~	~~Phil Day~~
~~Infrastructure Project Manager~~	~~Dean ParsonsMike Brady~~
~~Platform & Storage Architect~~	~~Jason Clark~~
~~POLFS Migration Design~~	~~Chris Credland / Joseph Diffin~~
~~Network & DNS Design~~	~~Dave Haywood~~
~~Security Architect~~	~~Jim Sweeting~~
~~Network Design~~	~~Dave Tanner~~
~~Network Design~~	~~Andrew Oram~~
~~Test Architect~~	~~Peter Robinson~~
~~SSC & Time Services Design~~	~~Pat Carroll~~
~~Design~~	~~Tom Northcott~~
~~Branch Database Architecture~~	~~Nasser Siddiqi / Gareth Jenkins~~
~~IS Project Manager~~	~~Pat LywoodMark Walsh~~
~~Unix/DBA/NT Support~~	~~Andrew Gibson*~~
~~SMC~~	~~Ian Cooley~~
~~Network Support~~	~~Dave Jackson~~
~~Migration Architect~~	~~Brian RidleyJeremy Worrell~~
~~**ALL PPD OWNERS**~~	~~Brownsword I, Carroll P, Comer G, Gosnold J, Haywood D, Holmes A, Kalenic Z, Latif S, Macdonald D, Mills C, Mital U, Morris T, Noad P, Olubor J, Siddiqi N, Stewart P, Stock M, Swain J, Sweeting J,Taylor N, Tomlinson P, Walton L, White N, Williams A, Wright M~~
Issued for Information – Please restrict this distribution list to a minimum	
Position/Role	Name
~~HNG-x~~HNG-x Implementation Manager, RMG	Will Russell*
~~Programme & Project Manager – Solutions Group~~	~~Andy HeathMike Jackson~~

(*) = Reviewers that returned comments

0.4 A

Associated Documents (Internal & External)

Reference	Version	Date	Title	Source
PGM/DCM/TEM/0001 (DO NOT REMOVE)	1.0	13/6/06	Fujitsu Services Post Office Account ~~HNG-X~~HNG-X Document Template	Dimensions
ARC/SOL/ARC/0001			~~HNG-X~~HNG-X Overall Solution Architecture	Dimensions
REQ/CUS/STG/0001			~~HNG-X~~HNG-X Migration Strategy - Agreed Assumptions and Constraints	Dimensions
COM/CUS/SCH/0011	1.0	31/8/06	Schedule B2 - Business Continuity (CCN 1200)	Dimensions
COM/CUS/SCH/0014	2.0	25/1/07	Schedule B3.3 - HNG-x Central and Telecommunications Infrastructure. (CCN 1200)	Dimensions
~~FS/BAU/SPP/001~~SVM/SDM/SIP/0001			Business Continuity Framework	Dimensions
SVM/SDM/PLA/0001			HNG-X Business Continuity Support Services Test Plan	Dimensions
SVM/SDM/PLA/0002			HNG-X Business Continuity Services Test Plan	Dimensions
SVM/SDM/PLA/0003			~~HNG-X~~HNG-X Business Continuity Operational Test Plan	Dimensions
SVM/SDM/SD/0003	1.1	5/3/08	Data Centre Operations Service: Service Description	Dimensions
ARC/GEN/REP/0001			~~HNG-X~~HNG-X Glossary	Dimensions
ARC/PER/ARC/0001			System Qualities Architecture	Dimensions
ARC/SEC/ARC/0003			Security Architecture	Dimensions
ARC/PPS/ARC/0001			Platforms and Storage Architecture	Dimensions
ARC/APP/ARC/0005			Branch Database Architecture	Dimensions
ARC/APP/ARC/0008			Online Service Architecture	Dimensions
ARC/SYM/ARC/0001			~~HNG-X~~HNG-X System and Estate Management Overall Architecture	Dimensions

Formatted Table

Formatted: Not Highlight

Formatted: Not Highlight

Formatted: Not Highlight

Formatted: Not Highlight

Formatted Table

Formatted: Not Highlight

Formatted: Not Highlight

Formatted: Not Highlight

Formatted: Not Highlight

Formatted Table

Formatted: Not Highlight

Reference	Version	Date	Title	Source
ARC/SYM/ARC/0003			~~HNG-X~~HNG-X System and Estate Management Monitoring	Dimensions
ARC/GEN/STD/0002			[TITLE * MERGEFORMAT] [NB Ref may change]	Dimensions
TST/SOT/HTP/0006			~~HNG-X~~HNG-X: ITU VI Business Continuity High Level Test Plan	Dimensions
DES/MIG/HLD/0001			~~HNG-X~~HNG-X Migration High Level Design for Branches	Dimensions
DES/MIG/HLD/0002			~~HNG-X~~HNG-X Migration High Level Design for ~~Data Centre~~Data Centres	Dimensions
DES/SYM/HLD/0015			~~HNG-X~~HNG-X Backup and Recovery High Level Design	Dimensions
DES/APP/HLD/0020			~~HNG-X~~HNG-X Branch Database Design	Dimensions
DES/PPS/HLD/0009			~~HNG-X~~HNG-X Platform Type List	Dimensions
DES/PPS/HLD/0007			Storage High Level Design	Dimensions
DES/PPS/HLD/0003			Active Directory High Level Design for ~~HNG-X~~HNG-X	Dimensions
DES/PPS/HLD/0025			~~HNG-x~~HNG-x BladeFrame and PAN High Level Design	Dimensions
DES/SYM/HLD/0001			MON - Supporting Platforms	Dimensions
DES/SYM/HLD/0002			MON - Supporting Agents	Dimensions
DES/SYM/HLD/0003			MON - Horizon Support	Dimensions
DES/SYM/HLD/0004			MON - Usability	Dimensions
DES/NET/HLD/0006			Domain Naming System High Level Design	Dimensions
DES/NET/HLD/0007			~~HNG-X~~HNG-X SAN High Level Design	Dimensions
DEV/INF/LLD/0004			~~HNG-x~~HNG-x Storage Design	Dimensions
DES/NET/HLD/0008			~~Data Centre~~Data Centre LAN Design	Dimensions
DES/NET/HLD/0009			~~HNG-X~~HNG-X Wide Area Network Design	Dimensions
DES/NET/HLD/0014			~~HNG-x~~HNG-x Branch	Dimensions

Formatted Table

Formatted Table

Reference	Version	Date	Title	Source
			Access HLD	
DES/NET/HLD/0015			~~HNG-x~~HNG-x Transit LAN Design	Dimensions
DES/NET/HLD/0010			~~HNG-x~~HNG-x Branch Router Design	Dimensions
DES/NET/DPR/0002			Design Proposal for Branch Router	Dimensions
DES/NET/HLD/0010			Branch Router Network High Level Design	Dimensions
DES/NET/HLD/0012			Network Management System High Level Design	Dimensions
DES/NET/HLD/0013			Time Synchronisation High Level Design	Dimensions
DEV/INF/LLD/0041			HNG-x Data Centre LAN LLD	Dimensions
DEV/INF/ION/0002			HNG-x VLAN Mappings	Dimensions
DEV/GEN/ION/0001			FTMS Configuration for the TIP Gateways	Dimensions
DEV/GEN/ION/0002			FTMS Configuration for the EDG Gateways	Dimensions

Formatted Table

Formatted Table

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

0.5 Abbreviations

Formatted Table

Abbreviation	Definition
A&L	Alliance & Leicester
AD	Active Directory. An implementation of lightweight directory access protocol (LDAP) that provides central authentication and authorization services.
ADSL	Asymmetric Digital Subscriber Line
AP	Automated Payment
AP-ADC	Automated Payment – Advanced Data Capture
APOP	Automated Payment Out-Pay
APS	Automated Payment Service. Also used for the name of the Oracle database that supports this service.
ATM	Asynchronous Transfer Mode. A form of network communication that does not rely on each packet of data being acknowledged before the next is transmitted.
BAL	Branch Access Layer

BCT	Business Continuity Test
CDL	EMC Clariion Disk Library. Brand name for a type of VTL.
C&W	Fujitsu Services Backbone Network, a private managed network operated by Fujitsu Services
CAPO	Card Account Post Office. An external client providing banking services.
Cisco ACE	Application Control Engine. ACE polls a number of possible service providers and advertises an external virtual address for the service based on pre-determined selection criteria. This is a common means of providing resilience for web services. The full term is used to avoid ambiguity as ACE is a common abbreviation.
COTS	Commercial off the shelf
DC	~~Data Centre~~Data Centre
DCS	Debit Card Service
DHL	Definitive Hardware List. Also a courier company used to ship components to site.
DHS	Definitive Hardware Store
DMX	EMC "Direct Matrix". DMX3 is the latest generation of the EMC Symmetrix range of disk arrays.
DMZ	Does actually stand for de-militarized zone, but is used here in the networking security and firewall sense to mean an intermediate zone between two networks which affords protection to the main servers. For example one would find a mail proxy in the DMZ between the internet (the Rest of the World) and the mail server on the ~~data centre~~Data Centre LAN.
DNS	Domain Name Service
DR	Disaster Recovery
DVLA	Driver and Vehicle Licensing Authority
DWDM	Dense Wave Division Multiplexing. A means of passing many optical signals simultaneously along a single fibre optic path. Typically used for long distance links where the installation cost is very high, it allows several customers to be offered a service where each appears to have their own dedicated link.
ECC	EMC Enterprise Control Centre. A Storage Management System.
EDG	Electronic Data Gateway
EMC	Disk Array manufacturer
EPOSS	Electronic Point of Sale Service~~; HNG-X service that supports Retail functions in Branches. I'm not convinced that the term is relevant to HNG-X.~~

Formatted Table

Formatted Table

Formatted Table

ETU	Electronic Top-Up
FAD	Finance Accounts Division, part of Post Office Ltd
FC	fibre channel (see glossary)
FRU	Field-Replaceable Unit. A part or component of a device or system that easily can be replaced by a skilled technician without having to send the entire device or system to be repaired.
~~C&W alpha order?~~	~~Fujitsu Services Backbone Network, a private managed network operated by Fujitsu Services~~
FTMS	File Transfer Management Service; ~~HNG-X~~HNG-X process that provides configurable file transfer services between Horizon and Post Office Ltd. Clients. Services available include data compression and encryption
GPRS	General Packet Radio Service. A generic term covering a number of technologies used to provide internet connectivity for example using a mobile phone.
HA	High Availability
HBA	Host Bus Adapter
HLD	High Level Design
~~HNG-X~~HNG-X	Horizon Next Generation – Plan X
IP	Internet Protocol
IPMP	IP multi-pathing. A Solaris driver to provide resilient network connections.
ISDN	Integrated Services Digital Network
LAN	Local Area Network
LINK	The organisation responsible for branded and shared network of cash machines and self-service terminals of certain member banks and building societies in the UK, which enables services from one member bank or building society to be available at cash machines of all member banks and building societies.
LPAN	Logical Processor Area Network. A subset of a PAN (see below)
LST	Live System Test. A pre-production test rig built and operated like live.
LTPDB	Long Term Performance Data Base. A repository for capacity planning and performance management statistics.
LUN	Logical Unit Number
MTAS	MID / TID Allocation System.
NAS	Network-attached Storage. An appliance based mechanism for presenting storage as generic shares. NAS is very flexible, and allows many servers to share a file store, but does have

Formatted Table

	performance limitations compared to dedicated storage.
NBS	Network Banking Service – one of the A&L, CAPO or LINK Authorisation Services
NBU	Symantec NetBackup
NIC	Network Interface Card
NPS	Network ~~Banking ~~Persis~~is~~ten~~ce~~t ~~Service~~Store
ntp	Network Time Protocol
Oracle RAC	Oracle Real Application Cluster. The full term has been used rather than the abbreviated "RAC" as this might be confused with the Network Banking Request Authorise Confirm model.
OS	Operating System
PAF	Postal Address File
PCI	Payment Card Industry. A consortium of companies which has developed standards for processing electronic payments, especially relating to the security and storage of card-holder details. This term will usually appear as "PCI compliance"
PCI	Peripheral Component Interconnect. A replacement developed by Intel for the early personal computer bus which has migrated to ~~data centre~~Data Centre class computers because of the ability to share mass-produced peripheral components (e.g. network interface cards) and the software drivers for those components. Although this seems a retrograde step it should be remembered that today's desk top computers are considerably more powerful than a super computer from 1985 which cost many millions of pounds. This term will usually appear as "PCI bus"
PAN	Processor Area Network. A term used to describe the overall set of BladeFrame resources. The alternate meaning of Primary Account Number is not used in this document.
PIM	eGenera Power Input Module
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 based in Fujitsu Services ~~Data Centre~~Data Centres.
POL-MIS	Post Office Ltd Management Information System based in the Northern ~~Data Centre~~Data Centre
RAD	Real-time Active Dashboard. Event filtering service based on Tivoli Netcool used to provide a business view of events.

Formatted Table

RHEL	Red Hat Enterprise Linux
RPO	Recovery Point Objective
RTO	Recovery Time Objective
SAN	Storage Area Network based on fibre channel protocol
SAS	Secure Access Server
SCW	Security Configuration Wizard
SDH	synchronous digital hierarchy standard developed by the International Telecommunication Union (ITU), documented in standard G.707 and its extension G.708
SLT	Service Level Target
SOX	Sarbanes-Oxley
SPOF	Single Point Of Failure
SRDF	EMC proprietary term for storage replication between sites "Site Remote Data Facility".
SRRC	Service Resilience and Recovery Catalogue. A document describing the failure modes, business impact, events raised, and recovery mechanism for each service.
SSH	Secure Shell
SSL	Secure Socket Layer
TES	Transaction Enquiry Service
TPS	Transaction Processing Service; Horizon service that formats data for transmission to POL-MIS and POL-FS and other places
TWS	Tivoli Workload Scheduler. A batch scheduling system. This is the new name for the latest version of the Unison Maestro scheduler used in Horizon.
V&I	Volume and Integrity Test Rig. A full-scale test rig used for performance testing and non-functional testing.
VIP	Virtual IP
VLAN	Virtual LAN. Larger switches can operate as if they were a number of separate logical switches. This allows a larger number of services to be taken down when a switch fails.
VSAN	Virtual SAN. Ditto. Although IP and fibre channel differ, the approach used by Cisco has strong analogues for virtualisation.
VPN	Virtual Private Network
VTL	Virtual Tape Library
WAN	Wide Area Network

Formatted Table

Formatted Table

Formatted Table

0.6 Glossary

Term	Definition
trunking, trunked	Used of network connections. Trunked describes a switch to switch interface, where several VLANs may be shared between the switches using a single physical network port. This allows great flexibility in shaping the physical network architecture using core-edge toplogies, and in effect this is how BladeFrame is leveraging network virtualisation.
kB, MB, GB, TB mbps, gbps	There are no ISO standards for "byte" and "bit". This document will use the terms on the left to indicate kilobytes, megabytes, gigabytes, terabytes where: kB = 1024 bytes MB = 1024 kB GB = 1024 MB TB = 1024 GB Note that disk salesmen often quote gigabytes as 1000 MB mbps = megabit per second gbps = gigabit per second Normal Ethernet speeds are 10/100/1000 mbps. 100 mbps roughly equates to 7 MB/s Normal fibrechannel speeds are 1/2/4 gbps 1 gbps roughly equates to 70 MB/s The "rounding errors" are down to things like frame headers, checksums and queuing on shared links. Other common network speeds are 34 mbps (E3 or T3 leased line) and 155 mbps (ATM over SDH). Readers who wish to explore this further should follow the links on Wikipedia which are easily accessed via the page on Fibre channel. A knowledge of these terms is not essential for reading this document.
fibre channel	**Fibre Channel** is a gigabit-speed network technology primarily used for storage networking. Despite common connotations of its name, Fibre Channel signaling can run on both twisted pair copper wire and fibre-optic. It is also possible to run IP connections over fibre channel, but the advent of gigabit ethernet has made this less common. http://en.wikipedia.org/wiki/Fibre_channel
switched fabric	Commonly abbreviated to fabric, and specifically used here to mean fibre channel switched fabric (FC-SW). A fabric is a network of devices connected by switches. Each is identified by a worl~~k~~d wide node name (WWNN). Many VSANs can be created in each fabric. In a typical resilient deployment **mirrored fabrics** are used where each VSAN has a counterpart in the other fabric, and resilient paths between hosts and storage devices are provided through both of the mirrored fabrics.
N+1	This is a term commonly used in describing resilience. If two servers are required to provide adequate performance, and any server can perform all the tasks required,

Term	Definition
	then deploying three servers provides one more than is required (N+1), and the failure of a single server does not reduce performance below the level required.
Hydra	The term hydra is used to avoid ambiguity which the term "Horizon Branch Services" would otherwise introduce during migration. During the counter migration phase a "Hydra" component will remain to support Horizon counters. This will continue to operate Active/Active, but will be removed once the last counter has migrated to the HNG-x application. Hydra ends when there are no more Horizon counters, but note that SYSMAN2 is required until all NT systems are upgraded or removed.
stateless	Used specifically in this document to mean a system which does not need to store information about its previous state. The opposite term stateful is used to define a system which has some data that would need to be recovered in the event that the system was lost. In designing recovery solutions stateless servers are very flexible as they merely need to be rebuilt or restarted, whereas stateful systems need to be failed over or recovered.

0.7 Changes Expected

Changes
Network and migration design are still work in progress, and both have an impact on DR and HA design, however no substantial changes are expected. ~~The list of platform types is still subject to change.~~The design of BMX is not complete.

0.8 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.

0.9 Copyright

1 Introduction

~~I still need to produce the~~An overview is provided which details the platforms and summarises the DR and resilience models in tabular form.

Formatted: Not Highlight

DES/GEN/ION/0002	TECHNICAL OVERVIEW OF DESIGN FOR BUSINESS CONTINUITY	*THE DOCUMENT IS INTENDED TO INFORM APPLICATION DESIGNERS ABOUT THE RESILIENCE AND DR SOLUTIONS, INCLUDING BACKUP & RECOVERY, SO THAT THE SYSTEM QUALITIES CHAPTER IN THE APPLICATION HLD PROVIDES THE INFORMATION REQUIRED BY THE BACKUP DESIGNER AND THE DR AN*

~~1.1~~ Business continuity is a complex subject. This document is not intended to cover every aspect of

Formatted: Bullets and Numbering

business continuity design, planning or operation, but to outline the various technologies chosen, to cover all services at a high level, and to direct the reader to those areas of architecture, design and business continuity planning where the appropriate detailed view may be found.

~~1.2~~1.1 Scope

This document covers two aspects of the high-level architectural design as it relates to the **steady-state** hosting of the ~~HNG-X~~HNG-X solution – resilience and disaster recovery.

This document describes the key architectural elements which support both resilience and disaster recovery, and defines a number of classes of hosted service.

For each class of service, this document outlines the high level procedures for recovering from failures, and for recovering service back to the primary ~~data centre~~Data Centre in case of a disaster which forces service to the secondary ~~data centre~~Data Centre.

Disaster recovery during the ~~data centre~~Data Centre migration will be covered at a high level. It is expected that related groups of services will be migrated together, and that each group of services may use Horizon or ~~HNG-x~~HNG-x DR mechanisms independently.

~~1.3~~1.2 Not in Scope

This document assumes a general knowledge of the ~~HNG-X~~HNG-X architecture.

Backup and Recovery design covers the recovery from data corruption.

Enhanced Agent and Correspondence Server Resilience and Recovery (EACRR) for Hydra will be covered in a separate design document DES/PER/HLD/0003 HNG-x Branch Trading Resilience Design.

Resilience of an application is the responsibility of the application designer. Guidelines for designing a resilient service and a number of technology options are presented in this high level design, and a summary of the approach adopted by each service, but the reader is referred to the application high level design for further details.

This document does not cover the provision of secondary sites for staff, such as the SMC at Stevenage and the SSC at Bracknell. The need to provide such provision is covered by the Schedules for those services (this includes the Development and Integration functions) and the only requirement that is in scope for this document is that the network architecture should allow any of these sites to connect to either data centre. These are covered by the Business Continuity Plans for each service, and are substantially unchanged from Horizon.

The requirements of these secondary sites are principally driven by the need to keep operating the live service (SVM/SDM/SD/0003) and the decision as to whether we require a "warm" or "cold" option is driven by risk and cost rather than design. As far as possible I have tried to encourage access from standard corporate workstations for support staff which reduces the cost of these provisions and maximises flexibility, but there are clearly cases, either for reasons of bandwidth or security, where this is not possible.

~~1.4~~1.3 Changes from Horizon

In Horizon the ~~data centre~~Data Centres operated in an Active/Active mode, the counters connect to the Riposte asynchronous messaging store. They could connect to servers in either site

In ~~HNG-X~~HNG-X the ~~data centre~~Data Centres operate in a Production/Test mode, where the Standby site is used for system testing. This is facilitated by the deployment of BladeFrame which simplifies the definition of virtual sets of servers that exist in the SAN rather than in a physical implementation.

2 Resilience and Disaster Recovery

These two topics, supported by backup and recovery and estate monitoring, define how the overall Business Continuity Plan is to be implemented. SVM/SDM/PLA/0003 is the Business Continuity Plan.

Requirements are principally derived from the following sources:

COM/CUS/SCH/0011 - Schedule B2 - Business Continuity

COM/CUS/SCH/0014 - HNG-x Central and Telecommunications Infrastructure

SVM/SDM/SD/0003 - Data Centre Service: Service Description

Note that other services may have expressed requirements in DOORS which relate to these requirements. This document does not attempt to reconcile all such requirements, and for a particular service the service HLD should perform this function.

Although there are many other schedules and the scope of business continuity is considered in its widest sense the use of a web-services architecture with a central data repository means that the Data Centre service description serves to capture the requirements of the end-to-end service quite neatly.

Schedule B2 is met by the Business Continuity Framework and Data Centre Service Description. No further direct reference will be made to it here.

ARC/SOL/ARC/0001 Overall Solution Architecture, especially Chapter 6, provides an architectural context for this design.

ARC/PER/ARC/0001 System Qualities Architecture summarises the requirements of the service for availability and performance.

DES/SYM/HLD/0015 Backup & Recovery High Level Design deals with recovery from data corruption at an application level. This is outside the scope of DR, as the corruption may have also corrupted the data at the secondary site and DR is not an appropriate means of recovery. It is generally part of application resilience.

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

2.1 Definitions

Resilience relates to elements of the hosting solution which provide tolerance to faults within the primary ~~data centre~~Data Centre; the main design goal being that any single failure will not prevent the application from continuing to work within the primary ~~data centre~~Data Centre

Disaster Recovery relates to the elements of the hosting solution which allow hosting of the solution from a secondary ~~data centre~~Data Centre in event of a catastrophic failure at the primary ~~data centre~~Data Centre~~, with very little or no loss of data~~

2.2 Approach to Resilience

The aim of providing resilience within the ~~data centre~~Data Centre is to allow the solution to survive the failure of a single part of the infrastructure. As far as possible, this should be extended so that multiple failures can be withstood, to a commercially reasonable level.

In addition to this, the design of the resilience should also provide for simplified maintenance and minimisation of any outage required to replace or fix failed parts.

In the event of a major environmental problem, such as a major fire at either site, the remaining site is designed to be N+1 resilient for all business critical components.

Section 6.2.3 details single points of failure that are declared in the design. These are covered by risk management, and are typically services that can be bought if a disaster really occurs or which would be too expensive to justify the risk that is mitigated.

2.3 Approach to Disaster Recovery

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

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

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

Formatted: Font: Italic

2.4 Approach to Development and Testing

For any DR solution at least 80% of the success depends on procedures and processes rather than the technical implementation. In addition key operational staff need to stay familiar with these processes, and the processes themselves need to be kept up to date as the solution develops during its life cycle, or as problems and work-rounds are found during live use or testing.

Fujitsu Services operate a series of business continuity tests under the management of the Business Continuity Manager reporting to the Business Risk Manager and various Service Managers. These tests are run in the live environment, and are designed to be as realistic as possible without putting the service deliberately at risk.

The ~~HNG-x~~HNG-x solution has a lot of elements that will be very familiar to Horizon support staff, such as the Legacy Batch Server, Oracle RAC based HA databases, and the POLFS system. There are other elements such as network banking agents which look similar but have subtle differences, and yet other areas such as the deployment of BladeFrame server virtualisation, deployment of Oracle DataGuard and Streams, and operation in Production/Test mode which are new.

A separate project known as "Pathfinder" is being performed in situ in Belfast to prove the migration approach for the POLFS service. Initially this will connect the four sites to give a single SAP Landscape. The existence of this project allows the POLFS resilience and DR approach to be validated in Belfast, and this validated solution will be reused for the main Solaris Oracle database server (DAT) which has similar characteristics to the POLFS XI subsystem.

The production infrastructure is being deployed a few months in advance of when it is required. This de-risks a very complex deployment, and also permits a representative "classroom" to be assembled to provide support staff with on-site training run in conjunction with Fujitsu-Siemens Professional Services to introduce them to some of the new concepts.

The Data Base Administration team in Belfast are also running a representative development environment on the pre-production estate to assist the Branch Database designers in producing the Branch Database High Level Design (DES/APP/HLD/0020). This will ensure that the designers follow Infrastructure Services best practices for implementing Oracle DataGuard and Oracle Streams, and that the processes used to support the ~~HNG-X~~HNG-X Solution are familiar to the DBA Team, and similar to processes used on other accounts. Fujitsu Services are taking advantage of their experience in offering similar services (albeit not on such a large scale) to other customers in ensuring that a reliable and effective service is delivered for ~~HNG-X~~HNG-X.

The advance infrastructure will then be available for initial DR Development work, which will mainly consist of ensuring that the basic principles demonstrated in the classroom are turned into rigorous processes and scripts for the actual solution.

Finally a series of failover tests has been built in to the V&I Cycle testing. These tests will be run by ITU and operated by the actual support staff. The number of tests allows all staff to be made familiar with the solution, and for F3 and F4 (as the final two tests are known) to be run in cooperation with the Business Continuity Manager and their counterpart from the Customer and serve as acceptance tests. The V&I HLTP (TST/SOT/HTP/0006) should be referred to for a detailed view of this testing.

These tests will be preceded by a procedural walk-through to familiarise all the Service Delivery Units with their role in the new solution. This is similar to the processes involved in existing business continuity planning and rehearsal.

In addition is expected that V&I will include a programme of non-functional testing designed to test the N+1 resilience features and recovery from backup.

It is also a design goal to make the eventual Test environment representative enough that many of the N+1 resilience features and their operation may be tested in the Test environment in a realistic enough manner to satisfy many of the requirements of Business Continuity Testing without impacting on the Production service.

2.5 Ongoing Business Continuity Testing

A series of walks-through are also performed, especially for those tests where there is considerable inter-disciplinary interaction. These exercises are actually highly valuable in keeping the processes fresh, and also as training exercises prior to a full business continuity test, especially where new staff are being trained. They also provide an opportunity for junior staff members (who may find themselves at the sharp end one night) to manage their team's input in a non-threatening environment.

The Business Continuity Framework (~~FS/BAU/SPP/001~~SVM/SDM/SIP/0001) and Business Continuity Plan (SVM/SDM/PLA/0003) cover normal "Business as usual" testing, and a schedule is published annually to coordinate this.

Inevitably a test of site DR will involve some impact to the normal service. In a Production/Test campus pair this will also involve loss of the Test Service for the duration of the test. This is particularly noticeable for the counters which in Horizon are able to operate autonomously as isolated Riposte nodes (although temporarily without network banking), but at ~~HNG-x~~HNG-x must contact the Branch Database in the ~~data centre~~Data Centre in order to transact any business and will suffer an outage of up to two hours as the failover occurs.

In order to minimise the impact to the business it is recommended that business continuity tests be scheduled to fail over on Saturday at 0200 and fail back on Sunday at 0200, which will cause the Saturday business day and main overnight batch processing to be run from IRE19, and allow a contingency window on Sunday to ensure readiness for Monday morning. Monday morning is when the peak business transaction rate occurs.

3 Core Architectural Elements

3.1 ~~Data Centre~~Data Centres

There are two ~~data centre~~Data Centres, based in Northern Ireland, denoted by the Fujitsu location codes IRE11 and IRE19.

Under normal operating conditions, IRE11 offers the Production ~~HNG-x~~HNG-x service and IRE19 offers the Test service (strictly a number of Test services).

Production components of the overall solution are active at both IRE11 and IRE19 simultaneously in order to facilitate failover, and Test components are located at IRE11 to allow Test systems to fully test changes to the DR solution once in service.

The network is considered active at both ~~data centre~~Data Centres and is managed as a single Production entity. This is required to support the rapid failover from IRE11 to IRE19 in a disaster recovery situation.

The direct distance between IRE11 and IRE19 is 6.3 kilometres (3.9 miles). Via the main roads, the distance is 15.8 kilometres (9.8 miles).

3.2 Intercampus Link

The intercampus link is a high speed fibre link between the primary and secondary ~~data centre~~Data Centre hosting sites, comprising two redundant, diversely routed fibre links which are DWDM multiplexed to form a number of usable logical links. The DWDM end points are separated by at least 5m at each site. There is a detailed description in the SAN High Level Design (DES/NET/HLD/0007).

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

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

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

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

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

3.3 Storage

The Platform & Storage Architecture (ARC/PPS/ARC/0001) Section 2.5.3.1 describes a number of service classes for storage, based on performance and resilience requirements. The Storage High Level Design (DEV/PPS/HLD/0007) provides more detail on the storage implementation.

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

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

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

Storage Service Class 4: Near line file system storage – Asynchronous replication or no replication (RPO > 24hr, RTO > 48hr).

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

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

Storage Service Class 7: Tape for off-site storage

These have been mapped physically onto a number of different systems, which happen to all be provided by EMC except for Class 7, although other vendors provide similar functionality.

EMC were selected in a rigorous bidding process as the best overall supplier. This judged the ability of the vendor to support the service long term as well as meeting the requirement, and was not simply a lowest cost bid, although the alternative bids were used to secure a competitive price from EMC. There are advantages from a support resolution perspective of having a single supplier provide all components in a SAN.

More detailed information can be found on EMC's website for each of these systems. This section is only intended to give a very high-level overview.

DMX (Direct Matrix) Disk Array is the latest generation of the Symmetrix range (Symm7). The Symmetrix provides an extremely robust and resilient storage array, fronted by a large cache for performance, and capable of synchronous or asynchronous hardware replication, known by EMC as Site Remote Data Facility (SRDF), over considerable distances, although large distances do introduce latency on disk writes. The disk array provides RAID-1 and RAID-5 protection internally, and using the TimeFinder product enables clone or split-mirror backups, known by EMC as Business Continuance Volumes (BCV), to be managed by the array, and avoid the need for host CPU cycles to be dedicated to this task. The Symmetrix has an internal battery backup, and even in the event of a complete site power failure (e.g. the fire brigade shut off the power) the data will be de-staged from the cache to disk before automatically shutting down.

Clariion Disk Array (formerly made by Data General) is a slightly less resilient disk array, which is still capable of high performance, but not on the same scale as DMX. Limitations on the number of clones and site replicated volumes that Clariion can manage mean that Clariion cannot carry all Class 2 & 3 storage.

~~Celera~~Celerra Network Addressable Storage (NAS). A small appliance connects to both SAN and LAN and allows filesystems to be presented in a heterogeneous way either to NFS based systems such as linux and Solaris or to CIFS based systems such as Windows 2003. The use of such a facility means that the storage may be presented directly to many systems which need to share data, and the inter-dependence of those systems is broken, allowing lower resilience targets or longer backup windows to be specified.

Centera Content Addressable Storage (CAS). This is used by EMC to provide a low cost and more reliable alternative to tape for audit and archive solutions. The resilience of each array is relatively low, but by having one at each site and allowing the audit application to manage replication a solution with high overall resilience may be built. The IXOS archive solution deployed for POLFS actually allows users to query data in the archive in practically real-time which would be impossible with a tape or optical juke box system.

EDL 4100. Formerly CDL (Clariion Disk Library) but renamed EMC Disk Library as the latest generation allows a DMX to be used for greater scaleability. A small linux server is attached to the disk array, and presents virtual tape drives to the SAN to emulate a variety of physical tape libraries. In the ~~HNG-x~~HNG-x solution we have chosen to emulate a StorageTEK L700 with LTO3 drives, as this both allows a simple migration path from Horizon which used a real StorageTEK L180 with LTO1 drives, and allows test systems to use StorageTEK libraries such as the L40 without modifying the backup application.

TX24 LTO3 Autoloader. A small and very simple tape library with a single LTO3 drive, but capable of writing to as many as 24 tapes without operator intervention. This is provided only to allow very rare physical exports of data from the ~~data centre~~Data Centre. Normally the size of the EDL and Centera mean that all data is stored within the solution, using the pair of sites for resilience.

The primary storage elements in each ~~data centre~~Data Centre consist of:

- Two EMC Symmetrix DMX3
- One EMC Clariion CX3-80
- One EMC ~~Celera~~Celerra NAS
- Two EMC Centera CAS systems (one for Audit, one for POLFS Archive)
- One EMC EDL 4100 Virtual Tape Library
- Two Cisco MDS9509 Directors [configured as independent switched fabrics]
- One Fujitsu-Siemens FibreCat TX24 LTO3 autoloader.

For systems which are hosted on BladeFrame, all data including the boot drive will be hosted on the SAN.

~~There will be an extra CX3-80 at the IRE19 site for the use of Test systems only. This is not replicated to IRE11 and this does not form part of the reduced testing capability available following loss of IRE19.~~

All significant business data from discrete servers is also stored on the SAN (these servers use local boot drives).

3.3.1 Storage Design

The storage design for this project is detailed in DES/PPS/HLD/0007 and ~~HNG-X~~HNG-X SAN High Level Design (DES/NET/HLD/0007). A brief overview is provided here to aid the reader's understanding.

[SHAPE * MERGEFORMAT]

Different VSANs are used to segregate data paths either for performance or security reasons. The primary separations have been labelled here as vSAN A etc.

Storage is synchronously replicated between the primary and secondary sites using ~~SDRF~~ SRDF via the Intercampus link described earlier. Before a write is acknowledged to a server, it is written to the storage at both the primary and secondary sites. This ensures consistency of data between the primary and secondary sites.

The Symmetrix DMX disk arrays are operated as two independent pairs. One will contain the Branch Database, and the other the Branch Standby DataGuard copy. In the unlikely event that a fault develops with the storage continuity of service is ensured, as at least one copy of the Branch Database will be available. Note that this also requires the BAL servers and the BranchDB RAC Node boot disks to be balanced across both DMX's. This will be discussed in more detail in the chapter later for that service.

The CX3-80 Clariion also supports an area of disk for Oracle RMAN backup. This is the primary recovery point for Branch Database corruption, and online storage is provided to allow recovery to be made in as timely a manner as possible.

[DN: This has not been adopted for NPS, so we would lose banking while the problem was resolved if that DMX happened to be the one with a fault.]

3.4 BladeFrame

The Platform & Storage Architecture (ARC/PPS/ARC/0001) and the BladeFrame High Level Design (DES/PPS/HLD/0025) discusses BladeFrame deployment in detail. A summary is provided here.

BladeFrame from Fujitsu-Siemens consists of a chassis with up to 24 stateless processing blades (pBlades), two control blades (cBlade) and two switch blades (sBlade) in a cabinet with a foot-print similar to a normal server cabinet.

- A processor blade or pBlade contains processors and memory. There are 24 per chassis (or frame).
- A switch blade or sBlade manages the internal switching of fibre-channel (storage) and Ethernet packets and provides I/O to the pBlades via the backplane.
- The control blades or cBlades comprise a cluster presenting the PAN Manager service. This provides out-of-band management of the server instances (power on, power off, console), and also controls all the external connections to the network and storage. A single view of resilient network and storage connections is provided to the pBlades.
- From PAN Manager 5.1.3 Xen virtualisation is offered natively by BladeFrame, and a hypervisor may be started on a pBlade to provide a number of virtual blades or vBlades. The amount of memory used by all the vBlades cannot exceed the physical memory in a pBlade, and this controls the amount of virtualisation that may be offered. Provided that the virtual CPU threads do not exceed the number of physical CPU cores little performance degradation is noticed.
- An LPAN is a named set of resources. In the ~~HNG-x~~HNG-x implementation this has been chosen to correspond to the Platform Set (also known as Rig). The disks in the SAN have a static relationship with the LPAN, but other resources such as pBlades may be shared allowing a form of work load management.

The cBlade runs the PAN Manager software. This software allows logical sub-sets of resources (LPAN) to be defined and virtualised server instances known as pServers to be associated with resources. This permits one to specify a set of services in a virtual way, and therefore have an area of storage (which we know how to replicate using hardware) as the DR mechanism.

Since all the pBlades are stateless, all data including boot disks are stored on the EMC arrays. All disks except disks holding swap or temporary information will be replicated to the secondary site.

The PAN Manager service is hosted on the master cBlade, and the cBlades operate as a cluster to provide a resilient PAN Manager on a VIP.

If a pBlade were to fail, the PAN Manager software can automatically start up the pServer on a spare pBlade in the server's primary pool, or on a pBlade in a designated failover pool with minimal loss of service (the time for a reboot). It is up to the application to recover. Although tools do exist to allow application management by PAN Manager these have not been used at ~~HNG-x~~HNG-x.

~~*[DN: This last statement may change as a result of resilience testing, but it is fairly unlikely as SYSMAN is performing this function in the HNG-x context]*~~

During normal use, the Control Blades share external traffic for the pBlades in a balanced state, allowing the BladeFrame to harness the capacity of both simultaneously. If an individual cBlade should fail, the other continues to manage resources for the entire BladeFrame.

A pair Switch Blades (sBlade) are provided to route traffic from the pBlades to the cBlades. At system startup or pBlade insertion each pBlade makes a data connection to both sBlades, which in turn make a connection to both cBlades. Egenera uses the term GigaNet for this service, which operates both LAN and SAN traffic for the pServers on the pBlade. In service these paths are balanced in a least busy fashion. Upon failure of either an sBlade or a cBlade the pServer simply retries and uses one of the remaining paths.

In principle when an sBlade or cBlade has failed the number of data paths, and therefore the throughput, are halved, but in practice it is rare for all pServers to be operating at maximum capacity. A single cBlade (LAN and SAN) can pass 16gbps of traffic which is a huge amount, and other infrastructure constraints are likely to be of more concern.

Due to the exceptional self-managing of this infrastructure component and abstraction of the physical processors from the storage which holds the boot images and data, BladeFrames are exceptionally well-suited for use in a site-failover scenario. It is expected that all services will be hosted on BladeFrame technology except where some feature explicitly inhibits this.

There is an additional benefit that the number of "spare" boxes sitting about goes down from one per service to one per ten services reducing the overall running cost of the solution.

BladeFrames may operate in "farms" of up to three chassis, allowing a single point of control, and also allowing services to fail over between frames. One PAN Manager Service is designated as the Frame Master, and all farm members operate as an extended cluster.

Early experience has shown that Farm Manager failover is not as well behaved as PAN Manager failover, and that if the chassis with the Farm Manager fails control of the farm is lost until it is restored. This may be rectified in newer versions of PAN Manager.

This experience, plus physical limitations on the number of LUNs that may be presented to each PAN Manager (2000) have ruled out the use of farms for ~~HNG-x~~HNG-x. The loss of the Farm functionality is not very important to ~~HNG-x~~HNG-x, but it does introduce a loss of flexibility. Services may still be moved between frames, but this is a manual process that involves storage reconfiguration.

Highly available services (like BranchDB) still require special solutions (like Oracle RAC) which BladeFrame is particularly well suited to host because of the way storage is virtualised. If a cluster is deployed it is important to make sure that the pBlades used for the cluster are in different power domains so that the failure of a power module will only affect one cluster member. This approach is applied more generally to other physical resources, but it must be recognised that there are so many services in an individual frame that complete failure of an entire frame will be likely to trigger DR.

One feature of this virtualisation of storage is that the pBlade is not SAN aware, and certain agents and tools from storage vendors will not be able to resolve or manage storage presented to these servers. This may change in future releases.

The BladeFrame chassis are deployed in pairs, one at each site, with the same physical arrangement of pBlades in corresponding pairs. Whilst this is not strictly necessary it makes assurance of performance upon DR more straightforward. Each member of a pair will have the same LPAN definitions, typically a Production LPAN and one or more Test LPANS. Services within these LPAN may be stopped and started as required provided external resources such as storage and networks are being presented to the LPAN at that site.

The Power Modules each have two 3-phase supplies and supply power to 6 chassis slots for pBlades. They are labelled A B C & D. A also supplies cBlade 1 and switchBlade 1, and B supplies cBlade 2 and switchBlade 2.

If the PIM partially fails (one line in) but it does still work on the second line, we still have to schedule a PIM replacement and that means taking 6 PBLADES and possibly 1 control blade and 1 switch blade out of action while the PIM is replaced.

If the PIM fails causing a total loss of power on a power domain it will result in 6 PBLADE crashes (in a fully populated frame) and single control blade and switch blade failure if it is in the A or B power domains.

While a PIM is off-line there is not enough spare resource to restart all the failed services on the remaining pBlades, and services have been distributed in each BladeFrame and across BladeFrames such that PIM failure causes loss of resilience not loss of service (see DES/GEN/SPE/0007 for detailed layouts).

Loss of an entire chassis is unlikely. Depending on which chassis is lost an assessment would need to be made of whether to initiate DR immediately or whether enough service was being offered to delay the DR until outside the normal operational day.

Formatted: Font color: Auto

Formatted: Font color: Auto

Formatted: Font color: Auto

Formatted: Font color: Auto

Formatted: Font color: Auto

3.5 Network

The network approach to resilience is defined in Network Technical Architecture (ARC/NET/ARC/0001) and developed for each major subsystem in high level designs: the ~~Data Centre~~Data Centre LAN Design (DES/NET/HLD/0004), the Wide Area Network Design (DES/NET/HLD/0009), the Branch Access HLD (DES/NET/HLD/0014), and the Transit LAN Design (DES/NET/HLD/0015) which presents models for connecting to third parties such as the financial institutions.

The network is based on four core Cisco 6513 switches (two per site) and four Access 6513 switches with resilient ASA5540 based firewalls between the Core and Access layers. Each switch is internally highly redundant, and is connected via ISL to its partner. Each core switch will also have a firewall module and an ACE (Application Control Engine) module.

3.5.1 Cisco VRRP

The router uses Cisco VRRP to provide a virtualised routing service. If one of the 6513 switches fails, then the other switch will take over as being the next hop gateway for each VLAN that it supports.

3.5.2 Cisco ACE (Application Control Engine)

Cisco ACE module is a load balancing module which fits within the Cisco 6513 chassis.

The module monitors services on the network that are in an N+1 configuration. If a client uses one of these services then ACE will be used to direct the client to the least loaded instance, or in the case of active/standby systems to the active instance.

Configurable probes allow the ACE module to check an application is available rather than just the server is available – so in addition to simple "ping" monitoring, scripts can be configured to check the health of a particular web application hosted by a server, and remove the server from the load-balancing pool if it returns an incorrect response. The application designer is expected to discuss how ACE is used and the detection strategy in the network section of the application HLD.

Inter-chassis: An ACE in one Cisco Catalyst 6513 is protected by an ACE in a peer Cisco Catalyst 6513.

Examples of systems monitored by ACE:

Branch access layer (Interstage)

Other web services

POLFS Production instance (PLP)

The list is not exhaustive, but there is a service-by-service overview later in the document that will describe whether a system uses ACE, and if so whether it uses ACE load balancing or some internal load balancing.

There are other mechanisms for load balancing, for example using the Oracle TNS*Listener service, and it is up to the application designer to choose the most appropriate technology in terms of performance and resilience. Choosing ACE at least means that the application designer can hand off the operation of resilience to the network but the time taken to detect and advertise may not be suitable in all cases.

3.5.3 Network Management Systems

HP OpenView, Cisco Works and a number of diagnostic probes are located in each ~~data centre~~Data Centre. The network is treated for management purposes as entirely Production, with isolated areas of testing. This is to ensure that the secondary site network is always ready to act as the DR target.

HP OpenView gathers ~~snmp~~SNMP events from network equipment, filters the events, and forwards them to Tivoli (SYSMAN). OpenView also actively probes for managed devices, such as servers, and raises an alert if the server cannot be contacted.

Alarm Point is used to forward pager alerts for critical failures to ensure that a "flashing light" is less likely to be overlooked. This is especially important when a critical alert is received during a period when a number of less severe alerts are being dealt with. Alarm Point will also be used by the SYSMAN systems to raise alerts.

Loss of these systems is analogous to losing the instruments in a car. It does not put the system into immediate jeopardy, but there is a danger that a warning will be missed. If an event has occurred, then the unavailability of these systems means that diagnosis will be slowed down considerably.

The systems, collectively known as NMS, are described in the Network Management System HLD (DES/NET/HLD/0012), and their migration is covered in the migration design (DES/MIG/HLD/0001).

3.5.4 Directory Services

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

Secondary authentication is integrated with AD and managed transparently for applications. This is described in Strong Authentication High Level Design (DES/SEC/HLD/0001).

Although at first sight these are not very critical services, their non-availability may prevent staff from accessing the estate to effect a timely repair. It is therefore essential that under all failure circumstances where it is possible to offer a business service that a working authentication mechanism is available.

AD generally operates as a peer-peer service, but has a number of "Flexible Single Master Operations" or FSMO Roles. These roles are not all offered by the same server, and tend to be spread out amongst the members. The roles are only required to make changes to the domain, but such changes include adding members and password resets, so any aspect of DR that requires such an operation may need to wait for AD FSMO Role transfer to complete, and no critical services should depend on this.

3.5.5 DNS

A primary and secondary DNS service is provided based on dedicated linux servers. DNS is designed for resilience across sites, and therefore the resilience model is also the DR model. The Windows Active Directory domain controller infrastructure is also a secondary DNS server, and provides DNS services to the Windows platforms in the estate. The design of DNS is covered in the Domain Naming System High Level Design (DES/NET/HLD/0006).

As with AD there is a "master" for making updates, and a number of slaves offering the look-up service. Critical services should not rely on~~ DNS~~ updating DNS entries~~es~~ as part of the DR process as the master may not be available at the time of failover.

3.5.6 Time Synchronisation

There will be four dedicated time servers, two at each ~~data centre~~Data Centre. Like AD and DNS, time synchronisation is designed to be inherently resilient and the terms resilience and DR are synonymous. The design of time synchronisation is covered in the Time Synchronisation High Level Design (DES/NET/HLD/0013).

3.6 Oracle

3.6.1 Oracle RAC

The Branch Database and the NPS database both use Oracle RAC to provide high availability.

This configuration provides load balancing of client requests during normal operation. Additionally, if one node in the RAC cluster fails the other nodes will take over the load.

Therefore, the system capacity is managed in an N+1 configuration, such that the RAC cluster can handle peak load even with one failed node.

Branch Database N=3 (four nodes normally active)

NPS N=1 (two nodes normally active).

For Branch DB each branch will normally access the same node, if a node is unavailable then the failed node's branches will be spread across the remaining nodes.

NPS supports several services e.g. ETS, DCS, ~~NWB~~NBS as if they were separate applications. A branch connection will go to the same node for any one of these services, but the branch will not necessarily use the same node for all services.

The mechanism for ensuring persistent connections and load balancing and for managing failed nodes is described in detail in the Branch Database HLD (APP/ARC/HLD/0020), with high level context in the Online Services Architecture (ARC/APP/ARC/0005) and the Branch Database Architecture (ARC/APP/ARC/0008)

3.6.2 DataGuard

Oracle DataGuard is an Oracle feature which allows one or more standby databases to be maintained in a transactionally consistent fashion to a primary database. This is achieved by applying changes from the primary database to a secondary copy of the database.

In the implementation for ~~HNG-X~~HNG-X, the production Branch Database will be the primary database, and a secondary copy of this database will be maintained at the **same site**. The only purpose of this secondary database is to guard against physical corruption of the primary production database, i.e. some form of unrecoverable I/O error resulting in a loss of database integrity.

DataGuard allows for control of lag target in a number of ways, including fully synchronous. Because of the manner in which the changes are propagated as application level applied changes it is very unlikely that a data corruption will be propagated, and in the event of the primary being unavailable due to a corrupt block failover to the standby is extremely fast. The primary can then be repaired.

Whether the database is operated flip-flop or by failing back is a design choice.

3.6.3 Streams

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

In ~~HNG-x~~HNG-x the Branch Database will use Streams to send data to the Branch Support database. This allows the SSC to perform diagnostic queries without connecting to the Branch Database, and also allows batch SLA reporting to avoid the Branch Database.

The consequence for resilience and DR is that a DataGuard failover to the Branch Standby must continue to replicate via Streams to the Branch Support database.

3.7 Site Consistency

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

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

At the time of writing no applications have elected to operate asynchronous storage replication. Replication is either synchronous using the EMC SRDF or MirrorView storage replication mechanisms, or is managed by the application.

3.7.1 Storage

A number of classes of storage are proposed in the Platform & Storage Architecture (ARC/PPS/ARC/0001). Applications which require zero data loss on failover must specify suitable storage, which in this case is EMC DMX with synchronous SRDF enabled. Applications which can recover following a failover where some amount of data has been lost, either because the loss of the data is not significant or because the recovery may be effected from journals or upstream systems, may use a lower storage class.

Storage presentation is very important. LUNs that are replicated between sites should have the same SCSI IDs so that the BladeFrame at the failover site can recognise the storage without need for reconfiguration. Additionally device ID's within the storage array should correspond between sites for replicated devices. There are a number of mechanisms for coping with asymmetry, but they all require considerable development and testing effort, and are confusing for support staff. To achieve the goal of operating an efficient and cost effective service such asymmetries are denigrated.

In addition, to allow the workload of any pBlade to be run from any BladeFrame, each LUN that is required by any pServer will be presented to the cBlades in every BladeFrame. There is a limit to the maximum number of LUNs that may be presented to an individual frame at any one time, and this will be controlled by LUN Masking.

The process of changing LUN Masking is somewhat tedious and error-prone. It will therefore be scripted. As far as possible DR will not require changes to LUN Masking. If this becomes unavoidable it must be recognised that the process of changing the masking is time consuming and subsequent steps in the DR must not be allowed to proceed until the masking has been confirmed as having been changed successfully.

3.7.2 Network

All VLANs that are required by any pServer will be trunked to all of the BladeFrames. Internally the pServer instances connect to vSwitches. The vSwitches will be associated with a~~nd~~ VLAN ID and named l<rig>pbfsNNNN~~vSwitch_XXX~~ where ~~XXX~~ NNNN is ~~a mnemonic for the Security Domain which t~~the VLAN ID and rig is the two letter identifier, e.g. pr for production.~~resides in.~~

Formatted: Font: Italic

The presentation of vSwitches to LPANs prevents test systems and production systems from accidentally sharing the same VLAN, and the naming convention assists operators in identifying production switches.

The association of VLAN ID and Security Domain is enumerated on a rig by rig basis in ~~the Data Centre LAN Design~~HNG-x VLAN Mappings (DEV/INF/ION~~LLD~~/0002~~14~~).

There is a distinction to be made between server instances which are BladeFrame based and have a production/test mode of failover, and those server instances which are active/active or active/standby, which may include both BladeFrame and discrete servers.

Server instances which use the production/test mode of failover will appear on the secondary site with the same IP as they had on the primary site. These servers will use 172.18.0.0 subnets. In order to prevent confusion these subnets will normally be inactive at the secondary site. During DR they will be disabled at the primary site and enabled at the secondary site. To simplify this process Cisco Works scripted tasks will be made available for Network Operations staff to use.

Note that although POLFS internally considers itself Production/Test, from the point of view of service delivery (See Section 10.1.8) it is treated as active/active, as the "Test" systems are used by external customers. It is thus possible for the POLFS Production service to fail over without any impact on the rest of the Data Centre. As now, network load balancing (CSM in Horizon, ACE in HNG-x) will advertise a VIP to the external customer for the POLFS service.

~~*[DN: Will we update POLFS which is discrete and uses production/test]*~~

As far as possible one pair of frames will be dedicated to hosting active/active services such as AD, DNS, and SAS which are required to support the initiation of disaster recovery..

Server instances which use the active/active or active/standby model must be present at both sites simultaneously, and will use 172.16.0.0 subnets for IRE11 and 172.17.0.0 subnets for IRE19. Services on these servers are likely to require ACE to present a service VIP to clients.

A Security Domain which contains server instances of both types will require a minimum of four VLAN IDs, two per site.

4 Dev/Test

In order to provide limited testing capability in the event of a Data Centre DR, it is planned to take copies of the LST test rig (from IRE19) and store these in IRE11. The same DR approach can be used irrespective if IRE11 or IRE19 fails. Apart from ensuring that the LST data is available, no additional design work is needed in advance.

Formatted: Justified, Space After: 6 pt, Adjust space between Latin and Asian text, Adjust space between Asian text and numbers

In the event of a DR event, once the production system has been established, an LST configuration will need to be determined sufficient to test specific fixes; this configuration may be able to run within the spare production capacity. In some cases it may be necessary to reduce production capacity in order to accommodate LST testing, for example outside normal PO hours when online systems can safely be run with reduced resilience.

In order to define limitation of scope two sections from Schedule B3.3 are reproduced here for convenience.

B3.3 Section 1.2.6 The HNG-x Services shall use the facilities of the DR Data Centre to provide a testing environment, which shall be able to concurrently support as a minimum either:

(a) two small testing configuration for performance and volume testing (which may support up to 50 per cent of the business volumes as set out in the CCD entitled "Horizon Capacity Management and Business Volumes (PA/PER/033)) and resilience testing of platforms and applications; and

(b) two small testing configurations for functional testing,

or:

(c) a single testing configuration only to support performance and volume testing (which may support up to 100 per cent. of the business volumes as set out in the CCD entitled "Horizon Capacity Management and Business Volumes (PA/PER/033)) and resilience testing of network components.

Any additional test configurations that are required to support changes to Post Office's business shall be dealt with through Change Control Procedure.

B3.3 Section 3.3.1 (g) In the event that the DR Data Centre needs to be used to run the live service or if the DR Data Centre itself is unavailable, there will be no significant test environment available. In this scenario, limited testing (sufficient to test minor fixes needed to keep the live service operational) will be available at a Fujitsu Services development site. However such testing facilities will not be sufficient to test releases.

In normal use the servers located on the ~~secondary site~~DR Data Centre will be used for software testing. In the event of having to perform site failover, testing would cease and the servers will be re-configured for live running.

In order to allow best use of the physical resources available the Platform & Storage Architecture (DES/PPS/ARC/0001) mandates that all pServer instances will be hosted on vBlades except where some specific requirement, e.g. use of 64 bit operating system, makes this impossible. This allows Test to overload pServers onto pBlades with the same configuration and build as live except for the amount of memory & CPU offered.

Horizon platforms use NT4SP6A which it is not possible to run on many modern systems owing to lack of suitable hardware drivers. NT4SP6A systems will either be run on like-for-like hardware sourced from equipment brokers, or where possible hosted on Microsoft Virtual Server Host in the BladeFrame. As VSH instances already offer virtualisation these will be run on pBlades rather than vBlades.

It is possible to present both Xen and MS-VSH virtualisation external to the BladeFrame for example on FSC Primergy RX300 type servers. Provided no SAN storage is required this is relatively uncomplicated, but these systems are not as straightforward to manage.

~~*If the service at IRE19 is unavailable, then testing will take place at a designated Fujitsu Services Development site. There is no requirement to have testing facilities at IRE11*~~

Formatted: Font: Italic

~~*[DN: The full scope of what is required for test is still being determined.]*~~

The BladeFrames, console management systems, network & SAN switches, routers and storage arrays at the secondary site are all considered live production equipment and any management interfaces are connected to production networks not test ones. This is to ensure that readiness to fail over is never compromised.

~~It is expected that some~~Some components of test systems will exist at IRE11 in order to test cross-campus features, as shown in ARC/SOL/ARC/0001 Section 6.2. In order to reduce the cost of maintaining representative development environments these systems may also be used for production support to develop fixes for example for backup and DR. This will be under the control and management of the LST Test M~~m~~anager.

Test services (known as "rigs") may be shut down when not in use. The only requirement is that they be safe during a period of DR or DR testing, and that they can be restarted when required. This will be managed by controlled shutdown of test services during any DR, managed by the business continuity plan. The start-up is exactly equivalent to a failback of the production service.

In order to meet the requirement of Schedule B3.3 Section 3.3.1 (g) the data that provides the main pre-production rig, known as Live System Test or LST, will be periodically copied to IRE11 storage under the direction of the LST Test Manager. This will use a backup system already being used to preserve "start of cycle" images for the System Test Rig. In the even of running from a single Data Centre there are sufficient pBlades defined in the Production service in DES/GEN/SPE/0007 v4.4 to enable an appropriate LST LPAN to be defined. If necessary testing will be performed at a time when more resources are available, e.g. outside core hours.

A CP has been raised to enable the alternate site for LST Testers to operate concurrently with BRA01. Any failover of LST will be transparent to the site the testers are located at.

In order to minimise on-going support costs the operation of LST in this way is considered an emergency process and will not normally be exercised. Any limitations to operating LST in this way will be overcome operationally in the event of having to invoke the process, but in fact the process is similar to the procedure for moving services between BladeFrames, and the risk is not considered to be great.

Formatted: Font: Not Italic

5 Server classes

There are different classes of server. These are described in the Platforms & Storage Architecture (ARC/PPS/ARC/0001).

- BladeFrame, which is SAN-attached
- Discrete Windows 2003 R2 servers Fujitsu-Siemens Primergy platforms (principally RX300)
- Discrete Linux RHEL 4.0 servers on Fujitsu-Siemens Primergy platforms (principally RX300)
- Discrete Solaris 9 Servers on various Fujitsu-Siemens PrimePower platforms (principally PW450)
- Discrete Solaris 10 Servers on various platforms (principally Fujitsu-Siemens PrimePower PW250, PW650 or Sun SunFire V125)

All of these servers are "~~data centre~~Data Centre class" systems and offer substantial internal resilience such as mirrored boot disks, multiple network interfaces and N+1 power supplies.

Where application resilience requirements permit they may be deployed in a less resilient mode, for example with a single boot disk, but in practice the savings achieved in deployment are usually lost the first time a system rebuild is required as a result of a failure.

5.1.1 Production/Test BladeFrame

There are three production BladeFrames at the primary site running the Production LPAN. This service will only be available at one site at a time. The BladeFrames at the secondary site will be used for Test LPANs during normal operation. In the event of failover the Test LPANs will be shut down and the Production LPAN started at the secondary site.

The PAN Manager service will at all times be in a Production state to facilitate failover, and therefore the cBlade connections will be to the Management VLAN at each site.

Test provisioning systems that need to communicate with the PAN Manager Service will do so through a firewall using an LPAN Administrator role for the test service they are provisioning.

The detailed design of the LPANs is described in the BladeFrame HLD (DES/PPS/HLD/0025).

5.1.2 Active/Active BladeFrame

There is one active / active BladeFrame pair which will host services which are required at all times at both ~~data centre~~Data Centres. These may need to be available to support site failover, or the application model of resilience and DR may be implicitly the same.

There is no equivalent of this in Test, and the test rig design must make due allowance. Many of these services are active/active in order to provide resilience which is not part of the requirement for that test rig, and the extra systems may be omitted unless resilience is a specific part of the test.

Examples of such services include domain controllers and DNS servers, and also elements of the Hydra Branch services.

5.2 Without BladeFrame

The Platforms and Storage Architecture mandates that a platform is only permitted to be outside BladeFrame if:

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

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

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

Note that although NT4 systems are not supported directly by eGenera, the use of virtualisation permits their hosting in the BladeFrame, e.g. on a VSH platform.

6 Availability/Resilience

6.1 Service Level Targets

The migration agreed assumptions and constraints strategy REQ/CUS/STG/0001 defines relaxations to DR requirements during migration.

There are a number of availability requirements presented in Service Level Agreements between the customer and the account. There are also a number of contract schedules. Typically the contract schedules should only reflect the penalties for missing a target and not define targets themselves.

An overall view of the availability requirement is presented in the Systems Quality Architecture (/ARC/PER/ARC/0001).

6.2 Design Criteria

Resilience targets are met by the application design.

This section lists and clarifies some principles when designing for resilience.

6.2.1 Application Design

Host Applications Database Design and Interface Standards (DES/GEN/STD/0001) lays out standards for application designers that direct them towards delivering an application that is inherently recoverable, and that conforms to standard mechanisms for raising alerts to the estate monitoring system.

The effect of component failures, the alerts raised and the recovery mechanism should be detailed in the Service Resilience and Recovery Catalogue (SRRC).

For database applications the effect of each object in a database becoming unreadable and the mechanism for recovery should also be detailed.

The failure of upstream and downstream systems, any alert raised and its effect on service should also be considered in the design and covered in the Application Support Guide. Manual intervention should not normally be required to recover from such failures.

Many services are designed to cope with the loads of a peak transactional day, and the effects of failures at other times may not be as severe. The SRRC should provide guidance on whether graceful degradation in performance or actual loss of service is expected as the result of a failure.

6.2.2 Alerting and event management

~~Need to review DES/APP/HLD/0007 Host Applications Monitoring~~

Component failures will result in events being raised. A single component e.g. a network interface, may result in a number of events being raised independently, e.g. the network switch and the server may both raise an event.

The System and Estate Management: Monitoring Architecture (ARC/SYM/ARC/0003) describes how events are gathered and analysed.

The events raised must be identified in the SRRC to enable the event filtering system design to present a relevant interpretation of the event to the Systems Management Centre, to assign an appropriate priority to dealing with the event, and to direct the appropriate support teams to respond to the event.

Realtime Active Dashboard (RAD) is designed to filter these events and provide a consolidated business view. The model for this is given in ARC/GEN/STD/0001.

6.2.3 Reduce Single Points of Failure (SPOFs)

Single points of failure will be eliminated where possible.

Examples of remaining SPOF and their mitigation include:

- BladeFrame chassis failure – Internally each frame is highly resilient. The chassis is just the steel structure holding the individual components plus the backplane which is just copper. Excluding deliberate physical attack and events which destroy the ~~data centre~~Data Centre these components are unlikely to fail.
- Storage arrays – Internally each array highly resilient. Business critical systems are replicated to at least two physically separate arrays.
- Branch router – this causes the loss of only one branch
- Counter – this causes the loss of only one trading counter
- C&W link to ~~data centre~~Data Centre – network triangulation via the secondary ~~data centre~~Data Centre C&W and inter-site link will ensure the service will continue if the C&W to the primary ~~data centre~~Data Centre fails.
- IRE11-TH1 has a single generator. The likelihood of this failing whilst there is also a power cut is mitigated by regular testing and servicing. Only services which are non-critical or able to fail over to IRE19 independently will be located in TH1 until the site is upgraded by the commissioning of the old TH2 generator.
- IRE19 is only supplied by a single sub-station. This is mitigated by the ability to run for an extended period on generator.

Formatted: Bullets and Numbering

Formatted: Indent: Left: 1.27 cm, Bulleted + Level: 1 + Aligned at: 1.9 cm + Tab after: 2.54 cm + Indent at: 2.54 cm, Tab stops: 1.9 cm, List tab + Not at 2.54 cm

A number of systems are not fully N+1 resilient in the event of loss of the primary ~~data centre~~Data Centre. Examples include HP OpenView where the licensing cost precludes having four systems just to provide full N+1 resilience, and the POLFS Central Instance server which is being migrated "as is" from Bootle & Wigan.

A number of other systems, e.g. the EMC Remote Gateway server and Supplier Access Server do not have local N+1 resilience as there is adequate time to rebuild these systems should they fail.~~.~~

[DN: Should there be a schedule in the SVM/SDM document set that details these items? Basic process is to identify areas and bring to attention of risk manager.]

6.2.4 System redundancy

All systems will be fully redundant internally.

The BladeFrame already has these capabilities.

Discrete servers will have the following:

- Dual HBAs (independent cards, as opposed to dual-port cards), attached to two different Director switches, if SAN attached
- Dual NICs (independent cards, as opposed to dual-port cards)

- Disks setup in a redundant fashion with either RAID-5 or RAID-1
- Redundant power supplies

In Horizon many servers were deployed that did not comply with this approach because the overall solution deployment could cope with the failure of a single server. On balance experience has shown that the support cost is lower if more resilient servers are deployed, especially where this reduces overnight call out.

6.2.5 Cluster

Clusters are suitable for systems that require very high availability. Such systems include the Branch Database and Network Persistent Store. Clusters are typically very complex to manage and generally expensive to implement.

Use of BladeFrame and Oracle RAC has used frameworks that make cluster deployment relatively straightforward, but the resilience of agents and interface layers such as the Banking Authorisation Agents and Branch Access Layer still demands specialist design expertise.

It should be noted that a cluster is simply a shared database, and whilst this may provide for a very high availability in the event of server failure it provides no protection in the event of deliberate or accidental data corruption.

6.2.6 Redundant systems

Externally this is provided by two or more systems which all offer the same service, either as peers or managed by Cisco ACE or similar load balancing.

A similar effect may be achieved by placing servers in the BladeFrame where the failure of the processor or memory results in what appears to be a simple reboot as the server fails over onto a new pBlade.

This is not suitable for protecting against events such as boot disk corruption, but as the pServers all boot from SAN it is possible to provide clone images for quick recovery (less than ten minutes).

6.2.7 Warm standby

Warm standby is suitable for systems requiring failover in between 30 minutes and 2 hours

The standby server will be placed into a standby mode, for example Solaris can stop at run level 2, where it is ready to take over the service, but the application disks are not mounted and the applications are not started. If the primary host were to fail, a script would be run to mount the disk(s), start the application(s) and present any service addresses.

6.3 Server placement within the BladeFrame

Ideally for high availability as many servers as possible within an N+1 configuration should be placed in different BladeFrames to spread the load in case of a BladeFrame failure.

However, the internal network switching of the BladeFrame is much more performant than inter BladeFrame. The Oracle RAC databases for the Branch Database need to have a very low latency interconnect to run with acceptable performance, and this may be achieved by having all the Oracle RAC branch database servers in the same BladeFrame chassis.

Although this appears to be less resilient as the loss of the BladeFrame chassis would affect all instances, the BladeFrame chassis is highly resilient, and in any case the loss of an entire chassis would be likely to trigger DR.

The BladeFrame has four power modules or PIMs labelled A through D. Each of these powers a different set of six pBlade slots, and the A and B PIMs supply cBlade1 & sBlade1 and cBlade2 & sBlade2 respectively. Services which reside within the same frame should consider the effect of PIM failure on the service, for example cluster members should reside in different power domains to avoid widespread disruption of service upon PIM failure.

The failure of LAN and SAN connections is handled automatically by PAN Manager. In the event of multiple failures or failure of a cBlade followed by subsequent failure of a LAN or SAN interface in the remaining cBlade service may degrade to the point where DR is required.

The Platform Hardware Instance List DEV/GEN/SPE/0007 provides a detailed view of the layout and may be used for example to determine the effect of a PIM failure.

6.4 Storage

6.4.1 Storage Arrays

Both the Symmetrix and Clariion storage arrays have a high level of internal resilience.

Storage arrays are provided with power from separate ~~data centre~~Data Centre power supplies, and these are themselves supplied through independent uninterruptible power supplies and separate breaker (fuse) panels. Internally the storage arrays have many features the mean that failure of a single component is unlikely to affect the ability of the array to continue offering a service. Where application resilience demands it redundant data paths are provided to the disk volumes through separate mirrored SAN fabrics.

Within individual disk arrays RAID is used to ensure data integrity in spite of the loss of a disk drive (either a RAID-1 mirror or RAID-5 parity stripe).

The disk arrays also allow point-in-time copies of data to be maintained as snapshots or cloned copies. These enable rapid recovery from corruption, but the management of making the clone copy and also recovering from it is the responsibility of the application designer.

Storage design is covered in detail in DES/PPS/HLD/0007 and DEV/INF/LLD/0004. The actual mapping of storage to platform instances is in DEV/INF/LLD/0043.

The majority of use of clones is through the backup solution as discussed in the Backup & Recovery High Level Design (DES/SYM/HLD/0015).

6.4.2 SAN Fabric

The SAN Fabric is built around two fibre-channel switches (directors) at each ~~data centre~~Data Centre. The two switches are independent, that is, they form two separate SAN fabrics.

This provides at least two forms of redundancy – firstly, the two fabrics allow for failure of any element in any one fabric (assuming both server and storage are connected to both fabrics). Secondly, a (human) error during a configuration change on one fabric will not affect the other independent fabric, typically allowing the change to be corrected before any adverse results are encountered.

Details of the SAN configuration are presented in ~~HNG-X~~HNG-X ~~-~~SAN Design & Patching Schedule Spreadsheet (DEV/INF/ID/0003). SAN High Level Design is in DES/NET/HLD/0007.

6.4.3 Host Systems

All systems that connect to the storage arrays will have two HBAs, each HBA will connect to a different FC director, and therefore to a separate fabric. Multi-pathing is managed either by the control blades in the BladeFrame, or by host based multi-pathing for the relevant platform. This may be Solaris leadville, Symantec (Veritas) Dynamic Multi Pathing (dmp) or EMC PowerPath. The choice for each platform foundation will be detailed in the platform foundation high level designs, Windows 2003 DES/PPS/HLD/0001, Red Hat Enterprise Linux (DES/PPS/HLD/0002) or Solaris 10 (DES/PPS/HLD/0012). There are some historical Solaris 9 platforms. These will use dmp.

If one of the Storage ports, HBAs or FC directors fails or if there is a cabling problem, it will not cause the server to lose its connection to the storage. Connectivity should be designed so that this will also not affect performance.

BladeFrame may either allow many paths to a single device, or the paths may be ~~grouped~~ segregated so that many more devices may be presented. If there are four path~~ groups~~s (the maximum, each with one cBlade HBA) then 1024 devices may be presented, 256 to each path group. Note that if four path groups are defined there is no resilience to dual failure, and this should normally only be configured for services that are not business critical.

The BladeFrame High Level Design (DES/PPS/HLD/0025) will detail the setting up of ~~path groups~~storage connectivity.

Note that the BladeFrame uses a "least busy path" algorithm. Presenting too many paths to a disk will have a detrimental effect on performance. The optimum is 2, 4 or 8 paths.

6.5 Network

All network components are deployed in pairs at each ~~data centre~~Data Centre. The only exception as outlined in Section 6.3 is the C&W link which has a single CE router and single HO router as having pairs would not substantially improve service availability.

Every discrete server that connects to the network will have at least two NICs. Each NIC will connect to a different network switch. The NICs will be configured in an Active/Passive configuration (not load balanced) with Switch 1 as the preferred switch.

The Platform Foundation for each discrete platform type, including the BladeFrame HLD will state how this is implemented for each platform. Applications may have special requirements, and these will be covered in the HLD and highlighted in the PPD for such platforms.

Some appliances may only provide one NIC. Where such appliances are business critical there will be adequate performance available from the remaining services following the failure of a single catalyst switch.

Horizon systems will be migrated as is, and no upgrades will be made. As the majority are being hosted on MS Virtual Server Host in the BladeFrame this is not relevant to most of the migrated Horizon Branch systems (Hydra).

~~*[DN: Implies some sort of coordination between the Platform Foundation Design and the Network Low Level Design and Patching Schedule, which I don't think is clearly expressed in either.]*~~

6.5.1 Basic Topology

This section is only attempting to provide an overview. For further details the Network Technical Architecture (ARC/NET/ARC/0001) should be referred to.

The network is split into major subsystems which are defined in ~~Data Centre~~Data Centre LAN Design (DES/NET/HLD/0004), Wide Area Network Design (DES/NET/HLD/0009), Branch Access HLD (DES/NET/HLD/0014), and Transit LAN Design (DES/NET/HLD/0015) which presents models for connecting to third parties such as the financial institutions.

For details of the Data Centre LAN implementation please see DEV/INF/LLD/0041 which shows the various routing protocols in use and the interfaces to Transit and Branch LANs.

The figure below gives a very high level view of the switch connectivity that forms the basis of providing a resilient ~~data centre~~Data Centre network service with DR capability. The inter-site links are leased dark fibre from Virgin (NTL) with an FTEL service. The links do not share any single point of failure, and there is a minimum 10m component separation at all points along the route.

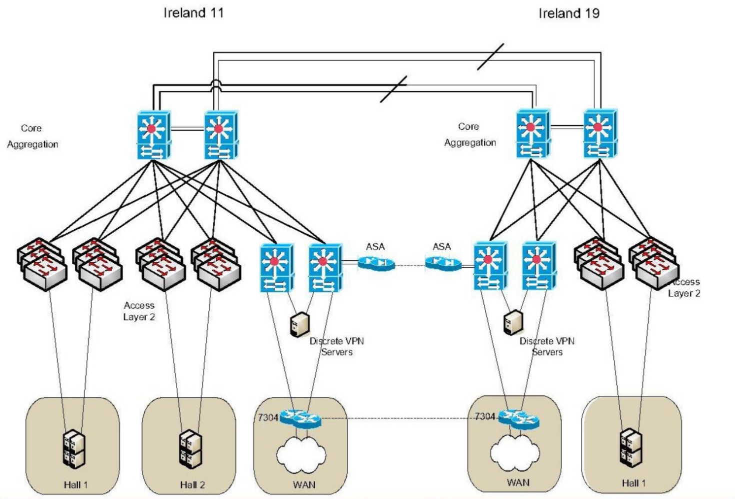
Ireland 11
Ireland 19
Core
Aggregation
Core
Aggregation
ASA
ASA
Access
Layer 2
Access
Layer 2
Discrete VPN
Servers
Discrete VPN
Servers
7304
7304
Hall 1
Hall 2
WAN
WAN
Hall 1

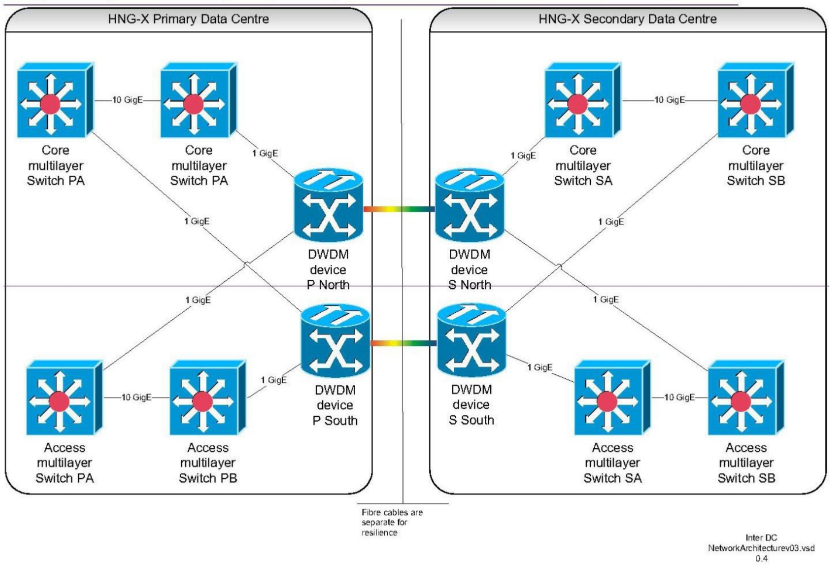
HNG-X Primary Data Centre
HNG-X Secondary Data Centre
Core
multilayer
Switch PA
10 GigE
Core
multilayer
Switch PA
1 GigE
1 GigE
DWDM
device
P North
1 GigE
DWDM
device
S North
Core
multilayer
Switch SA
10 GigE
Core
multilayer
Switch SB
1 GigE
1 GigE
1 GigE
1 GigE
DWDM
device
P South
DWDM
device
S South
1 GigE
Access
multilayer
Switch PA
10 GigE
Access
multilayer
Switch PB
Access
multilayer
Switch SA
10 GigE
Access
multilayer
Switch SB
Fibre cables are
separate for
resilience
Inter DC
NetworkArchitecturev03.vsd
0.4

Connections from outside the ~~data centre~~Data Centre are via dedicated Cable & Wireless 155 mbps circuits, one to each ~~data centre~~Data Centre. Resilience of client connections is provided via the inter-site link.

The segregation is such that a DR Business Continuity Test should only need to consider failover of the core switch layer. The Access layer is covered by normal component resilience testing.

6.5.2 Server Connections

6.5.2.1 BladeFrame

All server instances on BladeFrame reside on the Core switches, and the BladeFrame is not physically connected to the Access switches.

It is necessary to distinguish between the management interface, which is a single connection to each cBlade, and the network ports available for pServer instances. Each BladeFrame will attach on cBlade to the header and one to the footer. For out of band support an Aurora Console Tower system (CON) will be provided that allows secure, managed serial connections to the serial port on the control blade.

Each BladeFrame has two management ports, one on each cBlade, and these are connected to different ~~data centre~~Data Centre switches. The PAN Manager service has a virtual IP which fails over to the master cBlade. The cBlades provide proxy arp resolution for each other, and simple ping tests are an unreliable way of tracing network problems.

Each cBlade has four on-board and four PCI based gigabit NIC for use by the pServers. These are grouped into resilient Ethernet interfaces (rEth) which may also span BladeFrame chassis as mega-rEth (mrEth) if a BladeFarm has been formed to allow a pServer to move between chassis.

Virtual switches are created within the PAN, and these are identified with the NIC that traffic is passing through and the security domain of the VLAN ID, e.g. vSwitch1_DB or vSwitch15_SAS. The LPAN Administrator will ensure that servers are only permitted access to those switches associated with VLANs in which that server resides. The PAN Administrator will ensure that Test VLANs are not visible to Production systems and vice-versa.

6.5.2.2 Discrete Linux

Discrete Linux servers on RX300 will utilise one on-board Intel interface and one Broadcom 57xx NIC interface. Broadcom drivers will offer an active/standby service.

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

For details please refer to the Linux Platform Foundation HLD (DES/PPS/HLD/0002)

For out of band support an iRMC port will be provided, and a Fujitsu-Siemens KS1621 KVM that supports connections over TCP/IP.

6.5.2.3 Discrete Windows

Discrete Windows servers on RX300 will utilise one on-board Intel interface and one Broadcom 57xx NIC interface~~s~~. ~~, and run in~~ Broadcom ~~SLB [Smart Load Balance] mode~~drivers will offer an active/standby service.~~ This is a feature of the Primergy RX300 platform rather than Windows.~~

For details please refer to the Windows 2003 Platform Foundation HLD (DES/PPS/HLD/0001)

For out of band support an iRMC port will be provided, and a Fujitsu-Siemens KS1621 KVM that supports connections over TCP/IP.

6.5.2.4 Solaris

Each Solaris server will have at least dual NICs presented one to each switch. It is usual on larger servers to deploy two quad port ethernet cards (fjqe) so that in the event of port failure a separate port is readily available and the replacement is a simple card swap.

The PrimePower PCI bus only has two slots cable of taking gigabit class cards, and these are usually occupied by the HBAs for the SAN, so PrimePower systems will not normally be connected via gigabit connections. It is however getting common for systems even as small as the PW250 to have two gigabit ethernet connections on the motherboard. The use of these provides a certain level of resilience e.g. against switch or cable failure, but is not fully resilient.

Solaris ipmp will be used to provide network multi-pathing. This presents a base IP address for each card, and a third virtual address which fails over to the active card. Multiple virtual addresses may be overloaded onto a single interface.

Solaris 10 provides Layer 2 detection as well as ping detection which is preferred. In this case the base addresses are not needed. The Solaris Platform Foundation HLD (DES/PPS/HLD/012) defines this as the method to be used for ~~HNG-x~~HNG-x.

Solaris 9 systems (POLFS) that migrate to Belfast will continue to use the three address ipmp method.

For out of band support an Aurora Console Tower system (CON) will be provided that allows secure, managed serial connections.

6.6 Counter Access

Branch Access HLD (DES/NET/HLD/0014).

6.6.1 Counter network access

6.6.1.1 VPN

DEV/INF/LLD/00??

Formatted: Highlight

Horizon counters connect via a single gateway machine in the Branch, and encryption is provided by Utimaco VPN. The longer term goal is to replace this with Branch Router and SSL encryption, but the VPN solution needs to be retained until the counter migration is complete and the operating system has been upgraded to XP.

There are enough VPN servers at each site that a single site has N+1 capability.

~~[DN: The longer term picture is still not clear. We may have VPN for some time]~~

Formatted: Highlight

6.6.1.2 Branch Router

Branch Router HLD (DES/NET/HLD/0010)

The branch router is an ‘off the shelf’ router. Its primary purpose is to provide seamless failover to a backup GPRS network should the standard (usually ADSL) network not be available.

It has no high-availability features built-in. If the router were to fail, an engineer will visit the site and install a new one. This is no different to the Gateway PC failing in Horizon, but it is estimated that the likelihood of appliance failure is lower, and general availability should be improved.

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

6.6.1.2.1 ADSL

The router connects to one of six LNSs

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

6.6.1.2.2 ISDN/PSTN

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

[DN: ISDN was supposed to have been done away with at ~~HNG-x~~HNG-x but in some areas it is still the only feasible mechanism]

6.6.1.2.3 GPRS/3G

Two GGSNs (Gateway GPRS support node) in different locations each connected to different C&W POPs

6.6.1.2.4 VSAT

[DN: Looks like we'll still have them.]

A small number of remote sites use satellite connections. The long term goal is to replace these connections as ADSL coverage improves.

6.6.1.3 Network connection to ~~data centre~~Data Centre

The counter traffic is passed through the C&W cloud and makes a connection to an end-point in the ~~data centre~~Data Centre Branch DMZ. This connection gets authenticated by a RADIUS server. There are RADIUS servers at both ~~data centre~~Data Centres operating in active/active mode. One reason for this is that in the event of a site failure the counters will start polling to reconnect, and if no RADIUS server is present they may continue polling for an extended period, which puts an undue load on the external supplier.

~~*[DN: Need to consider site failure. Do we then need two RADIUS servers to provide N+1]*~~

6.6.1.3.1 Network triangulation

If the C&W link to the primary ~~data centre~~Data Centre fails, then traffic will be re-routed via the secondary ~~data centre~~Data Centre and across the intercampus link.

6.7 Power

The power provided to the ~~data centre~~Data Centre sites can be summarised as follows:

The Primary Site is at IRE11. This has two feeds from independent substations to onsite transformers. The site is capable of running for three days at peak load when on generator power, or longer provided fuel deliveries to site are possible. IRE11 has two computer rooms (or Tech Halls) which are physically separate buildings. TH2 is the site for the majority of the ~~HNG-x~~HNG-x equipment, and is fully N+1 resilient for UPS and generator power.

TH1 is connected via separate paths of single-mode and multi-mode fibre-optic to allow the use of a row of cabinets for RMGA equipment. TH1 has N+1 resilience in UPS, but only a single generator. A second generator has been freed by the upgrade of TH2, and after servicing it is planned to connect this to TH1 to provide full resilience. In practice the double failure of NIE supply and generator is highly unlikely, and there is a programme of generator and UPS testing each year to ensure that the generator is serviceable.

The Secondary Site is at IRE19. This has a feed from a single substation with the transformer offsite. The site is capable of eleven days at peak load when running on generator.

Neither site shares a common substation and Northern Ireland is a resilient part of the National Grid with inter-connectors to Scotland and England.

There is an on-going programme of site improvement which is customer driven.

7 Disaster Recovery

Basic Principles

In ~~HNG-X~~HNG-X the ~~data centre~~Data Centres will be run in a~~n~~ Production/Test mode. All processing normally takes place at the primary ~~data centre~~Data Centre. In the case of a disaster all processing activity will be transferred to the secondary site.

A disaster is an event which renders a ~~data centre~~Data Centre incapable of providing the service.

Business Continuity Planning (~~FS/BAU/SPP/001~~SVM/SDM/SIP/0001) will define what events are deemed to be significant enough to trigger a decision to fail over. Whilst a number of events are clearly in this category, generally this process is non-deterministic and the Service Manager must make a decision and agree it with the Customer before invoking failover.

A regular schedule of testing is agreed each year with the customer both for component resilience and for site failover. As familiarity is gained from operational experience some tests may be scaled back to a procedural walk-through, but there will be at least on full site DR per year. This is driven by the Business Continuity Test Plan (SVM/SDM/PLA/0003).

Failover will cause any testing that is in progress to be halted.

If the primary event has not caused loss of service, then the failover surely will, and the actual start of failover may be delayed until outside core hours to minimise the business impact. For this reason also the normal start time of site failover tests or planned failover for maintenance will be around 0200 on a ~~Friday~~ Saturday with a failback at 0200 on a Sunday.

The secondary ~~data centre~~Data Centre will provide at least as good service as the primary one.

Exceptions:

Failure of the C&W connection into primary ~~data centre~~Data Centre is provided by connection to secondary ~~data centre~~Data Centre and inter-campus routing. In the event of a failure of the link to the secondary ~~data centre~~Data Centre DR is inhibited. In the event of DR the C&W connection is not N+1

Failure of connections to 3rd party services ~~is affected similarly as these~~that rely on C&W is affected similarly. Since no counters would be able to transact this is moot, but may trigger separate penalties.

[DN: Internet connectivity is the subject of CR01517. It may rely on similar triangulation to the C&W link]

Formatted: Font: Italic, Highlight
Formatted: Font: Italic, Highlight
Formatted: Font: Italic

Power supply following DR is not fully N+1.

7.1 Service Level Targets

The requirements for ~~Data Centre~~Data Centre Service Availability are given in the ~~Data Centre~~Data Centre Operations Service: Service Description (SVM/SDM/SD/0003).

The timings for failover design are assumed to run from the time that disaster recovery is authorised.

No time has been allowed for decision making, but until the design has enough detail to be able to determine how close it is to the 2 hour target for having branches operational it is not possible to quantify this as a risk. As previously noted, a decision to fail over may be deliberately delayed if the service is degraded rather than fully unavailable.

For the purposes of providing an end-point which has meaning to most people in all areas of the solution (architects, designers, operations, service managers, customer) a working definition has been used of the moment that the first counter is able to receive a network banking authorisation message. This is not the same as the strict service measure, and a means will need to be developed of assessing the actual service outage.

7.2 Disaster Recovery Process Overview

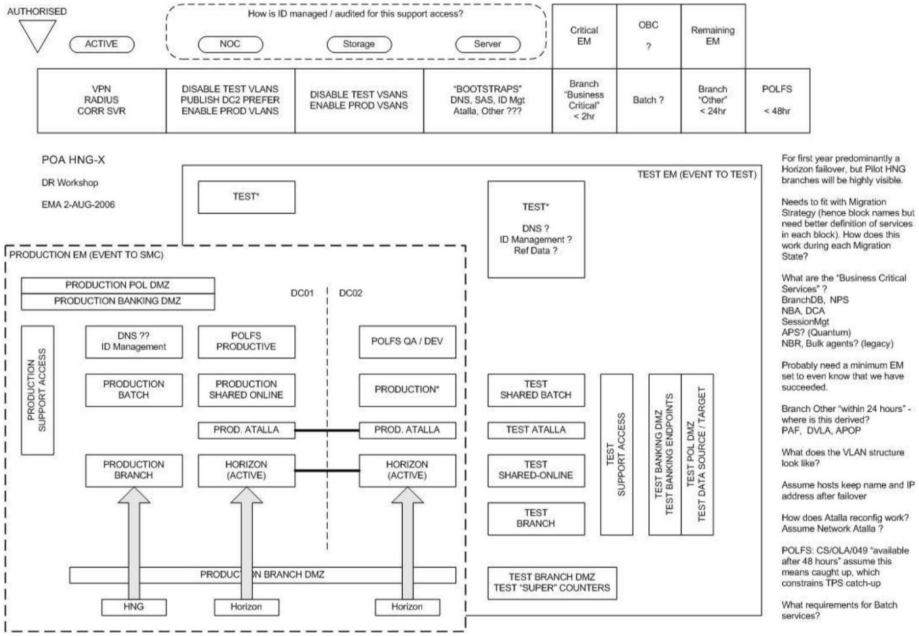

For first year predominantly a Horizon failover, but Pilot HNG branches will be highly visible.

Needs to fit with Migration Strategy (hence block names but need better definition of services in each block). How does this work during each Migration State?

What are the "Business Critical Services" ?
BranchDB, NPS
NBA, DCA
SessionMgt
APS? (Quantum)
NBR, Bulk agents? (legacy)

Probably need a minimum EM set to even know that we have succeeded.

Branch Other "within 24 hours" - where is this derived?
PAF, DVLA, APOP

What does the VLAN structure look like?

Assume hosts keep name and IP address after failover

How does Atalla reconfig work?
Assume Network Atalla ?

POLFS: CS/OLA/049 "available after 48 hours" assume this means caught up, which constrains TPS catch-up

What requirements for Batch services?

In order to allow business critical services to be given priority the conceptual overview presents a similar view to the migration design, and groups services into a small number of logical containers. These will be enumerated later. All that is important in this overview is that some services operate Active/Active, and some services have a permanent presence at one or other ~~data centre~~Data Centre, e.g. PRODUCTION*

It is impractical to make the ~~data centre~~Data Centre pair separate completely into Production and Test and still achieve such a high availability, and the design represents a compromise that minimises the amount of services that must be "warm". The design goal for eliminating these is that their equivalent must then also be present in TEST*, and this represents extra equipment that must be purchased, maintained and supported.

Systems required to permit or facilitate support staff access during DR must be available at the secondary site. These include the SAS terminal servers, AD servers, and systems to facilitate network and storage management. These will all be included in the PRODUCTION* set of systems, and will be available without interruption following failure of the primary site.

Until the AD FSMO Roles are failed over passwords cannot be changed, systems cannot change domains or new domains be added, and users cannot be added. Services which need to become available quickly should not rely on any of these features.

POLFS is effectively still deployed as an independent system capable of independent failover.

~~[DN: Not clear whether CS/OLA/049 is being changed as a result of discoveries during migration planning to provide higher availability for POLFS, but this has no practical impact on HNG-x service to the counters.]~~ It is possible for two Normal states to exist, one with POLFS Productive at IRE11, and the other with it at IRE19.

Formatted: Font: Not Italic

It is also possible that the failure at the primary site could be such that there is no loss of site, but rather loss of a major infrastructure component like an EMC array or a BladeFrame. In this case it is likely that the WAN triangulation would be left operational and POLFS Productive would remain in IRE11.

Formatted: Font: Not Italic

In the event of loss of IRE19 there will be a loss of the testing service, including POLFS QA. The service is expected to continue without interruption, although there will be a loss of resilience.

7.3 Disaster Recovery procedure

7.3.1 Site Failover

Site failover is covered by a Major Incident Process. Much of this is unchanged from Horizon, and the aim of this section is to give an overview of the process and identify new processes which need to be developed.

Some of these steps will be capable of being carried out in parallel. Major checkpoints will be included in the business continuity plan to allow coordination of steps which depend on each other.

Testers will be given as much notice as possible to stop their testing and shut down their systems cleanly.

A message will be put on the Help Desk phone to inform Post Masters and Mistresses of a major problem

Authorisation will be sought for failover

Availability and operation of support services will be confirmed.

If possible production servers and LPANs at the primary site will be shut down cleanly, followed by Test LPANs at the secondary site.

The test network will be disabled and the production network prepared for operation from the secondary site

The storage will be failed over.

The Production LPANs will be started at the secondary site

Business critical services will be started

Other services will be started

States of sites:

- Normal Running:
  - Primary – Production
  - Secondary – Test/Production*
- After Disaster:
  - Primary – non-functional / unavailable
  - Secondary – Test/Production*
- After Failover:
  - Primary – Unavailable
  - Secondary – Production / Limited Test
- After Failover and after primary site fault is fixed:
  - Primary – Standby
  - Secondary – Production / Limited Test

7.3.1.1 Authorisation and notification

Once Fujitsu has determined that site fail-over is necessary, the Customer will be informed through the normal mechanism for a major incident. Once approval to fail over has been received the process is irreversible. The SMC will manage the interaction of the various support teams, and keep Service Managers informed.

Note that there is a new requirement for SMC over Horizon, which is to inform the Test community if events are observed that indicate DR is likely to be invoked.

7.3.1.2 Network

It is important that the counters do not get directed to the primary site once failover has started. Although generally we will have caused Branch Database to be unavailable there are still conditions where the inter-site link ~~–~~may have become partitioned and we need to inhibit transactions to the site we are about to abandon.

The VIPs for all counter services will be advertised from the secondary ~~data centre~~Data Centre to the C&W with a lower cost.

[DN: Not sure that this is a guarantee. How do we guarantee it? Does it matter?]

Test~~ing~~ access to the secondary site ~~will~~ need not be disabled as the authentication domains are separated and testers will only be able to authenticate if their service is explicitly started.

[DN: Need to see final Network design]

Formatted: Font: Italic, Highlight

MSFC's supporting the extended VLANs ~~[DN: We may not need to worry about this. They have nothing to connect to anyway]~~

Formatted: Font: Not Italic

~~Subnets~~ used by BladeFrame hosted services will be ~~disabled at the primary site and enabled at~~failed over to the secondary site.

7.3.1.3 Storage

The Solaris backup servers will operate as Storage Management systems to enable failover of storage to be managed through scripts which will be called out by a master checklist.

All BladeFrame boot LUNs are synchronously replicated from the primary site to the secondary site.

All storage that requires zero data loss on failover is synchronously replicated from the primary site to the secondary site.

Replicated LUNs are normally Read/Write on the Primary site and Read-Only on the secondary site.

Storage is managed by permitting an LPAN to use certain disks. The Production LPAN can only use disks which have been allocated for Production, and each test rig can only use disks which have been allocated for that rig.

Discrete servers will not be moved from Test to Production without a complete rebuild, which will include changing the VLAN they are connected to and the storage that is presented to them. This storage is typically managed by LUN Masking, an operation which can only be performed from the Storage Management Server.

The storage "failover" command will be issued which will write disable storage at the primary site and write enable storage at the secondary site. This also inhibits replication until it is deliberately reinvoked.

At some point after failover an assessment will be made as to when and whether failback is possible.

At this point the direction of replication will be reversed to replicate changes made at the secondary site (now in a Production state) back to the primary.

In Horizon the only reverse replication mode available meant that full resilience was not achieved until after failback, but at ~~HNG-x~~HNG-x it is in principle possible to achieve a state where the Production service running at the secondary site is able to treat the primary site as a DR target. This allows a more extended period of failover running, for example if it is desired to perform failback only on a Sunday.

[DN: The risks of such an approach are difficult to quantify until more LLD work on LAN and SAN is complete. In any case this incurs a higher outage for Test systems which may have programme commitments for fixes or enhancements.]

It is also possible to put the storage into a split state in a controlled manner to maintain an image of the data at the secondary site, for example during migrations where a regression image may be required.

Scripts will be provided by the storage vendor to provide a toolkit to simplify:

Failover

Reverse replication

Failback

Campus split

Campus establish

Campus recover

These will be documented in DEV/PPS/LLD/00??

7.3.1.4 BladeFrame - Active

No action is required. It may be desirable to inhibit or shut down services at the primary site, or prevent traffic being routed to them.

7.3.1.5 BladeFrame - Production/Test

Testing on the secondary site will be stopped and all Test LPANs will be shut down. There is a lower risk of collateral damage to the test system if it is shut down cleanly, and the Test Manager will be consulted as soon as a major problem is identified to initiate shut down.

If access to the primary BladeFrame is available on the primary site, all Production LPANs running on the BladeFrame will be shut down. This enables storage to be in a "clean" state when failed over, which as well as being lower risk is also quicker as the filesystems do not need to be checked.

The Production LPAN will be prepared on the BladeFrames at the secondary site. This will involve reassigning pServers and starting Xen hypervisors using a script on the control blade.

The configuration of the BladeFarm is stored in a XML file. BladeFrame DR allows this to be saved to a SAN disk by an internal scheduler. In the event of difficulties the configuration of the primary LPAN may be recovered from the backup copy.

When storage failover has occurred the Production LPANs will be available to start.

Some resource configuration may be necessary. Such differences will be minimised and scripts will be developed for use during failover to speed up the process and minimise human error.

The individual servers in the Production LPAN will be brought up in a controlled manner in an order that achieves a minimal outage of the Branch service. This is enumerated in Section 7.4. In order to minimise the overall outage smaller, stateless servers may be set to start automatically with the LPAN.

As far as possible built-in LPAN start up ordering will be used, but it may be necessary to coordinate between frames, or to have a finer granularity. Unfortunately it is too late to put a requirement into HADDIS that services should start up and then poll, although many of them do behave in this way.

7.3.1.6 Legacy Batch Server

At the secondary site the standby legacy batch servers will not be used for testing. They are only in place for site failover purposes.

The server that provides N+1 resilience at the secondary site may be made available for volumetric testing. If it has been it will be rebuilt to provide Production N+1 resilience.

[DN: Actually doing this rebuild regularly as part of DR testing is no bad thing]

The standby server will be running [HYPERLINK "http://www.google.co.uk/search?hl=en&client=firefox-a&rls=org.mozilla:en-GB:official&hs=ZkB&sa=X&oi=spell&resnum=0&ct=result&cd=1&q=permanently&spell=1"] but will at a lower than normal 'run level'. When storage failover is completed the standby will be raised to the active run level. This will automatically mount the disks used by the applications, and start services as required, e.g. Oracle databases and the Oracle TNS*Listener. ACE will detect the TNS*Listener and advertise the VIP for the failed over service.

Note that this is substantially unchanged from Horizon and is similar for POLFS systems.

7.3.1.7 Discrete Servers

~~There will be few discreet servers apart from the legacy batch server and the POLFS system.~~

~~There are no other discrete servers which require failover.~~The remaining discrete servers fall into a number of patterns.

Boot from SAN (ECC)

Active at both sites with either able to offer a service (NMN, DXI)

"Like DAT"

Formatted: Indent: First line: 1.27 cm

These services will be dealt with in detail in the service summary. The most significant are actually the Hydra and SYSMAN2 platforms, and these are typically "Like DAT".

7.3.1.7.1 NetBackup Master Catalogue Service

The NetBackup realm consists of a number of media servers (servers attached to tape drives). In order to allow any of these to be used to restore any backup NetBackup provide a centralised Master Catalogue Service.

~~This is in effect a small database, and it is envisaged that it will reside on a replicated SAN disk, and that the server hosting the Master Catalogue Service will operate a similar type of run-level controlled active/standby state, where standby allows it to operate as an ordinary media server.~~

~~*[DN: I am raising a CP to try and get BSM in the BladeFrame to simplify failover]*~~This will be hosted on the BSM platform which is BladeFrame based.

7.3.2 Last Resort – Recovery from Backup

The backup images are provided in order to protect against application corruption. They have not been designed to provide a failover capability, but the very fact that an application image is available means that there is a further mechanism for restoring service.

If this method is invoked it is almost inevitable that some data loss may occur, so systems which require DR with zero data loss may not rely on a backup image.

It is also inevitable that if recovery from backup is invoked that the service outage will considerably exceed that allowed.

If this recovery is invoked it will be done so on the understanding that restoring any service is better than continuing with a service outage, but the ramifications particularly for the application support team in assessing the impact on audit and reconciliation should not be under-estimated.

7.3.3 Site failback

Failback has been alluded to earlier in 7.3.1.3 when discussing reverse synchronisation of storage. This is an essential precursor step, and if the outage has been lengthy this step in itself may take some time.

Following a failover, or more likely following a business continuity test, a point will be reached when the business is ready to failback. It is not possible to be prescriptive, as certain events like flooding can require considerable work to make the primary site available again, and the degree to which the infrastructure may be tested before failback may be limited.

It is likely that active/active services can be restored in advance. These are not business critical, but are essential to providing a secure managed service.

It is possible to state the following:

- Failback is a planned event.
- Failback will cause a service outage.
- Failback may have to be abandoned.

The high level process is very similar to failover, except that the Production services are guaranteed to be running:

- Ensure reverse replication has completed
- Confirm availability of support services at both sites
- Automatic message on help desk and maybe memo earlier during the day
- Shut down Production services at secondary site
- Restore Production Network to normal state
- Failback storage (write enable at primary site)
- Start production services at primary site
- Confirm services available
- Permit customer connections (point of no return)
- Enable Test network
- Start Test services

After the point of no return production services would be expected to remain at the primary site. Abandoning the failback will only occur prior to this point.

[DN: How do we confirm service is available without allowing an actual transaction?]

7.4 Service Start-up Ordering

[DN: Add an intro section to reiterate the SLA requirement for DR]

[DN: The list of services is not yet complete. As part of the Integration & Build process service dependencies are being established. This section will be updated in the next version if required.]

WAN, LAN, SAN (NMN, ECC, RSG)

PAN (CON)

AD, DNS, SSN

Above this line is very difficult to get going. It represents a minimal "core" service.

Formatted: Underline

EST, RADIUS, VPN

Key Management Service

BRDB [RIPOSTE]

NPS, NBAuth, Atalla, [NRA]

BAL

Other Auth Agents: DCS, ETU

Client connectivity: DCM, FTMS, C:D

Web Services: DVLA, PAF, APOP, Track&Trace, Kahala, Telecom, Online Training, Helpdesk

APOP, LFS, TES, DRS [AGE] - Note that TES Harvest from NPS can be disruptive in catch-up

SYSMAN3 / SYSMAN2~~[DN: The list of services is not yet complete. As part of the Integration & Build process service dependencies are being established. This section will be updated in the next version if required.]~~

~~[DN: This strictly lists platform types rather than services]~~ & EACRR - What is EACRR managing now?

TWS / MSH

RDMC, RDDS, TPS, APS, DWH

Backup

Need to get the top-level view of services from the customer perspective so as to say what we now are offering at each stage.

7.4.1 Supporting Services

Some systems need to be active at both sites because they are needed on the secondary site to support site failover, therefore the following will be running on at least one server at each site (now probably in BladeFrame):

SSN/SAS SAS are active/active and should already be available

DNS DNS will be operational as it is active/active with a master/slave relationship. If changes are required then the slave will need to be promoted.

ACD Active directory / Single Sign on / ID Management is active/active and should already be available. FSMO Roles need to be started for full functionality.

DOM Hydra domain controllers

RAU RADIUS Accounting authenticates connection to network devices

NMN Network Management

VPN VPN will exist beyond Hydra as the NT4 counter requires the protection of the Utimaco layer. The VPN systems will remain as in Horizon and will be hosted as guests in discrete Microsoft Virtual Server Hosts (VSD). Supporting services such as the policy manager (VPM) exception server (VEX) and loopback workstation (VDW) will be similarly hosted. VPN is logically located in the triangulated Branch LAN.

RADIUS There are several types of RADIUS server, but these are all active/active and should already be available. This is necessary to prevent counters continuously polling for a connection which appears like a denial of service attack on the branch network provider.

The Radiator product is able to cache authentication information and does not require EST to be deployed except for new connections.

7.4.2 Branch Critical Services (~~HNG-x~~HNG-x)

What do we need running to offer a banking service? What do we need running to prove we are running a banking service?

KMN	Key Management Server
BAL	Does BAL come up expecting BDB or does it poll?
BMX	May be required to manage BAL into service
~~VPN~~	~~Hydra system. active/active.~~
~~RADIUS~~	~~There are several types of RADIUS server, but these are all active/active and should already be available.~~
BranchDB	Oracle RAC in BladeFrame
NPS	Network Persistent Store. Oracle RAC in BladeFrame
NAA	A&L Auth Agent
NAL	Link Auth Agent
NAC	CAPO Auth agent (x2)
HSM	active/active networked Atalla key generator
BranchStandby	This is not strictly critical, but unwise to offer a service without it operational.

7.4.3 Branch Critical Services (Hydra)

These services are mainly active/active and no action should be required, but SYSMAN2 failover and Maestro Scheduler failover may require some sort of linking in activity.

OMD, TMR (Manage EACRR)

KMS

7.4.3.1 Active/Active

Formatted: Heading 4

Formatted: Bullets and Numbering

COR

AGE

NRA

~~KMS will still exist. It may be one of the external systems because of the hardware random number generator, but a CP is being suggested to replace with a software RNG.~~

~~KMS~~

~~ACF~~

~~OCM~~

~~OMD~~

~~SDS~~

7.4.4 Batch Services

DRS and TES are required to allow POL to have a view of transactions. TES should be no more than 15 minutes behind NPS Journals. DRS receives C12 messages which SSC use to alert on transaction codes for detecting failures.

[DN: Will the informal SSC detector be made formal in ~~HNG-x~~HNG-x?]

APS dealt with Quantum Emergency payments, but these should be defunct at ~~HNG-x~~HNG-x~~ Agreed – they've gone.~~

LFS deals with pouch deliveries and collections as well as planned orders.~~ Can certainly live without them.~~

RDMC may be passing Bureau de Change changes or memos Bureau changes are typically only once a day, but mem~~e~~os may be important if we want to communicate with Postmasters and occasionally there is an urgent Bureau update.

TPS, DW and RDDS are pretty much strictly batch.

Note that Batch systems have their own OLA and SLA, and it is possible that outside core hours this may take priority.

~~Will~~ SYSMAN2 will still operate primarily from IRE19~~ (it currently is active in Wigan with Bootle as the primary site for all other services)?~~

MSH is the Hydra maestro scheduler.

DAT ~~This~~ is a single host with a number of databases:

TPS

APS

LFS

TES

DRS

DW

RDMC

RDDS

Formatted: Indent: First line: 1.27 cm

Generally once the databases and Oracle listener are up support staff consider the job as complete, but there are then a number of external systems whose connectivity needs to be checked, especially the FTMS gateways and the TWS scheduling system.

Some of these databases offer a pseudo on-line service, such as LFS pouch tracking, and RDMC Bureau de Change rate changes, or support direct on-line services such as TES Query and DRS "F99" by BSU.

Add something to distinguish between systems which should get noticed automatically because they are always monitored, and systems which only run periodically and need a poke in the ear (although actually this means they are not being properly monitored).

7.4.5 Other Services

Branch_Support (STREAMS) - any SLA on reporting?

NetBackup Master Catalogue Service

Web Services: DEA/DCS, ETU, PAF, DVLA, MoneyGram, BWS, HWS, OWS

FTMS: EDG, TIP

CDG C:D Connect Direct Gateway for delivery of banking files.

DCM Interface to Streamline

~~??DCS;~~ APOP

Track & Trace

ACF - Hydra

Formatted: Portuguese (Brazil)

OCM - Hydra

DEL - Hydra

~~; MoneyGram~~

7.4.6 POLFS

POLFS operates with a 48 hour SLA on DR service outage (CS/OLA/049), primarily because of the time required to catch up on batch processing.

There is a DR facility for the Production service. This is provided by Fujitsu, and is designed such that it does not require the direct involvement of PRISM, but would benefit from their support In the event that the Production service cannot be operated from the platform within the Bootle Data Centre, the Test/QA service in the Wigan Data Centre will be closed down and the Production service built on the Test/QA Platform.

The Service Target is to have the system operational within 48 hours. POLFS is capable of failing over independently of a campus failover as it would be undesirable to invoke campus failover because of a failed SAP system, and because the Production/Test POLFS system is in effect a complete, independent service to Post Office and the Test system is not used by Fujitsu Services.

The DR facility is invoked by either Fujitsu or PRISM logging a Helpdesk call. Fujitsu and Post Office will then assess the request and obtain appropriate management approval. The service availability target states that no single outage should exceed ten hours. Thus this is a key criterion in assessing whether DR should be invoked.

There is an exercise going on as part of migration planning to determine just what business impact such an outage incurs, as many of the business fallback (manual) processes developed for the predecessor to POLFS will have changed radically.

There is a hand-off which is outside the control of Fujitsu Services, as once the SAP service is confirmed as available it is Prism support who manage the catch-up of the batch processing. The Service Target has not been well written in this respect.

~~I~~Prism Alliance has a 24 hour SAP outage every quarter which affects POL FS. Currently Fujitsu Services do not align to this; however it is likely to be considered as an option in the future.

In principle the initial database failover of POLFS is straightforward, and it is only the fact that this is generating activity when many more critical systems are being failed over that makes it desirable to put it in the low priority queue.

~~There is also a hand-off which is outside the control of Fujitsu Services, as once the SAP service is confirmed as available it is Prism support who manage the catch-up of the batch processing.~~

~~*[DN: Confirm whether the 48 hours is for catch-up or just for the Fujitsu Services component]*~~

8 Monitoring / EACRR Replacement

The overview and context is provided by ARC/SYM/ARC/0001.

ARC/SYM/ARC/0003 is the topic architecture for System and Estate Management - Monitoring.

Some terms appear in this chapter which are not in the Abbreviation or Glossary section of this document. This is because Chapter 10 is in effect a glossary of the platform types enumerated in DES/APP/HLD/0009.

8.1 Event gathering and collation

The framework can be described in several perspectives:

- Active monitoring – where agents are looking for known stimuli (for example service down, processor utilisation exceeded).
- Passive monitoring - where events are being raised at source, and can be classified, aggregated etc and forwarded to a collection layer for subsequent aggregation, display in a event viewer and possible forwarding to a business monitor service.
- Business service monitoring – which provides an aggregation of underlying events in business service terms

For each perspective we can describe the product architecture that delivers it as follows.

The Active monitoring perspective is provided by the ITM (IBM Tivoli Monitoring) product

The following is illustrative of the active monitoring infrastructure:

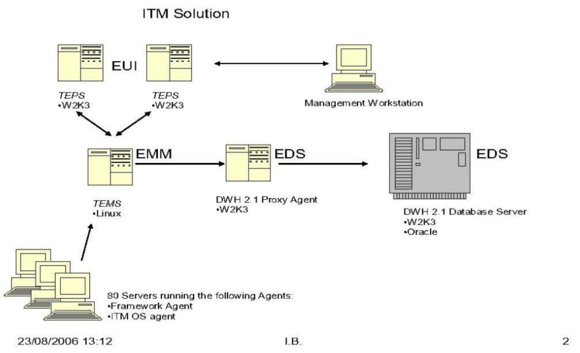

The products for the other perspectives are:

- Passive monitoring - IBM Ominibus products
- Business Service monitoring – IBM Netcool RAD.

The following is illustrative of the passive monitoring integration with business service monitoring.

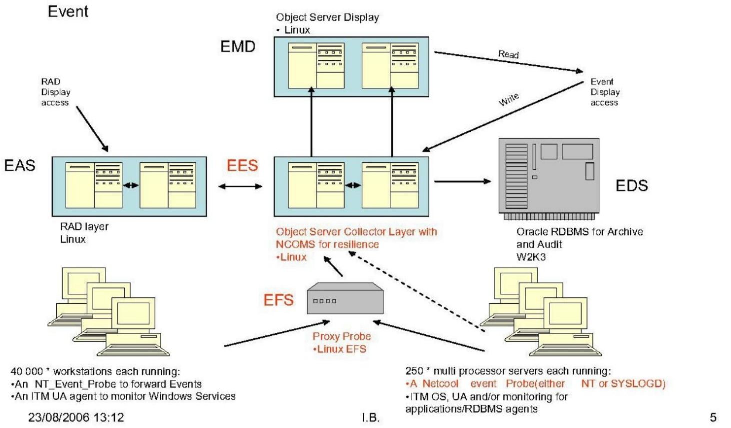
Event
Object Server Display
• Linux
EMD
Read
Event
Display
access
Write
RAD
Display
access
EAS
EES
EDS
RAD layer
Linux
Object Server Collector Layer with
NCOMS for resilience
•Linux
Oracle RDBMS for Archive
and Audit
W2K3
EFS
Proxy Probe
•Linux EFS
40 000 * workstations each running:
•An NT_Event_Probe to forward Events
•An ITM UA agent to monitor Windows Services
250 * multi processor servers each running:
•A Netcool event Probe(either NT or SYSLOGD)
•ITM OS, UA and/or monitoring for
applications/RDBMS agents
23/08/2006 13:12
I.B.
5

Links will be made from the EMD to the KEL and incident management system so that event displays can be enriched with KEL data and incident raised.

Events have business , operational and security significance . The monitoring solution will be configured and deployed to meet these needs . In some cases - such as security event monitoring - it will be augmented by specific tools “

The products can be deployed to construct a tiered solution managing a variety of domains including:

- Platform hardware –
- Operating system – Windows and Red Hat
- Bespoke Applications – both explicit application events and heuristic conclusions from transaction throughput
- Oracle –
- HP Openview – and ~~SNMP~~SNMP domains
- Storage solution –
- Middleware – Interstage
- High level batch scheduling
- SAP

Each domain and its managed objects are appropriately configured and introduced into the monitoring framework to populate event displays and higher level business service monitors, whose status is informed by the individual status of managed objects contributing to the service availability.

8.2 Service Control

[DN: Need to describe the various options, and how they tie in to the applications. A number of models exist in Horizon; should we be standardising on a few patterns; more likely to summarise the models in Horizon and make them available for patterns. A brief summary is presented here but it is not intended to be exhaustive and needs more architectural steer]

The focus in Horizon was on the technology being used to raise the event. The design goal in ~~HNG-x~~HNG-x is to focus on the managed object, and design the event filtering to ensure that SMC are presented with the business impact rather than the underlying cloud of events

8.2.1 Things not currently viewed as services

8.2.1.1 VIP / CSM

Services are monitored by the Cisco ACE which publishes VIPs and directs requests.

8.2.1.2 Interstage

Interstage will be monitored by JMX?, which will report into the System Management systems. Two main systems on the BAL, the Interstage services and the Java on-line routing agents. If Interstage were to fail, either completely or partly then the whole server should be restarted.

8.2.1.3 Systems

A variety of methods are used in Horizon

BMC Patrol (Oracle and Solaris and SAP)

Compaq Insight Manager

Fujitsu-Siemens ServerView

EMC Enterprise Control Centre, network device, and clusterware raising ~~snmp~~SNMP traps

Unix and Cisco syslogs

8.2.1.4 Applications

Applications raise events in a number of ways;

Directly to windows event viewer or unix syslog service

Indirectly to exception tables in the application database

Through an event to the Tivoli TEC or expedited TEC

POLFS (SAP) events are raised via BMC Patrol which is an SAP approved mechanism.

8.2.2 Services which are completely autonomous

There should be none i~~o~~f these. Every service should at least raise events that let SMC report on the state of the system.

In Horizon there were a number of such services, for example the BTI paging system which was offered as a print service on unix systems, and which sent pager alerts direct to the Unix Support Team. This meant that migrating components to SMC management was complex, and was a hang-out from early Horizon days when Tivoli monitoring of the ~~data centre~~Data Centre was not mature.

8.2.3 Services which are autonomous but report events

In unix these are "respawn" services managed by init.d, in Windows these are referred to as "locally Tivoli managed".

Examples are the Network Banking Authorisation agents. They may occasionally self-terminate, but an event is generated to explain why, and they then auto-restart ~~(in unix these are "respawn" services managed by init.d, in Windows these are referred to as "locally Tivoli managed")~~ and start offering the service automatically.

8.3 Interstage probably fits in here rather than above.

8.2.4 Services which are managed externally

There are several sub-classes of these.

The EACRR Horizon system is Tivoli managing the failover of a pool of agent services on a set of physical agent platforms.

The Maestro and TWS schedulers also perform this function. These are not generally event driven, but a class of scheduled jobs has developed (such as the rates file arriving) which trigger scheduled jobs.

There are also still some local unix cron or windows scheduled tasks. These should not be permitted, but certain systems e.g. cBlade PAN Manager do not allow agents (Tivoli or TWS) to be installed.

Tivoli can interact directly with systems through ssh. For example commands may be issued to PAN Manager as an LPAN Administrator.

9 Migration and Transition

9.1 Introduction

The migration strategy in ARC/MIG/STG/0001 foresees a protracted period where old and new ~~data centre~~Data Centres are running in parallel as sets of services are migrated on successive week ends. This is followed by a period of operation in the new ~~data centre~~Data Centre where the counter estate is entirely Horizon, although the ~~data centre~~Data Centre is substantially ~~HNG-x~~HNG-x. A small number of Horizon ~~data centre~~Data Centre services are migrated in order to support this until the last counter has migrated to ~~HNG-x~~HNG-x. In addition the counter will have an NT4 to XP migration, which may occur some time after the ~~HNG-x~~HNG-x application migration. Horizon ~~Estate~~ Systems Management systems (SYSMAN2) need to be retained until this period is completed as the ~~HNG-x~~HNG-x Systems~~Estate~~ Management systems (SYSMAN3) cannot provide all the required services to NT4 platforms.

Migration and Transition are periods between the current Horizon system and the final end-point running only on ~~HNG-X~~HNG-X.

Migration starts with Horizon running solely at Bootle and Wigan and ends with all services running from Belfast. This period is expected to last for 6 to 8 weeks. The final state is a Horizon service running in Belfast that is ready for ~~HNG-x~~HNG-x Pilot.

Transition is the period between the end of migration and there being no Counter using any Horizon components. This includes the upgrade from NT4 to XP at the counter. Transition is complex and protracted but is operating entirely outside the ~~data centre~~Data Centre apart from the demise of a few services during the transition period. As such it does not generate extra states to be considered.

~~It should be noted that once PCI compliance is achieved the Hydra DCS agent will be decommissioned, which is one of the few Hydra systems outside BladeFrame (it requires a cryptography key disk during start-up which cannot be redesigned cost effectively *[DN: This is from a design meeting with Dave Johns. Need to confirm if this is still the case, but the current BOM has this as an external system]*~~*[DN: Add refs for key disk replacement solution]*).~~ Don't we also need key disks for the Generic Agent servers, which will stay until the end of Hydra?~~

Formatted: Highlight

~~The Horizon counters will then be using the BranchDB to record the obfuscated PAN, so the **BranchDB becomes critical to network banking much, much earlier than it becomes critical to Branch operation**.~~

Formatted: Strikethrough

~~[DN: This is not generally very clear in the design, and I only picked it up from Jeremy's migration slides]~~

Formatted: Font: Not Italic

~~[DN: Need to word this section better. BDB is being used to obfuscate PAN that would otherwise be in clear in Riposte. I believe this was to save having the Horizon counter perform the obfuscation?]. I think it is NPS that is doing this PAN Encryption (not Obfuscation) rather than BRDB. (I think Jeremy's slide may be wrong!)~~The Horizon PCI Counters do not rely on the Branch Database. They use HNG-X style messaging (via the BAL and using NPS) for transaction authorisation, but report transaction outcomes via Riposte as normal.

Formatted: Font: Not Italic, Not Highlight

Formatted: Font: Not Italic

The BDB is not used to obfuscate PANs - that is done on the Counter. For the PCI Horizon Counters the necessary seed is transmitted as a software distribution encrypted under the GDK. For HNG-X Counters the seed is provided at log-on by the BAL

The Pilot will use real outlets. Whilst the number will be small the fact that they are involved in a highly visible Pilot programme will mean that they receive attention out of proportion to the impact of failures on the overall business.

9.2 DR During Migration

It is a pre-requisite for migration that the Tivoli Workload Scheduler (TWS), the ~~HNG-x~~HNG-x Systems Management environment and the backup system are functioning. It may be necessary to demonstrate SYSMAN2, SYSMAN3 and TWS DR prior to starting migration, although these are merely part of Production LPAN DR described in 10.1.1.3 which will have been formally tested during ITU testing.

The high-level DR architecture has been deliberately presented in a similar fashion to the Migration Architecture, with blocks of services.

9.2.1 POLFS

SAP is a German company. SAP staff tend to use the term "Productive" where an English speaker would use "Production" and the two terms should be regarded as synonymous in this context.

POLFS is migrated in "Weekend A", which is actually three separate sub-phases for Dev, QATest and Production. There may be more than one elapsed week between phases. Also Weekend A is now at the end after Weekend D!

All POLFS servers are Fujitsu-Siemens PrimePower. The Production Central R3 server, and it's DR counterpart are PW1500 systems, and the remainder are PW450 with 4CPU and 8GB RAM, except for the IXOS servers which only have 2 CPU. XI servers (identified by a platform code of nws) have 16GB RAM.

Data exists in the Athene LTPDB and in various SAP reports to show the usage of CPU and memory during the batch processing and on-line windows, and is not of interest to DR provided similar systems are available at the secondary site in a suitable state to become the DR target.

The database servers run Solaris9 and Oracle9iR2, but this is installed as part of the SAP install, and DBA services are provided by the SAP Basis team not IRE11 DBA Team. Prism provide application development and support to POL in Chesterfield.

~~It has not yet been decided whether a~~A "lift & shift" approach ~~or a "swing kit" approach will be~~has been adopted on cost grounds. This has a significant impact on the ability to provide timely DR and regression path during the migration~~. Swing kit may change the PW1500 to PW850 type servers.~~

. Softek TDMF is providing the "Host Based volume replication" functionality~~Different approaches to data migration for a very large database are being explored which will have an effect on this section if deployed~~.

Readers should not waste time commenting on the time taken to migrate which is being discussed in a separate forum specifically looking at POLFS migration. These sections are merely here to illustrate the states for DR during migration.

Formatted: Default Paragraph Font, Pattern: Clear

Formatted: Pattern: Clear

Formatted: Pattern: Clear

Formatted: Pattern: Clear

9.2.1.1 Dev: Bootle -> IRE19 Weekend A-2

Dev consists of two servers, one running an SAP/R3 instance PLD and one running an XI instance DXI. There are no separate Dialog servers and all DI instances run on the central system.

No requirement for DR of Dev has been stated in CS/OLA/049, and the Dev systems are not a DR target for any other system, so there is no impact on DR of moving this service to IRE19. In fact in the event of a real disaster the provisioning of replacement hardware and a restore from tape backup is the most likely recovery. Provisioning hardware can take eight weeks.

In spite of the apparent insignificance of Dev, Fujitsu Services are aware of the potential impact on a set of highly skilled and expensive users having the system unavailable, and also of the impact on any urgent fixes.

9.2.1.2 QATest: Wigan -> IRE19 Weekend A-1

QATest consists of one servers running SAP/R3 instances PLQ and PLE, three R3 Dialog servers, and three servers running an XI instance QXI.

The R3 central instance server is the failover target for Production.

The XI central instance server is the failover target for Production.

The QATest R3 dialog servers can be switched to point to Production.

XI has two dialog servers which are failover targets for Production.

When the QATest system has been moved to IRE19 (assuming no swing kit) the DR path for Production is from Bootle to IRE19. ~~This involves shipping the Wigan tapes to IRE19 and restoring them. A Solaris BrightStorEB server will need to be retained in Belfast for this period, and the likely time to recover the database is a total of 30 hours, leaving 18 hours for application recovery, which is fairly unrealistic given that the tapes from Wigan will have been from a day earlier, and because of the length of recovery at least three overnight batch runs will need to be performed before the service is available.~~

~~Without swing kit there is a non-compliance with the recovery target for POLFS which is difficult to mitigate.~~Softek TDMF will be used to provide data replication from Bootle to Belfast, and a DR test of POLFS from Bootle to Belfast is planned as part of the pre-migration activity.

9.2.1.3 Archive

Archive is based on the IXOS product and EMC Centera storage. This is a physically separate Centera from the POL Audit solution.

There are two IXOS servers, one in Bootle which normally runs a DSP database and archives Production, and one in Wigan which runs the DS database and is used for archive testing of QATest.

Both IXOS servers write to both Centera and the Centera contain identical images. If one is unavailable, then the missing images are automatically "caught up" once it becomes available again, and there is a manual process which may be invoked to force this.

It is therefore safe to move Wigan system, test functionality, and then move the Bootle system. There is a slight lowering of resilience during this period, but DR from Wigan tape to the IRE19 (migrated Test) system can be accomplished readily even from tape as it is a very small database.

9.2.1.4 Production: Bootle -> IRE11 Weekend A

Production consists of one server running a SAP/R3 instance PLP which is around 3TB in size with six Dialog servers, an XI instance PXI and three XI Dialog servers.

Failover is obviously a little unbalanced as regards Dialog servers, and by the time migration occurs the details may differ slightly, but this has no significant impact on this document.

Softek TDMF is being used to provide "Host Based volume replication" functionality. This will allow the data to be available at IRE11 fairly instantly, although there is still time required to transfer the server itself.

~~The Production migration may be very protracted, up to 4.5 days in total, which even allowing for two being at the weekend exceeds a 48 hour outage.~~

Once Production is in IRE11 the DR process is identical to that which operated in Horizon. It is intended that this would be exercised in Weekend A+1, but depending on some of the swing kit and migration proposals it could be exercised prior to migration. This is BC Test 28.1 in Horizon.

9.2.2 Batch Services - Weekend B

Batch services in Horizon operate on a single Solaris server with an active system in Bootle and a Standby system in Wigan. The rest of the Horizon estate is designed to allow the Batch server to fail over independently of site DR (BC Test 3 - Database Server).

Batch services are migrating from Solaris9 and Oracle8 or Oracle9 to Solaris10 and Oracle10gR2. There is a state test for migration that Horizon systems can address batch services via the Oracle10 listener, but this is a migration issue not a DR issue.

All services except APOP and Maestro are being migrated on Weekend B, so although the Horizon server will remain active, and will continue to be able to fail over Maestro and APOP services as normal, the remaining services will move to IRE11.

The ~~HNG-x~~HNG-x batch services will be addressed via a VIP. As part of the migration all Horizon systems that need to address TPS, APS, LFS, TES, DRS, DW, RDMC or RDDS will be updated to use this VIP.

Horizon services, and POLFS systems ~~now in IRE11,~~ will not be aware (apart from a short outage) of the Batch services failing over from IRE11 to IRE19, and this DR does not change significantly from BCT3 except for changes associated with Maestro.

FTMS services will migrate during weekend B. Their primary function is to transfer data from DAT to the customer.

~~During this migration the TES Query service may also move to Belfast, while the Horizon TES Query Server continues to support APOP queries, but there is no reason why this should not be deferred to weekend C, provided testing shows that the Horizon forms server can talk to Oracle 10gR2. This is not strictly a DR problem; the query servers are only required to address the correct service via a VIP.~~

9.2.3 Online Services - Weekend C

Many of the o~~O~~nline service components in Horizon are stateless and run in active/standby with automatic failover. This is strictly part of the resilience function not the DR function, but it does mean that DR of these services (e.g. network banking authorisation agents) themselves is new at ~~HNG-x~~HNG-x.

Web~~Some~~ services such as PAF and DVLA ~~effectively r~~run active/active across sites with ~~service preference~~load balancing providing both~~as~~ the resilience and DR mechanism in Horizon. At HNG-x these will run active/active at the primary site with load balancing providing the resilience mechanism, but will use BladeFrame failover for DR.

A number of services, notably NPS and APOP have traditional oracle database type DR, similar to the batch services.

~~It is not clear whether t~~The MTAS service ~~is part of this set (MTAS moves with DCSM and in particular the Payment / EMIS stuff for DCS, which I believe is Weekend C. However there is an argument for moving this stuff at Weekend B since it is primarily "batch"), nor when the cut-over from ACDB/OCMS to BCDB/BCMS occurs. I believe that this is much later.~~will cease in Horizon (on DCSM) and be offered from HNG-x (on EST)

With the exception of the DCS agent prior to PCI compliance all ~~HNG-x~~HNG-x online services components reside in the Production LPAN, and site DR is achieved as part of Production LPAN failover.

~~It is expected that the DCS agent will operate active/active whilst it is external. Once PCI compliance is achieved this agent will be decommissioned, and a Production LPAN resident service will take over, but this does not occur until many weeks after migration.~~The Hydra DCS agent will migrate in Weekend D along with NRA and other Horizon Branch services

During this phase there needs to be a trust between the ~~HNG-x~~HNG-x AD and Horizon domains.

~~*[DN: Need to clarify from Network HLD whether Production LPAN failover is transparent to Horizon Bootle/Wigan based systems]*~~

The operation of services in IRE19 or IRE11 will be transparent to systems remaining in Bootle & Wigan. As any service fails over the service address remains identical.

Formatted: Font: Not Italic

9.2.4 Branch Services - Weekend D

The ~~HNG-x~~HNG-x Branch Service only exists in IRE11 with a DR capability in the Production LPAN to IRE19.

The Horizon branch services are migrated on Weekend D. Only a small number of systems, notably the Riposte message store servers, Estate Management systems (ACDB and OCMS) and the Key Management Server, have significant data associated with them. The remainder, such as Generic Agents, ~~NWB~~NBS Routing Agents, VPN servers, boot server and RADIUS servers are stateless.

~~*[DN: Need to check RADIUS servers - where do they get auth from]*~~

After the Horizon Branch services have migrated they will be referred to collectively as "Hydra" to distinguish them from Horizon and ~~HNG-x~~HNG-x services. These are decommissioned at the end of Transition.

Once this migration is complete the estate is at the migration end point, and no special consideration is required post Weekend D.

Following the migration Horizon Branch services operate active/active from Belfast. This is a requirement of the correspondence server resilience model. The counters are already set up to prefer a particular correspondence server, and to try others in the same cluster at both sites in a round-robin fashion.

As each counter moves to ~~HNG-x~~HNG-x it will prefer IRE11 which is where the BAL is located. This will be similar as each counter becomes PCI compliant it will prefer IRE11 for DCS and NB Auth traffic.

[DN: Need clarification on SYSMAN2 migration, esp. EACRR which is strongly coupled to Hydra]

9.2.5 Audit

The Audit service migration requires careful management and the use of swing kit for EMC Centera.

The Audit system is designed to be down for up to three days, although this is undesirable as it would affect retrievals. The actual migration will be coupled to Weekend D to minimise traffic flows across the WAN (a large amount of Riposte audit data is collected each night).

Audit will operate, certainly initially during Transition, as an active/active system and as such DR is not any different to Horizon, and there is no special consideration required.

9.2.6 Network Management

The Network Management Systems, primarily HP OpenView and Cisco Works, plus some small diagnostic tools, collect and forward events to SYSMAN and provide backup and recovery services for switches. There will also be syslog servers for events forwarded from switches.

These systems all operate essentially active/active and there is no special DR consideration.

~~*[DN: Need to check this for OpenView]*~~

9.3 DR During Transition

During Transition, the Horizon components that are in use will use the same resilience and DR procedures as before.

These are:

- Correspondence servers.
- Generic Agents
- VPN Servers

Formatted: Bullets and Numbering

~~The correspondence servers will be running on virtual machines.~~

~~□Generic Agents~~

Formatted: Indent: First line: 0.63 cm

~~□VPN servers.~~

- Routing agents
- Boot server
- KMS

Formatted: Bullets and Numbering

For all systems except KMS t~~T~~his will be covered by the active/active PAN, and no special consideration is required.

KMS retains the current SRDF based database failover with some minor procedural modifications owing to the fact that a BladeFrame based system cannot interact with Symmetrix, and also that the SYMCLI version supported by NT4 is not supported with the generation of Symmetrix firmware in Belfast.

A component of SYSMAN2 known as EACRR (Enhanced Agent and Correspondence server Resilience and Recovery) manages services on the Generic agents. Consideration of how EACRR operates during DR will be described in the EACRR HLD and is not of interest to this design which is completed once the failed over infrastructure is available to applications.

10 System Parts

This section details each service. It will cover:

1. A brief overview of the service and its availability requirement

2. A brief description of the resilience mechanism

3. An overview of the service criticality and business impact of failover.

4. The HLD reference for this service.

Some of this section will appear to duplicate material that has been presented earlier. This is intentional, and is designed to make the description of each component easy to read.

Note that the list is by platform type. It is not always obvious when a platform runs several services.

Conflicts with earlier sections should be reported by reviewers. If the reviewer is the authority for that component they should also state the correct position.

Missing chapters should be highlighted by reviewers. This section is intended to cover every service and every platform. The general structure is intended to discuss the platform and then the services it supports, but in the case of very simple systems the platform and the service are combined into a single section.

In order to facilitate early publication for review some sections are not yet completed. Only the authority for that service should return a comment on that section, ideally as a suitable form of words for the section.

10.1.1 Discrete Core Services

This section describes services which operate in an active/active or active standby mode, but are supported on hardware external to the BladeFrame.

[DN: Need to add HLD references to each section?]

10.1.1.1 Network

ARC/NET/ARC/0001 - Network Technical Architecture

DES/NET/HLD/0008 - LAN Design

DES/NET/HLD/0009 - WAN Design

DES/NET/HLD/0014 - Branch Access

DES/NET/HLD/0015 - Transit LAN

Brief piece about switches operating in pairs

Bit about WAN and C&W cloud

Brief piece about resilient firewalls

Brief piece about how Test domains operate as sub-domains (VRF) and in effect Production is a sort of sub-domain of a Management super-domain.

[DN: There may be some testing issues as V&I wish to exercise control over rig time, but the Management domain and the V&I (Production) Domain have shared authentication and audit services. Details in V&I HLTP TST/GEN/HTP/0002 ?]

10.1.1.1.1DWDM Inter-site link

Ref: DES/NET/HLD/0004

The inter-site links are a shared service from Fujitsu Service Corporate Networks. A pair of Dense Wave Division Multiplexors are provided at each site, designated North and South. Leased dark fibre from NTL carries an FTel service to provide diversely routed links of approximately 28km and 48km. EMC have measured a ping response time of < 3ms which provides considerably lower latency than the similar links between Bootle and Wigan. Four 4gbps fibrechannel and four 1gbps ethernet links are provided for the SAN and Network switches respectively.

The inter-site link for the Branch DMZ provides the resilient triangulation to protect against loss of C&W link to either IRE11 or IRE19.

10.1.1.1.2Access Switch

The access switch provides a secure location to host VLANs for untrusted traffic. In order to communicate with a core switch the traffic must exit the access switch and travel via a resilient pair of external firewalls.

Deployed in pairs, one pair per site.

~~Inter-site ISL blocked on Switch 2 to prevent spanning-tree loops means that if inter-site link between Switch 1 fails manual action is required to restore full service between sites.~~

~~The inter-site link for the Branch DMZ provides the resilient triangulation to protect against loss of C&W link to either IRE11 or IRE19.~~

10.1.1.1.3Core Switch

Deployed in pairs, one pair per site.

Switch 1 is preferred for traffic (cuts down ISL traffic to track state)

Inter-site ISL is resilient per pair of switches and is blocked on Switch 2 to prevent spanning-tree loops means that if inter-site link between Switch 1 fails manual action is required to restore full service between sites.

[DN: Is this still the case?]

In the latest design a second layer (confusingly called the access layer) of Cisco 3750 switches is used to provide port scaleability, with the 6513 offering aggregation and routing services, including FWSM and ACE.

Formatted: Font: Italic, Highlight

Formatted: Font: Italic

10.1.1.1.4Cisco 2811 Router

This router is used to provide customer-edge (CE) and hand-off (HO) functionality. These routers are deployed as resilient pairs.

The 2811 only has a single power cord. Banks of 2811 have been deployed in "header" and "footer" comms cabinets, where each bank is connected to a different ~~data centre~~Data Centre PDU so that in the event of a PDU or UPS failure only one bank will fail.

10.1.1.1.5Cisco ASA5540

This "security appliance" is deployed in pairs to provide resilient firewall layers as required between access and core switches, and to protect HO and CE routers.

The 5540 only has a single power cord. Banks of 5540~~2811~~ have been deployed in "header" and "footer" comms cabinets, where each bank is connected to a different ~~data centre~~Data Centre PDU so that in the event of a PDU or UPS failure only one bank will fail.

10.1.1.1.6FWSM Firewall Service Module

The FWSM is a blade in the Cisco 6513 switch that allows VLANs to be separated by an enterprise class firewall.

The FWSM resilience is provided by the corresponding FWSM in the paired switch at the same site.

10.1.1.1.7ACE Load Balancer

The Cisco Application Control Engine is a load balancer (similar to the Content Switch Module) which is provided as a blade in the Cisco 6513 switch.

The ACE works as a pair across sites, with local N+1 resilience being provided by the locally paired switch.

The ACE watches for services being advertised, e.g. a particular server starting a service on port 80, and advertises a preconfigured virtual IP address (VIP) for that service. Each application owner that requires a VIP should describe how their service interacts with ACE in the network section of their HLD.

This is the standard method of providing virtual IP addresses in ~~HNG-x~~HNG-x.

~~10.1.1.1.10~~10.1.1.1.8 [NMN] OpenView

Formatted: Bullets and Numbering

Ref: /DES/NET/HLD/0012

HP OpenView on Sun V890

Active/standby? How does failover work?

One per site, although there is also a Test set in IRE19 which could be redeployed in the event of loss of IRE11.

Need to understand importance to service if there is a loss of this service (up to 1 day while it is repaired) following a loss of IRE19.

~~10.1.1.1.11~~10.1.1.1.9 [NCW] Cisco Works

Formatted: Bullets and Numbering

Ref: /DES/NET/HLD/0012

Hosted on Sun 280R.

Active/active.

One per site, although there is also a Test set in IRE19 which could be redeployed in the event of loss of IRE11.

Toolset to permit management of Cisco switches.

Primary purpose during DR is to provide a means for support staff to manage the 18 subnets through pre-scripted tasks.

This task can also be performed manually by a 3rd line Network Support person, but such manual intervention is likely to be slower and prone to human error.

~~10.1.1.1.14~~10.1.1.1.10 [NAP] AlarmPoint

Formatted: Bullets and Numbering

Ref: /DES/NET/HLD/0012

Discrete because a dial-out line is required.

AlarmPoint is the mechanism by which alerts are raised to pagers held by the second line support teams. There will be two AlarmPoint servers operating in an active/active mode.

There should also be an LST AlarmPoint server to allow testing of the event raising mechanism.

This is not the primary means of notification, but serves as a way of alerting support staff to serious events.

One per site, although there is also a Test set in IRE19 which could be redeployed in the event of loss of IRE11.

[DN: Horizon had three BTI systems to keep N+1 in the event of ~~data centre~~Data Centre failure. Should a third system be based in BRA01?]

~~10.1.1.1.13~~10.1.1.1.11 [NPC] Network Packet Capture

Formatted: Bullets and Numbering

Ref:

System to capture packets for traffic analysis and diagnosis. These are deployed ad hoc to resolve problems. Two are provided per site simply so that more than one problem may be analysed at the same time.

In a switched network environment traditional sniffers have a hard time. Cisco allow switch ports to be set to a diagnosis mode to capture traffic, and tools such as WireShark or Ethereal allow analysis of the captured packets.

Most sensitive packets on the RMGA estate already have the data payload encrypted, and this tool really only allows for analysis of flow and protocol problems, such as nfs mounts not succeeding.

~~10.1.1.1.16~~10.1.1.1.12 [NTP] Network Time Protocol

Formatted: Bullets and Numbering

Ref: /DES/NET/HLD/0013

A pair of Galleon NTS6000 servers are deployed at each site at ntp stratum 0.

These are able to take a time source either from GPS satellites or from the MSF time signal. The clocks can vote amongst themselves to provide an accurate signal, or indicate to a client whether they have a poor time. The ntp client in Solaris, RedHat and the Cisco IOS is able to handle this gracefully.

For Windows platforms the ACD platform will act as a stratum 1 time source.

~~10.1.2.3.6~~10.1.1.1.13 [VPN] VPN Servers

Formatted: Bullets and Numbering

Active/active.

12 per site allow each site to be independently N+1 resilient.

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

These will be deployed on four RX300 servers [VSD Platforms] running Microsoft Virtual Server Host, each hosting three NT4 VPN guests. They are discrete to allow connection to the Access switches.

The service is required for as long as there are NT4 counters in the estate which may be long after the Horizon application has disappeared.

Active/active. No failover.

Recovery by reprovisioning.

~~[DN: CP0039 is discussing the retention of VPN past the Hydra phase]~~

10.1.1.1.14 [VPM] VPN Policy Manager

Formatted: Bullets and Numbering

[DN: need some words]

Active/active. No failover.

Recovery by reprovisioning.

Deployed on a shared VSD with VDW and VEX

10.1.1.1.15 [VDW] VPN Loopback Workstation

Formatted: Bullets and Numbering

Active/active. No failover.

Recovery by reprovisioning.

Deployed on a shared VSD with VPM and VEX

Formatted: Normal

10.1.1.1.16 [VEX] VPN Exception Server

Formatted: Bullets and Numbering

Active/active. No failover.

Recovery by reprovisioning.

Deployed on a shared VSD with VDW and VDW

10.1.1.2 Storage Area Network

Ref: DES/NET/HLD/0007 SAN High Level Design

Ref: DEV/INF/ID/0003 SAN Configuration Document

Ref: DEV/INF/LLD/0043 Storage Mapping Document

10.1.1.2.1[SAN] MDS9509 SAN Switch

Ref: DES/PPS/HLD/0007 Storage High Level Design

There will be a pair of MDS9509 switches at each site operating as a mirrored fabric. Both switches make use of both inter-site links, North and South

Multiple line cards, multi-pathing, alias based zoning.

VSANs will be used to segregate logical sets of data, primarily for performance reasons.

LUN masking and internal BladeFrame LPAN presentation will ensure that Test and Production disks are only presented to the correct server instance.

10.1.1.2.2[RSG & RGP] EMC Remote Support Gateway

Ref: DEV/INF/LLD/0030

The EMC Remote Support Gateways (RSG) are on discrete RX300 as they face the internet. They provide a secure means of access for EMC engineers (using RSA SecurID tokens) as well as providing a generic mechanism for alerts from EMC equipment to be sent direct to EMC. This allows a timely response, e.g. to replace a failing disk, usually before any application is aware of the pending failure.

There are two gateways, one per site, operating in active/active mode. If support is required during a DR EMC must have a means of providing timely support.

The EMC Remote Support Policy Server (RSP) authenticates connections, but as the gateways have a short memory of recent connections it is not a service that requires high availability.

10.1.1.2.3[ECC & ECA] EMC Control Centre

Formatted: Portuguese (Brazil)

Ref: /DEV/INF/LLD/0029

EMC Enterprise Control Centre provides a centralised service for SAN and storage management. The ECC Server itself supports a database and a data collection service, and will fail over in DR.

There are a number of data collection and control agents. To unload the ECC Server and provide a level of performance scalability and resilience more than one ECC Agent server is usually deployed.

These services are not critical either to normal operation or to failover as other systems will alert upon failure, e.g. OpenView will report events from SAN Switches, and the attached servers themselves will report disk failures, but they allow a more effective diagnostic response, and they considerably simplify the SAN and storage management tasks.

10.1.1.2.4Storage Arrays

Formatted: Heading 5

Formatted: Bullets and Numbering

Ref: /DEV/PPS/HLD/0007

Ref: /DEV/INF/LLD/0004

EMC DMX-3 storage arrays are used to provide storage service classes 1 and 2. There are two pairs, A and B, one of which hosts BRDB and the other the Standby so that a storage array fault is unlikely to compromise the ability to offer a Branch service.

EMC Clariion CX3-80 is used for storage classes 2 to 6. There is only one system per site, but for these storage classes by definition the data is not mission critical and may be recovered from backup, or by holding separate copies on the Clariion at each site.

EMC Centera CAS is used to store Audit data and POLFS Archive data. These are existing solutions migrating from Horizon and are better covered in the sections on ARC and IXO platforms.

10.1.1.3 [PAN] BladeFrame PAN Manager

Ref: DES/PPS/HLD/0025 Bladeframe High Level Design

Ref: DEV/GEN/SPE/0007 Platform Hardware Instance List

Ref: DEV/INF/LLD/0066 BladeFrame Failover Low Level Design

The PAN Manager service permits resources to be allocated to logical groups known as LPANs.

Each PAN Manager service is hosted on a pair of control blades (cBlades) which provide a service address for that frame, and also provide resilience for the PAN connections to LAN and SAN. These cBlades are managed as appliances and are replaced as units by the hardware support engineer.

There are four BladeFrame chassis at each site. Three of these at the primary site normally operate the Production LPAN. Three equivalent chassis at the secondary site are available for the Production LPAN to fail over to, and normally support a number of Test LPANs. The fourth frame at each site is permanently assigned to Production use for active/active systems, and is expected to host Hydra systems running in virtualised environments, for instance, as well as services such as SAS, AD and DNS.

The 'frames operate in pre-designated pairs, so services in bf002 will always fail over to bf001 etc. The LPAN definitions on bf001 and bf002 are identical with the exception of VLAN ID's for VLAN tagging. It is thus possible to deploy changes to the LPAN configuration at the secondary site before applying them to the primary site.

In principle it is possible to operate Production and Test services concurrently, but in practice there are unlikely to be enough pBlade resources to support this.

Certain basic design rules are described in the HLD to maintain optimum resilience, for example:

Cluster members should not share a power domain

Cluster members should boot from diverse EMC cabinets

Services such as BAL which provide N+1 resilience should boot from diverse EMC cabinets

Dual cBlade failure is equivalent to the loss of that chassis, although there are circumstances where the PAN Manager service may be lost but the I/O virtualisation functionality continues. This need not necessarily result in a DR. A cold reset of the cBlades takes about 20 minutes, and is the primary recovery mechanism. This would happen in parallel with the escalation on loss of service.

The cBlades may be rebooted in turn, a process known as a "rolling reboot", in order to reset them e.g. in the event of major storage layout changes or after applying a patch. A rolling reboot does not impact on pServer operation, and is an allowable operation during normal working hours.

10.1.1.4 [VSH][VSD] Virtual Server Host

Ref: /DES/PPS/HLD/0004

It is getting very difficult to source equipment that will run NT4SP6A. To work around this a dummy platform with W2003 Enterprise Edition has been created to allow NT4SP6A instances to run under Microsoft Virtual Server Host.

This also permits several NT4 services to be hosted by a single platform. The majority of the hosted services are over five years old and have a relatively low memory and CPU requirement.

The Correspondence Servers have a fairly high I/O profile, but this still allows a smaller service such as a domain controller to be co-hosted.

This platform is available both for BladeFrame hosted and discrete systems. The discrete is known as VSD.

10.1.1.5 [BSM, BSS, BSL, BSW, EDL] Backup

Ref: DES/SYM/HLD/0015

The backup servers are active/active and do not perform DR.

The Solaris backup servers have EMC SYMCLI and NAVICLI installed and perform the additional function of storage management servers. Any scripts required to fail over storage will be run from these servers.

The Net Backup Master Catalogue Service needs to fail over in order to permit restores. This would be simplest if it was in the frame, but the current plan is to present an SRDF disk to the RHEL Media Server which will also run the master service in active/standby mode.

[DN: ~~A~~ CP0048 ha~~i~~s been~~ing~~ raised to move the master service into the frame on a separate platform [BSM] as the current DR solution is not supported by Symantec]

10.1.1.6 [MSH] Hydra Maestro Master

Ref:

A small Solaris server is required to continue running maestro as the TWS will not support NT4 and the old version of maestro does not run under linux.

This will be a SunFire V125 server which mounts a single SAN disk as /opt/maestro with a standby server in the secondary ~~data centre~~Data Centre. The same RL2 / RL3 as the existing Horizon batch solution uses is an appropriate design for making a platform active.

10.1.1.7 [ENT] Hydra RSA SecurID Server

Ref:

This server provides two-factor authentication for Horizon systems which will not be able to join the AD domain, and therefore will not be able to use the Vintella two-factor authentication.

There is a simple master/slave relationship with one server at each site, although later versions run as peers.

~~Whether this is a simple continuation of the Horizon system on Solaris 2.6 or is upgraded for the remainder of Horizon is not clear, but it is difficult to get servers which support Solaris 2.6, and a hybrid solution with a slightly updated ACE server on Solaris 10 has been piloted in INF1~~This will be a continuation of the Horizon solution on Ultra10 and Solaris 2.6. Retired systems will be retained to provide spare parts.

~~10.1.1.8 [KMS] Hydra Key Management Server~~

Formatted: Bullets and Numbering

~~Ref: /RS/MAN/013~~

~~This is a SQL-Server database running on Windows NT4 in a special security domain.~~

~~The database and any required file store are on the Key Management Server S: drive which is on EMC SRDF replicated storage. BCVs are not used, SQL-Server does a dump to disk and areas of the S: drive are backed up. There is a security concern over the possible theft of this database, and as a result cold backup images are rarely made.~~

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

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

~~It is doubtful whether EMC Solutions Enabler (SYMCLI) will work inside a VSH platform, and in any case the version of SYMCLI that was compatible with NT4 is not compatible with current EMC microCode versions on the DMX. The Booter service will probably need to be worked around in order to allow service startup, and the failover procedure will need to be modified to interact with the Storage Management Server (BSS, the Solaris Backup Server).~~

~~Resilience: fail over to secondary.~~

~~10.1.1.9~~10.1.1.8 [EBS] Enterprise Boot Server

Formatted: Bullets and Numbering

Ref: /DES/PPS/HLD/0024

The Enterprise Boot Server provides kickstart, jumpstart and PXE boot services during initial builds. It is not required to be available except during builds, and the most straightforward model is to simply have one at each site, active/active, and simply point builds at the preferred system.

~~*[DN: bootp is not clear]*~~

It is a manual build (it is the very base system in the provisioning solution), although in principle rebuilds could be via TPM, but that probably would require maintaining two platform types in Dimensions.

~~The current design seems to be that each service (System Test, SV&I, RV, LST, V&I/Production) has its own EBS.~~

The BladeFrame management interface needs to mount a share from the EBS to perform kickstart builds for RedHat.

~~10.1.1.10~~10.1.1.9 [NAS] Networked Storage

Formatted: Bullets and Numbering

Ref: /DES/PPS/LLD/0004

EMC ~~Celera~~Celerra is used to present SAN storage as network shares. In the present configuration this storage is hosted on the EMC Clariion CX3-80 and replicated between sites by MirrorView. On failover the MirrorView failover occurs first, and the slave ~~Celera~~Celerra at the secondary site then takes over presenting the shares.

As with all synchronously replicated storage this does not protect against corruption, and any data stored here should be backed up as necessary. This is straightforwardly achieved either via Clariion SnapView replicas presented to a backup server over the SAN, or for small repositories as a network backup.

The EFS share a repository with EPM. There is one huge read-only repository shared amongst all EPM instances (one per rig), plus a smaller rig-specific repository. These are delivered to via DXC by the Configuration Management Workstation as CM pass DPVB's to TPM for onward distribution.

The Branch Database nodes also use NAS to avoid the need for a clustered files ystem, e.g. to write audit data. Any node is able to write data, and the Audit Gatherer service can simply be pointed at the share rather than having to trouble a node for the data.

~~10.1.1.11~~10.1.1.10 [CON] Aurora Console Tower

Formatted: Bullets and Numbering

Ref: /DES/SYM/HLD0020

Aurora is used to manage all (serial) consoles for equipment such as cBlades, Solaris servers and Cisco switches in a ~~controlled and logged~~secure and auditable manner without physical access to the Data Centre being required.

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

Aurora1 manages bf001/cb1
Aurora2 manages bf001/cb2

Aurora1 manages core2 network switch (from mgmt LAN on core1)
Aurora2 manages core1 network switch (from mgmt LAN on core2)

Aurora1 manages Aurora2 and vice versa

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

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

There is no DR requirement for Aurora itself, but it is a critical component in Solaris DR to allow properly managed reboots. Emergency re-patching will allow a continued service for a limited number of servers.

~~10.1.1.13 [DXC] Corporate Data Exchange Proxy~~

Formatted: Bullets and Numbering

~~Ref: /DES/NET/HLD/0018~~

~~Active/Active. One per site.~~

~~Permits safe transfer of data from corporate networks to the HNG-x network, e.g. software packages released from Dimensions being transferred to the TPM Repository.~~

~~10.1.1.14~~10.1.1.11 [DXI] Internet Data Exchange Proxy

Formatted: Bullets and Numbering

Ref: /DES/NET/HLD/0017

Active/Active. Two~~One~~ per site. Secure Appliance WebWasher 1150.

Permits safe transfer of data from the Internet to the ~~HNG-x~~HNG-x network and vice-versa, e.g. software packages released from Dimensions being transferred to the TPM Repository.

~~10.1.1.15~~10.1.1.12 [SPS] Supplier Access Server

Formatted: Bullets and Numbering

Ref: /DES/SYM/HLD/0017

This is a solution to allow support by third parties, e.g. Fujitsu Siemens or Oracle.

~~Not well defined but basically a SAS to let 3rd party suppliers have managed access.~~A CP has been raised to move this from BladeFrame to discrete as it needs to sit in the internet facing Access layer of the network.

[DN: Currently deferred to Release 2]

Formatted: Font: Italic, Highlight

Formatted: Font: Italic

~~10.1.1.16[PMN] Patch Management Server~~

Formatted: Bullets and Numbering

~~Ref: /DES/SEC/HLD/0006~~

~~Not clear if this is external.~~

~~10.1.1.17[AVS] Anti-Virus~~

Formatted: Bullets and Numbering

~~Ref:~~

~~Not clear whether this will be an appliance~~

~~10.1.1.18~~10.1.1.13 [HSM] Hardware Security Module

Formatted: Bullets and Numbering

Ref: /DES/SEC/HLD/0002

Atalla AS10150 appliances replace the PCI cards in the Network Banking Agents.

There will be ~~four~~ three per site deployed active/active which caters for both site DR and resilience.

One~~Three~~ appliance~~s~~ will be deployed in IRE19 for Test services, which may reuse. The Atalla is primed by a key disk and these differ for Test and Live services.

There will also be smaller AS8150 in BRA01 and LEW02 for key generation.

~~10.1.1.19~~10.1.1.14 [VNS] Vulnerability Scanning Server

Formatted: Bullets and Numbering

Ref: /DES/SEC/HLD/0008

Foundstone FS1000~~ management~~ system~~ plus two FS850 per site.~~, one per site.

Each site is effectively an independent deployment.

There is not a very high availability requirement on these systems. As long as vulnerability scans are performed in a reasonably timely manner that is sufficient, so recovery by replacing with a spare is adequate.

The IRE19 set will be part of LST in normal circumstances, and be reconfigured in the event of losing IRE11.

The definitions of scans will be backed up so that they may be recovered in the event of a spare being provided.

[DN: Not clear how these are provisioned in the first place]

~~10.1.1.20~~10.1.1.15 [DAT] Legacy Batch Server

Formatted: Bullets and Numbering

Ref: [Platform HLD]
Ref: /DEV/INF/LLD/0065 Solaris Failover LLD

The DAT service is effectively hosted by a four node cluster of PrimePower PW650 servers, giving N+1 resilience at each site. The same model is used as for the POLFS XI Server as described in EA/DES/001. The data is shared via SAN storage, which also enables site failover.
Cisco ACE detects which of the four systems has started an Oracle listener and advertises a service VIP on behalf of that server. Only a server with access to the disks will be able to advertise.
The server hosts a number of services including eight Oracle databases and the master TWS scheduling service. These are discussed individually in the following sections.
All four servers need to talk to ntp and ssh on the SAS for management support, possibly also NBU client and TCA.
Audit gathering is from each application not from the underlying platform.
Not sure how DNS is configured for this as an overall service. In effect we have a four node cluster offering a single virtual service.
Because of the use of ACE there is the possibility that if the storage is "split" rather than failed over, such as may be done during a major migration to preserve a regression image, that both sites will offer a service. Procedures for splitting the ~~data centre~~Data Centre should ensure that the listener is inhibited.

For this reason also it is recommended that the default run level be made 2, and the service start up be a manual event.

[DN: This has an impact in situations where the server reboots, e.g. loss of a processor]

~~*[DN: will there be a HNG-x Dimensions reference for the POLFS resilience design?]*~~

~~10.1.1.20.1~~10.1.1.15.1 TSH Tivoli Workload Scheduler

Formatted: Bullets and Numbering

Ref: /DES/SYM/HLD/0016

The master Tivoli Workload Scheduler will be hosted on DAT as at Horizon. This has two advantages; firstly this is a resilient platform, and secondly most of the scheduled jobs actually run on DAT anyway, and much of the ancillary scripting from Horizon can be simply redeployed.

~~10.1.1.20.2~~10.1.1.15.2 DRS

Formatted: Bullets and Numbering

Ref: /DES/APP/HLD/0033

Data Reconciliation Service. Actually two services in one system, ~~NWB~~NBS and DCS.

Banking and Debit Card have different settlement periods, Debit Card being next working day, and Banking being same day.

Collects C12's

Produces CAPO and A&L REC files

Receives LREC file from Link

Produces Payment file for Streamline, and receives EMIS file.

Produces banking reports

~~10.1.1.20.3~~10.1.1.15.3 TES

Formatted: Bullets and Numbering

Ref: /DES/APP/HLD/0036

Transaction Enquiry Service

Supports queries from POL users in Huthwaite via TESQA.

Also provides C4/D feed to DRS that used to come from IBM NBE.

~~10.1.1.20.4~~10.1.1.15.4 TPS

Formatted: Bullets and Numbering

Ref: /DES/APP/HLD/0027

Transaction Processing Service.

Harvests EPOSS transactions.

Provides feeds for POLFS (BLE) direct via NFS

Provides POL-MIS (W_jjjnnn.TP_.pz~~R~~) via TIP FTMS Gateway

Provides ~~NWB~~NBS feed to DW via local file system.

Provides C112 data to DRS for reconciliation

Provides Client Transaction Summary (CTS) file to POL (W_jjj500.TP_.pz)

~~10.1.1.20.5~~10.1.1.15.5 APS

Formatted: Bullets and Numbering

Ref: /DES/APP/HLD/0026

Automated Payment Service

Harvests APS transactions.

Provides files to clients via EDG and a client transaction summary (CTO) to POL-MIS.

~~10.1.1.20.6~~10.1.1.15.6 LFS

Formatted: Bullets and Numbering

Ref: /DES/APP/HLD/0037

Logisitical Feeder Service

Ac~~This function could probably have been incorporated into the BranchDB at HNG-x. (but it hasn't been)~~

~~Ac~~ts as the interface between SAP/ADS and ~~Riposte~~ the Branch estate for planned orders, cash on hand, pouch transfers etc.

All communication with ADS is via TIP FTMS Gateway.

~~10.1.1.20.7~~10.1.1.15.7 DWH

Formatted: Bullets and Numbering

Ref: /DES/APP/HLD/0082

Data Warehouse.

Used to be a big SLA calculator. This is the only system in the estate truly designed to run asynchronously, and can catch up independently provided the input files have not been lost. Used to be on a separate platform but was consolidated at S50.

At Horizon is used to calculate some message delivery SLAs, but these really disappear at ~~HNG-x~~HNG-x (not at Hydra).

Produces some banking reports that could probably come from DRS. I think these are likely to be removed.

Some "extra" functionality at ~~HNG-x~~HNG-x but details not yet clear.

~~10.1.1.20.8~~10.1.1.15.8 RDMC

Formatted: Bullets and Numbering

Ref: /DES/APP/HLD/0004

Ref: /DES/APP/IFS/0004 ~~HNG-x~~HNG-x Ref Data Delivery

Reference Data Management Centre.

Receives updates from POL/RDS. These are tested on the RDT systems and then marked for release using the RDMC Workstation.

RDMC is also used for memos to PMs and for "urgent" ref data such as Bureau de Change rates which bypass the RDT checks.

~~10.1.1.20.9~~10.1.1.15.9 RDDS

Formatted: Bullets and Numbering

Ref: /DES/APP/HLD/0005

Ref: /DES/APP/IFS/0005 (~~HNG-x~~HNG-x Counters)

Ref: /DES/APP/IFS/0001 (BranchDB)

Ref: /DES/APP/IFS/0002 (BranchDB)

Ref: /DES/APP/IFS/0003 (SYSMAN)

Reference Data Distribution Service

Copies released data from RDMC and transforms it to be suitable for all client systems. Many clients read the data via ODBC database links, DW gets input files, and a Loader is run for Riposte.

Branch Database will probably use ODBC.

~~10.1.1.21~~10.1.1.16 SYSMAN2

Formatted: Bullets and Numbering

SYSMAN2 is the generic term for the Horizon version of Tivoli estate management. NT4SP6A will not run under the latest version of Tivoli, so SYSMAN2 must be retained for as long as there are NT4SP6A platforms requiring Tivoli management.

The resilience model is unchanged from Horizon.

It is expected that the SYSMAN2 Primary site will be IRE19 to continue the model established in Horizon that in the event of failure of the primary site (IRE11) SYSMAN2 is already available to manage the Hydra services at the secondary site.

Formatted: Not Highlight

Formatted: Not Highlight

~~*[DN: Not sure where this should be stated, but it is to allow the management system to be already up in the event of losing the primary site for the remainder of the service]*~~

Events ~~may~~ will be collected direct to SYSMAN3 from the counters to give a single management view of the branch estate. Retiring servers managed by EACRR will direct events to SYSMAN2.

~~[DN: Not clear whether this will happen]~~

Formatted: Font: Not Italic

~~10.1.1.21.1~~10.1.1.16.1 [OMD] Inventory Server

Formatted: Bullets and Numbering

Active/standby Oracle database server.

The OMDB database is critical to the management of the Horizon estate.

This is an SRDF based Oracle database, failover similar to DAT and KMS. It is discrete to permit continued use of BCV backups

~~*[DN: It is not yet clear whether all of the services listed below are required in Belfast]*~~

~~10.1.1.21.2~~10.1.1.16.2 EACRR Enhanced Agent & Correspondence Server Resilience & Recovery

Formatted: Bullets and Numbering

Ref: DES/SYM/HLD/0007

Formatted: Normal

This is a service effectively hosted on OMD which filters and interprets events to provide resilience in the agent services hosted on the AGE platforms.

Some of the services which EACRR detects or manages *are moving to the SYSMAN3 domain.*

Formatted: Font: Italic

Counter events from Hydra counters will go direct to SYSMAN3, but the data centre systems will continue to event to SYSMAN2 to allow continued operation of EACRR.~~[DN: The method by which SYSMAN2 and SYSMAN3 interact to allow EACRR to continue working in Hydra is not clear. It is being investigated by the SYSMAN2 Migration Team]~~

Formatted: Font: Not Italic

~~10.1.1.21.3~~10.1.1.16.3 [OMA~~SDC~~] Domain Controller

Formatted: Bullets and Numbering

Formerly the OMDB Archive server this platform now provides only the Domain Controller function for SYSMAN2.

Formatted: Normal

Primary in IRE19, Backup in IRE11, standard NT4 model.

~~10.1.1.21.4~~10.1.1.16.4 [DEL~~SDS~~] Delivery Server

Formatted: Bullets and Numbering

Ref: Software Distribution

Formatted: Highlight

~~Not clear whether this interacts with the NAS Repository or whether two delivery streams continue, one for~~ Horizon ~~and one for HNG-x~~retiring platform.

~~10.1.1.20.5~~10.1.1.16.5 [T~~S~~MR] Master TMR

~~10.1.1.21.6~~10.1.1.16.6 [SMT] Master TEC

Formatted: Bullets and Numbering

~~10.1.1.21.7~~10.1.1.16.7 [SCT] Client TEC

~~10.1.1.21.8~~10.1.1.16.8 [SEC] Expedited TEC

Required as long as EACRR is needed.

Formatted: Normal

~~10.1.1.21.9[SNT] SNMP TEC~~

Formatted: Bullets and Numbering

~~10.1.1.21.10[SRT] RADIUS TEC~~

~~10.1.1.21.11~~10.1.1.16.9 [SST] S-TEC

~~10.1.1.21.12~~10.1.1.16.10 [LGW] Login Gateway

~~10.1.1.21.13~~10.1.1.16.11 [~~S~~PGW] Post Office Gateway

~~10.1.1.21.14~~10.1.1.16.12 [SSG] Secure Post Office Gateway

Boot loader DMZ

Formatted: Normal

~~10.1.1.21.15~~10.1.1.16.13 [TGW] Campus Gateway

Formatted: Bullets and Numbering

10.1.2 Services in Active LPAN

Ref: ARC/PPS/ARC/0001

These are services which require an active/active DR model, or which use master/slave replication, and particularly those services which are required to enable DR of the Production LPAN.

Modelling of these services in Test may be problematical, and some e.g. any providing storage management or network security functions would not be modelled in Test.

10.1.2.1 [ARC] Audit Server

Ref: /DES/APP/HLD/0030 (Gathering)

Ref: /DES/APP/HLD/0029 (Retrieval)

SQL-Server on Windows 2003.

Horizon (and therefore Hydra) has separate gathering at the primary and secondary site to distinct EMC Centera CAS Arrays, with a separate index database at each site on the audit server.

Until all Horizon data has expired, a period of seven years after the final counter migration to ~~HNG-x~~HNG-x plus the time to resolve any outstanding court cases, there is not really an opportunity to redesign the service.

~~10.1.1.12~~10.1.2.2 [SPN] Metron Athene

Formatted: Bullets and Numbering

Ref: /DES/PER/HLD/0022

Performance & Capacity management reporting.

The database is stored on Class 2 storage, and is replicated by the application. Loss of this service would impact on the ability to provide monthly reports.

In general gaps in the performance history are not considered important.

~~10.1.2.2~~10.1.2.3 Support Services

Formatted: Bullets and Numbering

10.1.2.3.1[DXC] Corporate Data Exchange Proxy

Formatted: Heading 5

Formatted: Bullets and Numbering

Ref: /DES/NET/HLD/0018

Active/Active. One per site.

Permits safe transfer of data from corporate networks to the HNG-x network, e.g. software packages released from Dimensions being transferred to the TPM Repository.

~~10.1.2.2.1~~10.1.2.3.2 [SAS] Horizon Secure Access Server

Formatted: Bullets and Numbering

Provides a secure point of entry into the estate for the support of Horizon systems.

Two in each ~~data centre~~Data Centre Support DMZ in active/active state for core support staff.

Resilience model is to use one of the other servers.

Recovery is by re-provisioning.

Users logged into SAS are authenticated with their appropriate role against PWYDCS which is trusted by Hydra systems in BOPSS and WOPSS.

~~*[DN: I am not really sure why we still need this]*~~

~~10.1.2.2.2~~10.1.2.3.3[SSN] ~~HNG-x~~HNG-x Secure Access Server

Formatted: Bullets and Numbering

Ref: /DES/SYM/HLD/0017

Support Access Server. Provides secure point of entry into the estate for support staff.

Two in each ~~data centre~~Data Centre Support DMZ in active/active state for core support staff.

One in each POL DMZ in active/active state for SAP Basis support staff.

Resilience model is to use one of the other servers.

Recovery is by re-provisioning.

~~10.1.2.2.3~~10.1.2.3.4[DNP] & [DNS] BIND

Formatted: Bullets and Numbering

Ref: DES/NET/HLD/0006

Traditional DNS implementation with primary + secondary in IRE11 and IRE19.

Primary active/standby.

Secondary active/active.

Lookup services will be available during failover, but update services will not be available until the primary is failed over.

~~10.1.2.2.4~~10.1.2.3.5[ACD] Active Directory/DNS

Formatted: Bullets and Numbering

Ref: DES/PPS/HLD/0003

This platform also provides the two-factor authentication service as described in DES/SEC/HLD/0001.

It may also provide the Certificate Authority Service [CAN]

FSMO Roles are only active on one server at a time, so during DR basic authentication is available but updates will not be until the FSMO Roles have been transferred.

The interaction between Hydra NT4 Domains and ~~HNG-x~~HNG-x AD Domain is not yet clear. Most of the data transfer is by agent services logging in as Oracle database users which does not require any form of trust.

10.1.2.3.6[NRS] RADIUS servers

Formatted: Highlight

Formatted: Bullets and Numbering

Ref: /DES/PPS/HLD/0011

Ref: /DES/NET/HLD/0014

Ref: DEV/INF/LLD/0077 (CP0184)

RADIATOR RADIUS Service.

CP0184 changes the layout from each service having its own platform to a single platform operating as a load balanced pair at each site, running services BR-ADSL, BR-WWAN, BR-ISDNIN, BR-ISDNOUT

ACE RADIUS probes will be configured to test the authentication of the RADIUS servers within the server farm and control whether a particular RADIUS instance is functioning and hence be made available within the server farm. In the event that no servers are available for a given instance, the VIP will not be advertised through routing on the MSFC and the corresponding data center VIP will service requests.

10.1.2.3.7[NRM] TACACS

Ref: DEV/INF/LLD/0077

CP0184 introduces a resilient authentication service for network management. This is based on two platforms per site, N+1 at each site, on a similar model to DNS or AD.

10.1.2.3.8[IPS] Intrusion Protection System Management Server

Ref:

Two IntruShield 3000 probes are deployed at each data centre for Production use. There is a management platform at each site based in the active LPAN, either one of which may manage the service in a similar model to firewall management.

Formatted: Highlight

Formatted: Bullets and Numbering

Formatted: Highlight

Formatted: Not Highlight

Formatted: Normal

Formatted: Not Highlight

Formatted: Bullets and Numbering

Formatted: Highlight

Formatted: Normal

10.1.2.3.9[NFM]Firewall Manager

Formatted: Bullets and Numbering

Ref:

Hosted on BladeFrame

The firewall manager service is only required in order to update firewall rules. This is a relatively occasional requirement, and should certainly not be required during a DR.

10.1.2.3.10 [SYS] Syslog Server

Formatted: Bullets and Numbering

Ref: /DES/NET/HLD/0012

Active/active. All systems write to both syslog servers. The syslog servers will forward interesting events to Tivoli via NetCool probes, but it is likely that any events will have already generated an SNMP trap via OpenView.

The Branch Routers will also report to the syslog servers, so these are fairly high performance and run a dedicated syslog daemon in addition to the standard linux one that manages the platform.

Interesting events have already been sent to SYSMAN3 by the time the server fails. Opposite site system continues. In the event of prolonged site failure a second system would be provisioned.

~~10.1.2.3~~10.1.2.4 Branch Access - Hydra

Formatted: Bullets and Numbering

10.1.2.4.1 [KMS] Hydra Key Management Server

Formatted: Heading 5

Formatted: Bullets and Numbering

Ref: /RS/MAN/013

This is a SQL-Server database running on Windows NT4 in a special security domain.

The database and any required file store are on the Key Management Server S: drive which is on EMC SRDF replicated storage. BCVs are not used, SQL-Server does a dump to disk and areas of the S: drive are backed up. There is a security concern over the possible theft of this database, and as a result cold backup images are rarely made.

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

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

It is doubtful whether EMC Solutions Enabler (SYMCLI) will work inside a VSH platform, and in any case the version of SYMCLI that was compatible with NT4 is not compatible with current EMC microCode versions on the DMX. The Booter service will probably need to be worked around in order to allow service startup, and the failover procedure will need to be modified to interact with the Storage Management Server (BSS, the Solaris Backup Server).

Resilience: fail over to secondary.

~~10.1.2.3.1~~10.1.2.4.2 [COR] Correspondence Server

Formatted: Bullets and Numbering

Ref: DES/PER/HLD/0003 Branch Trading Resilience HLD.
~~*Ref: [DN: Need a ref for redesigned EACRR]*~~

Ref: SY/SPG/002

Active/active with local N+1 resilience. No failover.

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

The total message store has been manually split into four clusters, each of which contain approximately one quarter of the branch estate (both in terms of size and performance). Split is actually nearer 30%, 30%, 20%, 20%

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

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

The initial build will be from a .VHD file captured from the Horizon system. A rebuild will start by re-applying the .VHD followed by re-provisioning of any fixes.

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

~~10.1.2.3.2~~10.1.2.4.3 [AGE] Generic Agent

Formatted: Bullets and Numbering

~~Ref: [DN: Need a ref for redesigned EACRR]~~

Ref: DES/PER/HLD/0003 Branch Trading Resilience HLD.

Formatted: Font: Arial

The generic agent servers~~s~~ (which are generally known as agents) run services (confusingly also known as agents) which allow messages to be passed between Riposte and the back-end databases.

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

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

Stateless. No complex failover just restart.

Recovery by reprovisioning.

[DN: General note for Hydra systems - still need a reference for the provisioning process i.e. rebuild VSH, drop in original .VHD files and then apply ~~HNG-x~~HNG-x (Hydra) fixes]

~~10.1.2.3.3~~10.1.2.4.4[NRA] Network Banking Routing Agent

Formatted: Bullets and Numbering

Ref: NB/HLD/017

The **NBX Routing Agent** listens for [R1] and [C0] messages through a Riposte real-time message port, add a time-stamp and routes each message to the appropriate NBX Authorisation Agent. There is one Routing Agent instance for each Correspondence Server Cluster. The use of a Riposte real-time message port ensures that the Agent will only process 'fresh' messages.

Note that each physical NRA platform runs four instances of the Routing Agents. These run as active/standby pairs exchanging heartbeats via Riposte.

Resilience is unchanged from Horizon.

10.1.2.4.5[ACF] Hydra ACDB

Formatted: Bullets and Numbering

Ref: /DES/SYM/IFS/0001 (Connect DSL Interface)

Ref: TD/DES/150 (ACF Replay Design)

Hydra Auto-config database.

ACF Replay is the process of re-issuing all ACF to all counters, e.g. if there is a key compromise. The Hydra system (ACDB, ACS and SYSMAN2) needs to be capable of an ACF Replay until all counters are migrated to HNG-x.

[DN: May change to "KMS" model]

Formatted: Font: Italic, Highlight

Formatted: Font: Italic

10.1.2.4.6[OCM] Hydra OCMS

Formatted: Bullets and Numbering

Hydra Outlet Change Management System.

[DN: May change to "KMS" model]

~~10.1.2.3.4~~10.1.2.4.7 [BLS] Horizon Boot Loader

Formatted: Bullets and Numbering

This provides an initial point of contact for a replacement counter to make contact with the ~~data centre~~Data Centre and download its identity (the Auto Config File).

Although this does not sound like it has a very high availability requirement, because of the sheer number of counters in the estate the swap-out rate is relatively high, and the engineers have an SLA to replace the counter within 20 minutes of arriving at the branch.

Delays due to BLS being unavailable have a detrimental knock-on effect in scheduling of engineers for servicing other branches.

~~10.1.2.3.5~~10.1.2.4.8[BOO] VSAT Boot Server

Formatted: Bullets and Numbering

Boot Server acts as a domain controller for the Boot Loader. It also provides a boot loader function for satellite connected branches.

~~10.1.2.3.7 [VPM] VPN Policy Manager~~

Formatted: Bullets and Numbering

~~*[DN: need some words]*~~

~~Active/active. No failover.~~

~~Recovery by reprovisioning.~~

~~10.1.2.3.8 [VEX] VPN Exception Server~~

~~Active/active. No failover.~~

~~Recovery by reprovisioning.~~

Formatted: Bullets and Numbering

~~10.1.2.3.9~~10.1.2.4.9[DOM] Domain Controllers

Formatted: Bullets and Numbering

Ref: [DN: Need a ref for redesigned Hydra NT Domains]

Formatted: Highlight

The Hydra security domain is relatively complex. It is made up of the traditional NT4 PDC/BDC pairs, in some cases with a second BDC for extra resilience.

Many of the servers in DMZ were their own PDC, which has caused some confusion when selecting those that need migrating.

The design for Hydra Security Domains is still not clear, but as a minimum it is expected that the following will be required:

BOPSS - Bootle servers

WOPSS - Wigan servers

PWY~~KMA~~KMS - ~~KMA~~KMS servers and admin users

PWYSAS - SAS Administrators

PWYDCS - Support Users

~~Needs a security designer to comment and also to take ownership of the Hydra NT Domain HLD.~~

Note: technically SDC is a Hydra Domain Controller, but it is managed by a different team.

~~*[DN: PWYDCS just contains the support users, and could be possibly be done away with if BOPSS and WOPSS trusted AD. If I understood why this was difficult I might understand why we still need the Horizon SAS!]*~~

10.1.3 Services in Production LPAN

There will actually be three Production LPAN's; one on each of the three primary site frames, plus additional LPANs on each active/active frame.

Each LPAN will be allocated the necessary resources for the pServers that are contained within that LPAN. It is possible to operate several LPANs on a single frame, e.g. to limit each LPAN to a particular type of resource and prevent LPAN administrators from accidentally assigning an inappropriate resource.

These LPANs will be created on both the primary site PAN and the secondary site PAN, and in the event of DR, following the network and storage failover, plus reassignment of any pServers from Test LPANs the secondary Production LPANs may be started.

10.1.3.1 Support Services

These are support services which are not required to implement DR and can support a Production/Test failover model.

10.1.3.1.1[SSC] SSC Support Server

Ref:

Runs in parallel with the Audit system collecting evidence for problem analysis by SSC. Typically users will only report a problem sometime after an end of month report, so the evidence that led to the query needs to reside in the estate for several months, and the main processing systems do not have space to store data for this long.

This system also gets round any access or control issues of allowing SSC to gather certain evidence or make audit queries to retrieve the same data.

In Horizon this operated active/active and the two systems were lazy mirrored using Robocopy. At ~~HNG-x~~HNG-x this system must fail over and data replication is provided by MirrorView, so a backup must be taken for corruption recovery..

10.1.3.2 Systems Management

Ref: /ARC/SYM/ARC/0001

Ref: /ARC/SYM/ARC/0002

Ref: /ARC/SYM/ARC/0003

Ref: /DES/SYM/HLD/0034 - SYSMAN3 Backup, Availability & Disaster Recovery Design

All the systems management services reside in the primary Production BladeFrames. They may be deliberately spread amongst frames to distribute load or reduce sensitivity to power module failure.

~~[DN: Where does OEM live? ~~Oracle Enterprise Manager is raising Oracle events to replace BMC Patrol functionality, and is also hosting the RMAN Catalog Service. It will be hosted on EDS.]

Formatted: Font: Not Italic

Formatted: Font: Not Italic

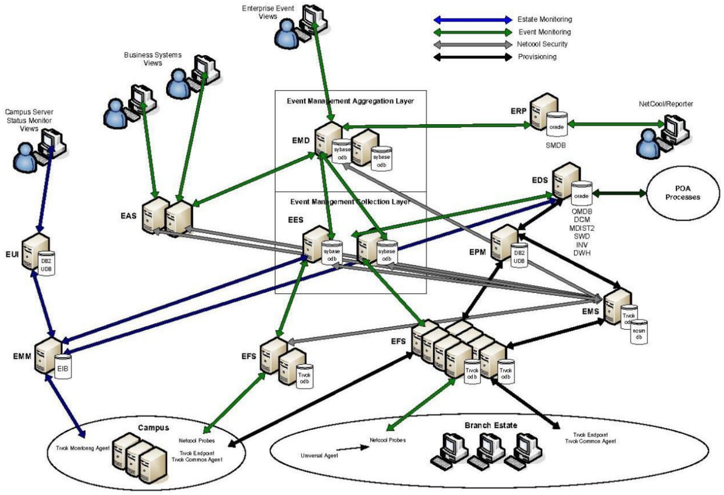

10.1.3.2.1 [EDS] Enterprise Database Server

Oracle database on W2003.

Standard Oracle database failover model. Database should perform automatic crash recovery after failover.

This provides the repository for Tivoli Provisioning Manager and Tivoli Configuration Manager.

10.1.3.2.2[EPM] Enterprise Provisioning Server

Software distribution management. Also provides the Provisioning functions, Branch Router management and Tasks for Campus support

10.1.3.2.3[EFS] Enterprise Fan-out Server

The EFS will provide the Event Concentrator. It also allows performance scaling for distribution to a huge estate. 40 EFS each handle up to approximately 500 counters.

10.1.3.2.4 [EMD] Enterprise Monitoring Display

Top level monitoring server, forming the "Aggregation Layer" of the Event Management Environment.

All events forwarded from clients across the Campus and Branch Estate, subsequently processed and forwarded from the EES Collection Layer servers, are made available here for view and action by the SMC (and/or Automation).

Configured with Tivoli NetCool/OMNIbus (NCO) all Events are processed and stored within its proprietary Object Server Database. This is an in-memory Sybase Database, a regular dump of which is taken via internal Netcool automation processes.

Events are written out to the ERP, to the relevant Oracle Database for access with the relevant Reporter Toolset.

10.1.3.2.5[EUI] Enterprise User Interface

This provides a portal service for real time views of monitored status of configured Servers.

10.1.3.2.6[EMM] Enterprise Monitoring Server

A component of IBM Tivoli Monitoring (ITM), the EMM fulfils the role of Tivoli Enterprise Monitoring Server (TEMS). Campus distributed ITM Agents route status and “situation” alerts into this server. This information is stored in a proprietary database called the EIB. From here, this information can be stored for Historical viewing by Tivoli DataWarehouse which is housed within the EDS. Data is also passed to the EUI Servers for RealTime view of Alerts by support groups.

10.1.3.2.7[EMS] Enterprise Management Server

Top Level Management Server, similar in function (and including some of the same components) as the MASTERTMR within SYSMAN2

2 Servers required, 1 active – 1 standby in separate BladeFrame, to provide N +1 resilience.

A Secondary instance of the Netcool Security Manager Database is required on the Secondary server, and this instance is automatically synchronised with the Primary instance.

10.1.3.2.8[EES] Enterprise Event Server

Second level servers, forming the “Collection Layer” of the Event Management Environment. All events from clients across the Campus and Branch Estate (Server and Counter Log Messages) are processed at these servers, after forwarding through the EFS NetCool Proxies.

Configured with Tivoli NetCool/OMNIbus (NCO) all Events are processed and stored within its proprietary Object Server Database. This is an in memory Sybase Database, a regular dump of which is taken via internal Netcool automation processes.

Audit level events are written out to the EDS, to the relevant Oracle Database. Action level events to be viewed at the Monitoring level are forwarded on to the “Aggregation Layer” NetCool/OMNIbus Object Server, the EMD.

10.1.3.2.9[EAS] Enterprise Availability Server

The~~wo~~ EAS Server~~s~~ will provide a Business Systems View of Events routed from the “Aggregation Layer” of the Event Monitoring structure (EMD) and ~~are p~~presented for view with Tivoli NetCool Realtime Active Dashboards (RAD).

10.1.3.2.10 [ERP] Enterprise Reporting Platform

Event analysis and statistics. Secondary Oracle Database Server, similar but smaller than the EDS.

Netcool/Reporter is installed and used for customised in-house reporting from the relevant Database schemas.

10.1.3.3 Estate Management

Ref: /ARC/SYM/ARC/0005

Ref: /DES/SYM/HLD/0030

~~*[DN: Design not yet complete]*~~

10.1.3.3.1[EST] Estate Management Server

~~What happened to BCDB ???~~Ref: DES/SYM/HLD/0039

The Estate Management System supports two databases:

EMBD the Estate Management Database is responsible for tracking the opening and closing of branches, address changes, etc. It is used by the OBC team.

The database is SQL-Server 2005 and replicates to a standby, much as ACDB and OCMS replicate today. This ensures availability in the event of a problem with the primary server.

Formatted: Highlight

The RADIUS servers use EST during authentication but the RADIATOR cache will allow previous authentications in much the same way as a domain login works on a disconnected laptop.

MTAS MID/TID Allocation Service. SQL-Server 2005. Hosted on DCSM at Horizon as this happened to be in the Network Banking DMZ. Provides input to Authorisation agents, and sends files to Streamline as MID and TID are allocated.

Does not have a particularly high availability requirement as rate of change is low, but as the banks simply drop a Request with an invalid MID or TID faults in MTAS (or at the Streamline end) can be tricky to track down.

[DN: Is EMDB crucial during counter swap-out or is this all handled by BPL and Tivoli?]

Formatted: Highlight

Formatted: Font: Italic

10.1.3.3.2 [BCS] Branch Change Management Server

Ref: /DES/SYM/HLD/0024 (autoconfig)

Ref: /DES/SYM/HLD/0026 (BCMS)

Ref: /DES/SYM/HLD/0031 (BCDB)

Ref: /DES/SYM/IFS/0002 (BCDB to BranchDB)

Branch Configuration Management Service replacing OCMS. This will be based in the corporate estate in Bracknell with a counterpart at the failover site.

[DN: Currently Lewes, but this may change]

Formatted: Font: Italic

~~10.1.3.3.3 [ACF] Hydra ACDB~~

Formatted: Bullets and Numbering

~~Ref: /DES/SYM/IFS/0001 (Connect DSL Interface)~~

~~Ref: TD/DES/150 (ACF Replay Design)~~

~~Hydra Auto-config database.~~

~~ACF Replay is the process of re-issuing all ACF to all counters, e.g. if there is a key compromise. The Hydra system (ACDB, ACS and SYSMAN2) needs to be capable of an ACF Replay until all counters are migrated to HNG-x.~~

~~10.1.3.3.4[OCM] Hydra OCMS~~

Formatted: Bullets and Numbering

~~Hydra Outlet Change Management System~~

~~10.1.3.3.5MTAS Mid Tid Allocation Service~~

Formatted: Bullets and Numbering

~~Ref: May be subsumed into EST_ not initially at least~~

~~MID/TID Allocation Service. SQL-Server (version? 2000). Hosted on DCSM at Horizon as this happened to be in the Network Banking DMZ. Provides input to Authorisation agents, and sends files to Streamline as MID and TID are allocated.~~

~~Does not have a particularly high availability requirement as rate of change is low, but as the banks simply drop a Request with an invalid MID or TID faults in MTAS (or at the Streamline end) can be tricky to track down.~~

~~10.1.3.3.6~~10.1.3.3.3[ACS] Auto Config Signing Server

Formatted: Bullets and Numbering

HYDRA Service based on VSH platform.

Counter Package Signing Server

~~10.1.3.3.7~~10.1.3.3.4[DSS] Dimensions Signing Server

Formatted: Bullets and Numbering

HYDRA Service based in BRA01 and LEW02

Formatted: Normal

Packages from Dimensions are signed before delivery to the estate. Tivoli (SYSMAN3) is able to verify the signature before using a package.

This system is outside the managed data centre service.

10.1.3.4 Branch Access - ~~HNG-x~~HNG-x

[DN: Is this list complete?]

10.1.3.4.1[BAL] Branch Access Layer

Ref: /ARC/APP/ARC/0004

Branch Access Layer based on ~~Interstage~~ grizzly from java.net.

Formatted: Strikethrough

There are 10 instances of BAL at the primary site. This number is to cope with a peak day transaction load, and the service is designed to degrade gracefully, so in practice as few as four or five servers will cope with normal day-to-day traffic.

Peak time is Monday and Tuesday mornings, with extra load just before major public holidays, although the change from OBCS vouchers to CAPO card withdrawals has smoothed this somewhat as a "double payment" is only a single banking transaction whereas two vouchers had to be encashed.

BAL also performs the authentication management for counter connections, and ensures that reconnecting counters are directed to the BDB node that they previously connected to.

In the event of a BDB node failure BAL redistributes the reconnecting counters amongst the remaining nodes. Approximately 4000 branches (average 8000 counters) will be reconnecting in this event.

10.1.3.4.2[BMX] BAL Management Server

Ref: /DES/SYM/HLD/0021 ???

Server with toolset to allow management of the services on BAL.

This server also collects statistics from the BAL platforms and processes them to provide SLA reporting information.

[DN: To where does it send them and how resilient is the mechanism? Are these "on-line" stats or historical?]

[DN: As Interstage is no longer on BAL do we need this platform any longer?]

Formatted: Highlight

10.1.3.4.3[BPL] ~~HNG-x~~HNG-x Boot Platform

Ref:

Strictly this is part of Estate Management.

This loads the ~~HNG-x~~HNG-x equivalent of the auto-config file to an ~~HNG-x~~HNG-x counter spare.

As for its Horizon equivalent, although this does not sound like it has a very high availability requirement, because of the sheer number of counters in the estate the swap-out rate is relatively high, and the engineers have an SLA to replace the counter within 20 minutes of arriving at the branch.

This is also responsible for the RCF for Branch Router Provisioning

~~*[DN: Does it have any other purpose?]*~~

~~*[DN: Need to see Estate Management Design. This is EM (counter provisioning) rather than Branch Access.]*~~

~~10.1.3.4.4 RADIUS servers~~

Formatted: Bullets and Numbering

~~Ref: /DES/PPS/HLD/0011~~

~~Ref: /DES/NET/HLD/0014~~

~~*[DN: Not sure whether the HLD covers all RADIUS servers or whether other HLD need to be referenced.]*~~

~~The basic model for RADIUS servers is an active/active pair. All except RAU are deployed in the Active LPAN. This does not provide N+1 resilience in the event of losing either site.~~

~~*[DN: Do we have enough kit to quickly deploy a second instance? What is the actual availability requirement and the effect of one being offline for an hour while it is rebuilt?]*~~

~~The reason for having these deployed active/active is that in the event of loss of a site (especially IRE11) the counters will immediately try to reconnect, and without a RADIUS server available will continue to retry. This puts a high load on the C&W network.~~

~~*[DN: What are they actually connecting to?]*~~

~~*10.1.3.4.4.1ADR ADSL RADIUS Server*~~

Formatted: Bullets and Numbering

~~*10.1.3.4.4.2GPR GPRS RADIUS Server*~~

~~*10.1.3.4.4.3RAD RADIUS Accounting Server*~~

~~*10.1.3.4.4.4RDD RADIUS Dialled (ISDN) Server*~~

~~*10.1.3.4.4.5ADT ADSL Test Server*~~

~~There is a suggestion that this could be managed some other way. Its function is to provide a file share that can be copied over the ADSL line to check the connection speed before installing a counter.~~

~~*[DN: This really belongs in the "Estate Management" section as it is part of branch provisioning.]*~~

~~10.1.1.1.8[RAU] RADIUS Accounting~~

Formatted: Bullets and Numbering

~~RADIUS Accounting Server.~~

~~This provides authentication for access to network switches for management support. It has been left as a discrete system in case BladeFrame connectivity problems prevent network support staff from authenticating to resolve the problem.~~

~~There will be one per site, active/active. In the event of long term loss of either site a second system would be procured to maintain N+1 resilience.~~

~~10.1.1.1.9[RDO] RADIUS Dial Out~~

Formatted: Bullets and Numbering

~~Ref: /DES/NET/HLD/0005 ???~~

~~Discrete, one per site active/active.~~

~~This system requires a PCI card to enable Tivoli to dial out to unresponsive ISDN counters. The call is then dropped, but the counter has been prompted to dial back.~~

~~It was thought that all ISDN connections would have been migrated to ADSL by the time of HNG-x, but this looks unlikely, and this service is required until the last ISDN counter is replaced.~~

~~10.1.1.1.12[NFM]Firewall Manager~~

Formatted: Bullets and Numbering

~~Ref:~~

~~Hosted on BladeFrame~~

~~The firewall manager service is only required in order to update firewall rules. This is a relatively occasional requirement, and should certainly not be required during a DR.~~

~~10.1.1.1.15[SYS] Syslog Server~~

Formatted: Bullets and Numbering

~~Ref: /DES/NET/HLD/0012~~

~~Active/active. All systems write to both syslog servers. The syslog servers may forward events to Tivoli via NetCool probes, but it is likely that any events will have already generated an snmp trap via OpenView.~~

~~The Branch Routers will also report to the syslog servers, so these are fairly high performance.~~

10.1.3.5 [BDB] Branch Database

Ref: /DES/APP/HLD/0020

BranchDB is a 4 node Oracle10gR2 RAC Database. The nodes are hosted in the BladeFrame

The mechanism for ensuring persistent connections and load balancing and for managing failed nodes is described in detail in the Branch Database HLD (APP/ARC/HLD/0020), with high level context in the Online Services Architecture (ARC/APP/ARC/0008) and the Branch Database Architecture (ARC/APP/ARC/0005)

As well as being highly resilient and providing a very high transactional throughput, the data partitioning scheme allow for more nodes to be added for scalability. In fact the data is typically partitioned 128 ways, so a corrupt table will only affect a small percentage of outlets, and there is every chance that Oracle Recovery Manager can repair the problem.

10.1.3.6 [BDS] Branch Standby Database

Ref: /DES/APP/HLD/0020

This is a second copy on a separate EMC storage array. The replication is using the Oracle DataGuard mechanism, which extracts changes from the transaction logs and applies them to the standby. This means that data corruption is unlikely to be transmitted.

[DN: Current thinking is that this is node 5 of the Branch Cluster (and we may need node 6) which operates two services (databases), the primary and the DataGuard standby. The structure of this section changes very slightly if this is the case]

10.1.3.7 [BRS] Branch Support Database

Ref: /DES/APP/HLD/0023

Branch Support provides two primary services. The first is SLA reporting via the TWS batch schedule, and the other is a historical view of data for SSC to analyse support problems.

Oracle Streams replication keeps BRSS within a few minutes of real time normally, but some catch-up may be required after local DataGuard failover or site failover.

BRSS is regarded as a lower priority service for recovery than BRDB, BRSTBY, BAL and other counter-facing services such as ~~NWB~~NBS, APOP, etc.

10.1.3.8 Network Banking

10.1.3.8.1[NPS] Network Persistent Store

Ref: /DES/APP/HLD/0017

Network Persistent Store. This provides a stateful store for the Banking Authorisation Agents, Debit Card Authorisation Agents and e-TopUp Authorisation Agents, and also a mechanism for the agents to heart beat.

NPS also supports the Track & Trace agent.

The NPS database is a highly available store for transaction journals created by the authorisation agents. The agents themselves are stateless to simplify the agent resilience model, and the NPS allows such features as transaction reversal to be managed by any agent.

The basic availability requirement is to recover within the time it takes a card to be re-swiped. Failed (unacknowledged) transactions time out after 30s, and typical recovery following node failure is <60s.

The database is hosted on RHEL based two-node Oracle RAC in the BladeFrame. The platform also co-hosts the APOP database.

Each Authorisation agent connects to both RAC instances upon start-up for critical processing threads. This places a high memory requirement on the server but reduces the failover time for the agent should a node crash.

The use of Oracle Recovery Manager (RMAN) for backup and recovery allows an even more timely recovery from corruption than at Horizon, but in practice this has not been an issue in service.

10.1.3.8.2[NAC] CAPO Authorisation Agent

Ref: /DES/APP/HLD/0009

Authorisation agent for Card Account.

Two agents run, an A and a B instance, to provide performance. There is one instance of each on hot standby, heart beating via NPS, thus a total of four NAC platforms in the ~~data centre~~Data Centre.

CAPO accounts for 85% of transaction volume, and outages are highly news-worthy.

Stateless. No complex failover just restart.

After DR a successful key exchange event indicates that the service is communicating with the FI.

10.1.3.8.3[NAA] A&L Authorisation Agent

Ref: /DES/APP/HLD/0009

Authorisation agent for Alliance & Leicester, the only bank to sign up for an individual service with POL.

Same resilience model as CAPO, but a single platform is able to host both A and B instances. Normally one platform will host the active A instance and the other will host the active B instance.

A&L accounts for <5% of transaction volume.

Stateless. No complex failover just restart.

After DR a successful key exchange event indicates that the service is communicating with the FI.

10.1.3.8.4[NAL] LINK Authorisation Agent

Ref: /DES/APP/HLD/0009

Authorisation agent for Link, the network covering most other major banks.

Same resilience model as CAPO, but a single platform is able to host both A and B instances. . Normally one platform will host the active A instance and the other will host the active B instance.

Link accounts for around 10% of transaction volume.

Stateless. No complex failover just restart.

After DR a successful key exchange event indicates that the service is communicating with the FI.

Note that whilst NAC and NAA initiate connections to the FI, NAL accepts. If there is a long outage, such as may occur during a DR, Link will need to be contacted to re-establish connection.

10.1.3.8.5[TWS] TES & APOP Query

Ref: /DES/APP/SPE/0001 (can't see an HLD)

Ref: ???? (APOP)

Oracle forms server to allow POL (Huthwaite & Chesterfield) users to query APOP and TES databases in a controlled manner.

There is also statement about availability:

> The Online Query Tool will be available under normal operation on a 24 x 7 basis, bar Daily maintenance for the On line Query tool may be taken overnight for up to 30 minutes during a possible 4 hour daily Maintenance. Fujitsu Services will advise on the start and finish times of the Daily Maintenance window so this can be communicated to the users of the TES (see NBX0116e)
>
> The On-line Query Tool will have an availability of 99.75% during 0700 to 22:00 measured annually and reported monthly on an exception basis.

Formatted: Font color: Auto

Formatted: Justified, Space After: 6 pt, Adjust space between Latin and Asian text, Adjust space between Asian text and numbers

Formatted: Font color: Auto

Formatted: Justified, Indent: Left: 1.5 cm, Space After: 6 pt, Adjust space between Latin and Asian text, Adjust space between Asian text and numbers

~~Oracle forms server to allow POL (Huthwaite & Chesterfield) users to query TES in a controlled manner.~~

~~Stateless. No complex failover just restart.~~

The APOP service is an Oracle forms server co-hosted with TESQA. The interface a~~A~~llows query and update of APOP for releasing and managing vouchers (e.g. Postal Orders). Each stream of vouchers has a different set of users.

Stateless. No complex failover just restart.

10.1.3.9 Other online services

10.1.3.9.1APOP Automated Payments and Out-Payments

Ref: /DES/APP/HLD/0011

Co-hosted on [NPS] Platform

Data repository for Automated Payments Out Payments service, which tracks vouchers such as Postal Orders to check for lost, stolen and forged vouchers, and to provide a position on the outstanding cash that POL has issued as vouchers but not redeemed.

Also provides stock tracking of vouchers, and a report on unredeemed (out of date) vouchers.

Primary interfaces are a bulk upload of voucher details as they are issued and an update by Post Masters as vouchers are transacted through [AWS].

Headquarters can also query and update records via a service on [TWS].

10.1.3.9.2 [KMN] Key Management Service

Ref: /DES/SEC/HLD/0003

[DN: Details to be confirmed]

Formatted: Font: Italic, Highlight

Formatted: Font: Italic

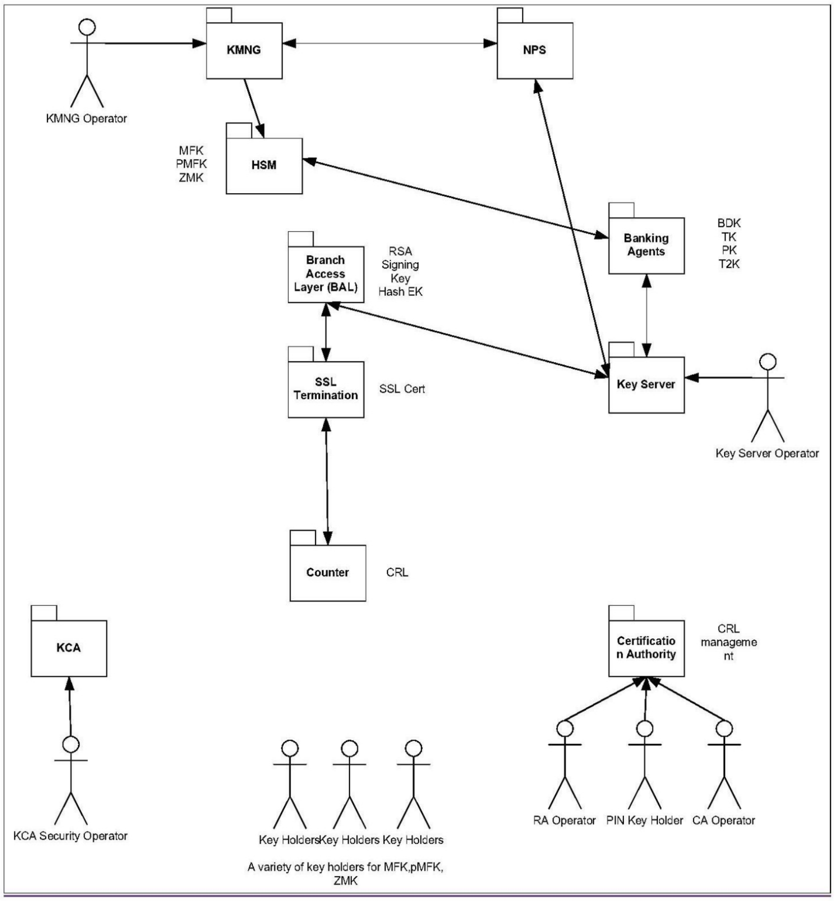

It is understood that the Key Server will operate as a load balanced pair with the NPS as a repository, which is analogous to the model used for other web services.

It is not clear whether the KMN platform is required during the DR process and therefore needs to be in the ACTIVE LPAN, but it is critical in allowing other services to start, and must appear early in the service start-up ordering.

10.1.3.9.3[DCA] Horizon Debit Card Agent

Prior to the completion of PCI Compliance Horizon counters will continue to use the Horizon DCS Agent, but re-hosted in Belfast.

~~Horizon agents will require crypto key disks during reboots, so in the event of the IRE19 agent failing approved staff will need to attend site to insert the key disk, and there will be reduced resilience for this period.~~

~~This is the only system in Belfast that will require physical key disks *[Ref: Dave Johns - is this still the case? - see also section 9.1]*. As stated earlier, I think Generic agents will too.~~

Post PCI Compliance and for ~~HNG-x~~HNG-x counters the ~~HNG-x~~HNG-x DCS Agent will be used.

This platform also hosts the Hydra ETS agent which continues for the duration of the Hydra phase.

Resilience will be the same model as now, with two agents per cluster, one at each site, in Active/Standby mode using Riposte heart beats to decide which is Active (a total of eight platform instances).

Stateless. No complex failover just restart.

10.1.3.9.4[DEA] Debit Card & eTop-Up Agent

Ref: /DES/APP/HLD/0007 - DCS

Ref: /DES/APP/HLD/0008 - ETS

Provides an authorisation service for debit card payments via an X25 link to the Merchant Acquirer NatWest Streamline, and a similar authorisation for eTop-Ups.

N+1 local resilience is achieved through a pair of agents running active/standby, heartbeating via NPS.

10.1.3.9.5 [DCM] Debit Card Management Server

Ref: /DES/APP/HLD/0055 DCS Bulk File Agents ~~???~~

Ref: /DES/APP/HLD/0078 (Streamline)

Ref: /DES/APP/HLD/0077 (eTop Up)

Send and receive settlement files. Input from a DRS share on DAT is converted into a payment file. The response is an EMIS file which is converted into a C4D file for DRS.

~~Actually this is just a format converter for the payment file on a DRS share before sending via ftp to Streamline.~~

With PCI compliance this server de-~~obfuscate~~crypts the PANs in the Payment file before transmission, and temporarily stores the result on encrypted file store for sending via ftp to Streamline. The DCM also encrypts the PANs in the EMIS file before passing the C4D feed to DRS.

Formatted: Not Highlight

In Horizon t~~T~~he bulk of the days transactions are sent at 1500 with a tail at 2000. The previous day's EMIS file is retrieved at 1500 as the payment file is sent.

~~Availability requirements are quite low. As long as the service is available at 1500 and 2000 to send the payment file it could be shut down for the rest of the day.~~ For ~~HNG-X~~HNG-X ~~we've~~this is reduced this to just the 21:00 payment file and the EMIS file is received when Streamline send it.

The availability requirement is rather low

Formatted: Font: Italic

10.1.3.9.6Web Services

Ref: /DES/APP/HLD/0010 (Generic)

All web services operate as a load-balanced pair. In the event of one failing the other is used automatically.

All of these systems are stateless and may be rebuilt if they fail. The existence of the partner system provides a continued service while the rebuild occurs.

10.1.3.9.6.1 [MWS] MoneyGram Web Server

Ref: /DES/APP/HLD/0014

Requires KMN to start up.

Authorisation of "international postal orders".

10.1.3.9.6.2 [PWS] PAF Web Server

Ref: /DES/APP/HLD/0015

Post Code look-up application. This has an internal "database" updated by controlled software release.

10.1.3.9.6.3 [DWS] DVLA Web Server

Ref: /DES/APP/HLD/0012

Interstage with link to external service for looking up vehicle license details.

Stateless. No complex failover just restart.

10.1.3.9.6.4 [OWS] Online Training Web Server

Ref: /DES/APP/HLD/0031

Provides a dummy service for training counters

10.1.3.9.6.5 [HWS] Help Desk Web Server

Ref: /DES/APP/HLD/0013

This service allows logging of problems with the help desk via a web service rather than via a phone call.

10.1.3.9.6.6 [AWS] APOP Web Server

Ref: /DES/APP/HLD/0011

This service provides authorisation for redeeming vouchers.

It also allows counters to update APOP voucher records as vouchers are sold and redeemed.

See also [TWS] which allows Headquarters to query and update APOP. There are two different "APOP Web Services" which is confusing.

10.1.3.9.6.7 [BWS] Telecoms Web Server

Formatted: Heading 6

Formatted: Bullets and Numbering

Ref: /DES/APP/LLD/0158

Ref: SVM/SDM/OLA/0002

Formatted: Portuguese (Brazil)

Ref: SVM/SDM/OLA/0003

Formatted: Portuguese (Brazil)

Ref: SVM/SDM/OLA/0004

Ref: SVM/SDM/OLA/0005

This service provides authorisation for selling broadband services at the counter. A number of Operational Level Agreements exist covering the various stages required to check the customer's creditworthiness and the availability of broadband for their address.

Formatted: Normal

10.1.3.10 File Transfer

10.1.3.10.1 [CDG] C:D Connect-Direct server

Ref: /DES/APP/HLD/0095

Ref: /DES/APP/HLD/0052 (PCI Agents)

The CDG decrypts PANs in the outgoing REC files and obfuscates PANs in the incoming LREC files.

Used to transfer files to and from Financial institutions as C:D is an industry standard.

It is similar to FTMS except that it is stateless and does not return receipts.

As with DCM the availability requirement of this platform is ~~relatively low~~zero outside the window when the REC and LREC files are being transmitted. Even within the window there is considerable margin.

[DN: Need to confirm SLA's on REC file delivery and normal delivery time]

10.1.3.10.2 [PLG] FTMS TIP Local

Ref: /DES/APP/HLD/0051 (FTMS)

FTMS was developed for Horizon to provide a system that could deliver files to a service boundary with proof for SLA reconciliation.

Although called TIP, TIP has long since been replaced by POLFS, and the old TIP feed (slightly modified) now goes to POL-MIS.

The TIP gateway is also used by RDS to send files to RDMC and ADS for two-way traffic to LFS and POLFS.

The critical SLA was delivery of the transaction files to TIP, but the equivalent (BLE) files go direct from TPS to POLFS now. There is an SLA on the PLO Early file from ADS being available to the counters by 0700, and failure of the system during the core day interrupts flows of data such as ADS pouch delivery and collection details and Bureau de Change updates and memos to Post Masters fed via the Reference Data system.

10.1.3.10.3 [FLG] FTMS EDG Local

Ref: /DES/APP/HLD/0051 (FTMS)

Ref: /DES/APP/HLD/0079 (EDG FTP Pull)

Ref: /DES/APP/HLD/0080 (EDG FTP Push)

Ref: /DES/APP/HLD/0081 (GIRO FTP Push)

EDG (Electronic Data Gateway) will have replaced individual remote AP clients before migration starts. Principally used for sending files for AP clients to POL for collection by the client from a POL external facing system (similar to Equifax), it is also used for transmitting other files between Northern ~~Data Centre~~Data Centre and Horizon because the TIP gateway hit a Windows limit on the number of services that could be run. POLFS has a number of feeds on this interface.

FTMS (GP) is the remote in Stevenage. One of the main transfers is reference data (outlet address changes) to the D1 Engineering Support System which sends out spare parts. GP may not be needed at ~~HNG-x~~HNG-x.

[DN: D1 is being shut down. Not sure what is replacing it. Graham Welsh to advise of any changes needed.]~~[DN: GIRO feed is as big as all the rest put together - will this also use EDG?] That's the current plan, but as an extra service.~~

~~10.1.3.10.4 [xxx] Track & Trace~~

Formatted: Bullets and Numbering

The T&T EDG Agents also run on the FLG. While there are Horizon Counters working there will be five such agents - one per Riposte cluster (represented in the NPS) and one for HNG-X

Ref: CR/CDE/018

Ref: CS/OLA/057

Ref: AS/DPR/013

Ref: DE/LLD/019 - ~~EGD~~ EDG Interface Agent

Formatted: Dutch (Netherlands)

Ref: DE/LLD/019 - Harvester Agent

Formatted: Dutch (Netherlands)

When a T&T barcode is applied to a package a record is forwarded to a 3rd party (which may be ParcelForce) via the EDG gateway.

The resilience model of the EDG gateway [FLG] assumes that fairly low availability is OK as it is just delivering files to APS customers. ~~???~~

The T&T agent may have higher availability requirements.

Stateless. No complex failover just restart.

10.1.4 Reference Data Test (RDT)

RDT operates as an isolated environment in Horizon. This could either be managed as a DMZ in Production or as an independent LPAN. The DR requirements are fairly relaxed and merely require working systems at a reasonably coherent point in time.

10.1.4.1 [RSH] RDT Solaris Host

Four per site, one for each service PL IV OV T. There is no planned data transfer between sites, as any of these systems should be able to reload from exports on the resilient RDT file share.

Formatted: Normal

10.1.4.2 [RLS] RDT Linux Server

BRDB, NPS and APOP may share a platform. In order to simplify DR these systems have been implemented in a way that allows EMC Clariion MirrorView or SAN Copy to be used. Any such requirement is likely to be once per week as in Horizon (tape transfer to LEW02 from BRA01).

Formatted: Normal

10.1.4.3 RDT Web Services

These systems are stateless, but must be kept up to date with any fixes. Replication of the Clariion disks through SAN Copy is the simplest mechanism.

Formatted: Normal

10.1.5 Services in Test LPANs

There will be one LPAN per test system, each named after the test system as per the naming convention in DES/PPS/HLD/0006 Naming Standard, e.g. ST for System Test. If a Test service is spread across several frames there will be one LPAN on each 'frame, each containing the appropriate set of resources for that frame.

A Test LPAN only has access to the resources which the global PAN Administrator has granted, so even if Testers are given LPAN Operator or LPAN Administrator roles they do not have access to other resources.

It is thus possible to safely trunk a frame to both core and access switch layers, and to present all VLANs to all trunked ports, as the creation of a vSwitch which maps to each VLAN is only permitted for the global PAN Administrator, who must explicitly assign the vSwitch to the LPAN. If a vSwitch has been assigned to several LPANs the LPAN administrator cannot tell that, he can only tell that the vSwitch is assigned to the LPAN being administered.

A user could be given a role of ST-LPAN-Administrator and RV-LPAN-Operator, and would then be able to see a wider range of resources, but still only assign ST resources in the ST LPAN. An Operator is allowed to start and stop pServers and see events, which is appropriate for SMC (2nd line) type users.

A vSwitch may be presented to several LPANs. If there is some infrastructure which needs to be shared between Production and Test systems (a NAS based software share, for instance) this can be managed in a controlled manner.

Each test stream runs in its own LPAN to prevent resource contention. These are generally scaled down copies of the services described above, but for many systems scaling down does not make sense.

There may be a need for the long term Test service to run full scale transactional volumes, at least for a single peak day. TBD.

In the event of DR these services will first be shut down

There is a requirement to operate a minimal test service to enable release of important fixes during DR. It is not yet certain how this will operate, but in principle there is not a problem operating this LPAN concurrently with any Production LPAN(s).

10.1.6 Test only services

Some components will exist in the ~~data centre~~Data Centre to support testing, for example there is a complete set of backup servers at both sites for Test. These are part of the managed service.
Some components of the test rigs will not be hosted in the ~~data centre~~Data Centre. This includes all the test counters, workstations, emulators, crypto devices, etc. which testers need to interact with directly.
The testers normally operate from the BRA01 site, but in the event of this being unavailable operations will continue from a secondary site (currently LEW02). All the required servers and workstations will already provisioned at this site, and the site will be tested for operational effectiveness as part of Business Continuity Plans.

10.1.7 Remote Services

10.1.7.1 Workstations

All remote sites have a DR equivalent. All workstations and network connections will be pre-built and available in the event of a DR. Periodic business continuity testing will exercise this equipment and the processes for accessing the secondary site according to the published yearly schedule of tests.

10.1.7.2 FTMS Remote

FTMS Remote platforms [FRG] and [PRG] operate both local N+1 resilience at the primary site, and have DR equipment located at the customer secondary site.

The DR equipment is exercised in concert with the customer DR testing for example the RMG NDC DR Rehearsal Plan.

The local resilience is tested according the published yearly schedule of tests. This may include the customer testing their own procedures for retrieving files from the secondary platform.

(The Northern ~~Datacentre~~Data Centre (Huthwaite) Disaster Recovery Test is currently taking place, OCP17279 refers.)

10.1.8 POL-FS

POL-FS will continue as under Horizon. This is being proved by the "Pathfinder" project.

Ref: [Need a PID ref from Chris Credland]

SVM/SDM/SD/0003 Annex B describes the service.

EA/DPR/005 describes the physical estate and the SAP landscape.

EA/DES/001 describes the resilience and failover.

EA/DES/002 describes the SAN, but this is superseded by DES/NET/HLD/0007

The IRE19 HNG-x systems are used by Fujitsu Services testers. The POLFS systems are used by PRISM testers (based in Huthwaite) so form part of a managed service offering to an external customer. Delaying an urgent fix to the Production system is the POLFS equivalent of not having LST operational.

Two services (HNG-x and POLFS) to coexist in the same infrastructure. This requirement comes from SVM/SDM/SD/0003 Annex B, and is due to the much longer DR time for POLFS. Scheduling the POLFS DR test to coincide with the HNG-x DC test is beyond the scope of this design.

We do not wish to invoke HNG-x DR simply because the main SAP database server had failed.

10.1.8.1 PLP

Production R3

10.1.8.2 PXI

Production XI

10.1.8.3 PLQ

QATest R3

10.1.8.4 QXI

QATest XI

10.1.8.5 PLE

Volume Test (talks to BTC7 in Horizon)

10.1.8.6 PLD

Development R3 (Prism)

10.1.8.7 DXI

Development XI (Fujitsu Services)

10.1.8.8 PLS

Support R3

Currently suspended as disks have been assigned to PLP. Will be resurrected as part of Pathfinder.

10.1.8.9 PLM

Monitor R3

Currently suspended as disks have been assigned to PLP. Will be resurrected as part of Pathfinder.

10.1.8.10 DSP

Production IXOS (OpenText) Archive. This is a simple Oracle database that is used to index the images stored on Centera.

10.1.8.11 DS

QATest IXOS (OpenText) Archive. Shares storage with Production but using a separate disk area.

11 Security

System fails because of network intrusion or malicious activity:

> This is an indeterminate failure. A damage assessment would need to be performed before permitting failover. This is mitigated by the security design, and the implementation of intrusion detection and response. Failover will only be implemented at the direction of an RMGA Service Manager who will have first sought Customer approval.

Ensure only users with the correct roles can initiate any of the failover steps:

> There is a need to avoid "back doors" to enable support or invocation of DR, and low level designs will be reviewed for this. This risk is then mitigated by normal security and access policies under ID Management.

Users need to be trained

> A programme of introducing support staff to the processes is outlined in Section 2.

Running incorrect scripts could damage systems:

> Any scripts must be fail-safe, and must make comprehensive tests that the system is in the correct state to initiate the action performed by the script. It will be possible to force invocation this should not be the norm and any warning issued should make this clear. This will be stated in the "design principles" section of any low level design document dealing with Disaster Recovery.

Separation of Production and Test domains:

> The external, shared infrastructure is managed by teams who have access and clearance for the Production system, and those management interfaces are essential to timely DR, and will be in the Production domain. Strict working practices and naming conventions will ensure that there is a very low risk of presenting Test systems to the Production estate.

12 Scenarios

The resilience scenarios should all be covered by the SRRC published for each service.

If this level of detail for the overall service is needed it should be in an LLD (equivalent to the aircraft industry Failure Modes, Effects and Criticality Analysis (FMECA) which is used to satisfy the certification authorities as to fitness for purpose - there are useful analogies here).

Disaster Recovery is an indeterminate problem, and a decision making process is more appropriate than a list of causes. The Business Continuity Framework (FS/BUA/SPP/001) is intended to guide the writing of such process documents.

In general one must review the degree to which service has been lost or is about to be lost and balance that against the definite outage that will be caused by a site failover.

~~13List of Platform Types as of 15-NOV-2007~~

Formatted: Bullets and Numbering

~~This is not part of the document but is attached for reference. The final list of platforms is still to emerge, subject to a number of CPs.~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
-		-	- - -	- -			
~~ADR~~		~~Radius ADSL Server~~	~~Radius Server for Authentication of ADSL connected branches.~~	~~Access and Authentication~~	~~Windows 2003 Radius Server~~	~~Migrated~~	~~Yes~~
~~ADT~~		~~ADSL Test Server~~	~~Used to test new ADSL connections.~~	~~Access and Authentication~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~
~~ARC~~	~~Y~~	~~Audit Server~~	~~The main application server for the Audit system.~~	~~Application~~	~~Windows 2003 SQL Server~~	~~Migrated~~	~~Yes~~
~~AWS~~	~~Y~~	~~APOP Web Server~~	~~Runs the APOP Web Service Application~~	~~Application~~	~~RHEL Interstage Server (Virtual)~~	~~Migrated~~	~~Yes~~
~~BLK~~	~~Y~~	~~ITU Test Bulk Loader~~	~~ITU Test Bulk Loader~~	~~Testing~~	-	~~Migrated~~	~~No~~
~~BWS~~	~~Y~~	~~Telecoms Web Server~~	~~Telecoms Web Server~~	~~Application~~	~~RHEL Interstage Server (Virtual)~~	~~Migrated~~	~~Yes~~
~~CDG~~	~~Y~~	~~ConnectDirect Gateway~~	~~Runs the Connect:Direct gateway, used to transfer data to and from financial institutions.~~	~~Application~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~
~~CDS~~	~~Y~~	~~Connect Direct Simulator~~	~~Simulates the Connect:Direct service.~~	~~Testing~~	~~Windows 2003 Server~~	~~Migrated~~	~~No~~

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~CON~~		~~Aurora Console Tower~~	~~Provides secure console access to other servers.~~	~~Access and Authentication~~	~~Solaris~~	~~Migrated~~	~~Yes~~
~~DAT~~	~~Y~~	~~Solaris Host~~	~~Main Unix database server for back-end applications.~~	~~Application~~	~~Solaris Oracle~~	~~Migrated~~	~~Yes~~
~~DCM~~	~~Y~~	~~Debit Card Management Server~~	~~Runs the bulk agents that pass data to and from Streamline for debit card processing, and the MID/TID allocation service (MTAS) database that holds debit card merchant and terminal identifiers.~~	~~Application~~	~~Windows 2003 SQL Server~~	~~Migrated~~	~~Yes~~
~~DVI~~	~~Y~~	~~ITU Test DVLA Injector~~	~~ITU Test DVLA Injector~~	~~Testing~~	-	~~Migrated~~	~~No~~
~~DWS~~	~~Y~~	~~DVLA Web Server~~	~~Runs DVLA web service, as a virtual platform on the EOS physical platform~~	~~Application~~	~~RHEL Interstage Server (Virtual)~~	~~Migrated~~	~~Yes~~
~~ECC~~		~~EMC Control Centre – Main Host~~	~~Controls the EMC storage.~~	~~Storage and Backup~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~
~~FLG~~	~~Y~~	~~FTMS EDG Local~~	~~The local server for the GP/EDG FTMS instance.~~	~~Application~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~
~~FRG~~	~~Y~~	~~FTMS EDG Remote~~	~~Remote EDG FTMS instance. Runs at the Post Office, not in the data centre.~~	~~Application~~	~~Windows 2003 Server~~	~~Migrated~~	~~No~~
~~GPR~~		~~Radius GPRS Server~~	~~Radius Server for authentication of GPRS connected branches.~~	~~Access and Authentication~~	~~Windows 2003 Radius Server~~	~~Migrated~~	~~Yes~~
~~INJ~~	~~Y~~	~~ITU Test Injector~~	~~ITU Test Injector~~	~~Testing~~	-	~~Migrated~~	~~No~~
~~LLX~~	~~Y~~	~~Lexcel Simulator~~	~~Simulates Link, A&L and CAPO online services.~~	~~Testing~~	~~Windows 2003 Server~~	~~Migrated~~	~~No~~
~~MIS~~		~~MIS Client~~	~~Workstation used by CS Management Support Unit for access to Data Warehouse, DRS and TESQ~~	~~Operational Support~~	~~Workstation~~	~~Migrated~~	~~No~~
~~MSW~~		~~MIS Support Workstation~~	~~Workstation used by support, primarily for support of Business Objects within Data Warehouse. Also hosts DRS workstation application.~~	~~Operational Support~~	~~Workstation~~	~~Migrated~~	~~No~~
~~MWS~~	~~Y~~	~~MoneyGram Web Server~~	~~Runs the MoneyGram Web Service Application~~	~~Application~~	~~RHEL Interstage Server (Virtual)~~	~~Migrated~~	~~Yes~~
~~NAA~~	~~Y~~	~~A&L Authorisation Server~~	~~Runs the agents that provide on-line access to Alliance and Leicester.~~	~~Application~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~
~~NAC~~	~~Y~~	~~CAPO Authorisation Server~~	~~Runs agents that provides on-line access to the Post Office Card Account service (CAPO).~~	~~Application~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~
~~NAL~~	~~Y~~	~~Link Authorisation Server~~	~~Runs agents that provides on-line access to the Link network for cash withdrawals.~~	~~Application~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~NPS~~	~~Y~~	~~Network Persistent Store~~	~~Database server shared across Network Banking, Electronic Top-up Service (ETS), Debit Card System (DCS), Track & Trace Agent, and APOP. Generally known as NPS.~~	~~Application~~	~~RHEL Oracle Database Server~~	~~Migrated~~	~~Yes~~
~~PFI~~	~~Y~~	~~ITU Test PAF Injector~~	~~ITU Test PAF Injector~~	~~Testing~~	~~-~~	~~Migrated~~	~~No~~
~~PLG~~	~~Y~~	~~FTMS TIP Local~~	~~The local server for the TIP FTMS instance.~~	~~Application~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~
~~PPT~~		~~PIN Pad Test Workstation~~	~~Used by Triage to test field Pin Pad spares.~~	~~Operational Support~~	~~Workstation~~	~~Migrated~~	~~No~~
~~PRG~~	~~Y~~	~~FTMS TIP Remote~~	~~Remote TIP FTMS instance.~~	~~Application~~	~~Windows 2003 Server~~	~~Migrated~~	~~No~~
~~PWS~~	~~Y~~	~~PAF Web Server~~	~~Runs the PAF web service~~	~~Application~~	~~Windows 2003 Interstage Server~~	~~Migrated~~	~~Yes~~
~~RAD~~		~~Radius Accounting Server~~	~~Radius Server for accounting of connections from dialled branches.~~	~~Access and Authentication~~	~~Windows 2003 Radius Server~~	~~Migrated~~	~~Yes~~
~~RDD~~		~~Radius Dial (Authentication) Server~~	~~Radius Server for authentication of dialled ISDN connected branches.~~	~~Access and Authentication~~	~~Windows 2003 Radius Server~~	~~Migrated~~	~~Yes~~
~~RDM~~	~~Y~~	~~RDMC workstation~~	~~Workstation to access the Reference Data Management Centre. Hosts APS workstation app.~~	~~Operational Support~~	~~Workstation~~	~~Migrated~~	~~No~~
~~SIM~~	~~Y~~	~~ITU Test Counter Simulator~~	~~ITU Test Counter Simulator~~	~~Testing~~	~~-~~	~~Migrated~~	~~No~~
~~SKG~~		~~PIN Pad Key Generation Workstation~~	~~Workstation for generating PIN Pad keys.~~	~~Operational Support~~	~~Workstation~~	~~Migrated~~	~~No~~
~~SSC~~	~~Y~~	~~SSC Server~~	~~Service used for various support tasks.~~	~~Operational Support~~	~~Windows 2003 Server~~	~~Migrated~~	~~Yes~~
~~SSW~~	~~Y~~	~~SSC Support Workstation~~	~~Used by SSC for support tasks.~~	~~Operational Support~~	~~Workstation~~	~~Migrated~~	~~No~~
~~TPP~~		~~Training PIN Pad Loading Workstation~~	~~Workstation where PINpads for Counter Training Offices (different from Live) are loaded with security keys.~~	~~Operational Support~~	~~Workstation~~	~~Migrated~~	~~No~~
~~TVW~~		~~Tivoli Workstation~~	~~Workstation used to manage the Tivoli environment.~~	~~Systems Management~~	~~Workstation~~	~~Migrated~~	~~No~~
~~TWS~~	~~Y~~	~~TES Web Server~~	~~Oracle application server that runs query application for the Transaction Enquiry System (TES) and some parts of the APOP application.~~	~~Application~~	~~Windows 2003 Oracle Application Server~~	~~Migrated~~	~~Yes~~
~~ACD~~	~~Y~~	~~Domain Controllers - Active Directory~~	~~Active Directory for HNG-X.~~	~~Access and Authentication~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~AUW~~	~~Y~~	~~Audit Workstation~~	~~Workstation to access the Audit system. Also connected to the HSM.~~	~~Application~~	~~Windows XP~~	~~New~~	~~No~~
~~AVS~~		~~Antivirus Server~~	~~-~~	~~Security Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~BAL~~	~~Y~~	~~BAL Server~~	~~The Branch Access Layer servers.~~	~~Application~~	~~RHEL Interstage Server~~	~~New~~	~~Yes~~
~~BCS~~	~~Y~~	~~Branch Change Management System Server~~	~~Server for the Branch Change Management System (BCMS), the HNG-X successor to OCMS~~	~~Estate Management~~	~~Windows 2003 SQL Server~~	~~New~~	~~Yes~~
~~BDB~~	~~Y~~	~~Branch Database Server - Main~~	~~The main branch database server.~~	~~Application~~	~~RHEL Oracle Database Server~~	~~New~~	~~Yes~~
~~BDS~~	~~Y~~	~~Branch Database Server - Standby~~	~~The standby branch database server.~~	~~Application~~	~~RHEL Oracle Database Server~~	~~New~~	~~Yes~~
~~BMX~~	~~Y~~	~~BAL Management Collation Server~~	~~Collects and monitors metrics including SLA Metrics exposed by the BAL~~	~~Operational Support~~	~~RHEL~~	~~New~~	~~Yes~~
~~BPL~~		~~Boot Platform (HNG-X)~~	~~HNG-X Boot platform~~	~~Estate Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~BRS~~		~~Branch Support Server~~	~~Historical data copied from the branch database server, used for support. Known as BRSS.~~	~~Application~~	~~RHEL Oracle Database Server~~	~~New~~	~~Yes~~
~~BSL~~	~~Y~~	~~RHEL Backup Server~~	~~RHEL NetBackup media server & master server.~~	~~Storage and Backup~~	~~RHEL~~	~~New~~	~~Yes~~
~~BSS~~	~~Y~~	~~Solaris Backup Server~~	~~Server for backing up Solaris servers. HNG-X version.~~	~~Storage and Backup~~	~~Solaris~~	~~New~~	~~Yes~~
~~BSW~~	~~Y~~	~~Windows Backup Server~~	~~Controls the backup of Windows servers under HNG-X.~~	~~Storage and Backup~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~CAN~~	~~Y~~	~~Certificate Authority Server~~	~~Creates certificates for HNG-X systems.~~	~~Security Management~~	~~Other~~	~~New~~	~~No~~
~~CNT~~	~~Y~~	~~HNG-X NT4 Counter~~	~~The new counter system for HNG-X, NT4 version.~~	~~Application~~	~~Counter PC~~	~~New~~	~~No~~
~~CTS~~		~~Certificate Server~~	~~A placeholder for whatever certificate server HNG-X requires.~~	~~Security Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~CXP~~	~~Y~~	~~HNG-X XP Counter~~	~~The new counter system for HNG-X, XP version.~~	~~Application~~	~~Counter PC~~	~~New~~	~~No~~
~~DEA~~	~~Y~~	~~DCS and ETS Authorisation Server~~	~~New HNG-X platform from DCS and ETU Authorisation Agents~~	~~Application~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~DNP~~		~~DNS Server (primary)~~	~~The primary domain name server.~~	~~Network Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~DNS~~		~~DNS Server (secondary)~~	~~The secondary domain name server.~~	~~Network Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~DSS~~		~~Dimensions Signing Server~~	~~Server for signing packages to be sent to the HNG-X Counter.~~	~~Software Distribution~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~DXC~~		~~Corporate Data Exchange Proxy~~	~~Server that provides a proxy for exchange of data between Data Centre and Fujitsu Corporate systems.~~	~~Network Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~DXI~~		~~Internet Data Exchange Proxy~~	~~Server that provides a proxy for exchange of data between Data Centre and External Internet systems.~~	~~Network Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~EAS~~		~~Sysman Availability Server~~	~~SYSMAN3 EAS provides the Realtime Active Dashboard (RAD) layer Server functionality~~ ~~Omnibus gateway (Incident integration)~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~EBS~~		~~Enterprise Boot Server~~	~~Server to help bootstrap of Linux, Windows and Solaris servers.~~	~~Systems Management~~	~~Solaris~~	~~New~~	~~Yes~~
~~ECA~~		~~EMC Control Centre – Agents Hosts~~	~~Controls the EMC storage.~~	~~Storage and Backup~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~EDS~~		~~Sysman Enterprise Database Server~~	~~SYSMAN3 Database repository service for TPM and TCM – supports Inventory , asset and software distribution and event audit service (replaces the SYSMAn2 OMDB)~~ ~~This is the DWH 2.1, event auditing database and TCM repository.~~	~~Systems Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~EES~~		~~Sysman Enterprise Event Server~~	~~SYSMAN3 EES provides the Collection Layer Omnibus Object Server functionality~~ ~~Aka Object Server. Also acts as the Omnibus gateway for Auditing.~~ ~~(Equivalent to client / server TECs)~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~EFS~~		~~Sysman Enterprise Fanout Server~~	~~SYSMAN3 multi-purpose server which provides the event concentration function for events flowing from individual platforms in the branch estate and data centre,~~ ~~Software depot for Provisioning / s/w distribution, and acts as a Tivoli framework Gateway (2 are login gateways, 2 for camopus servers and the rest are branch (client) gateways)~~ ~~[TCM S/W depot, Eventing Proxy Probe, RC Proxies)~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~EMD~~		~~Sysman Enterprise Monitoring Display~~	~~EMD provides the user access into the event data held by the system management solution This is in the aggregation layer.~~ ~~It is an Omnibus Object Server and runs the web top service.~~ ~~It's in the aggregation layer. (Equivalent to the sysman2 masterTEC)~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~EMM~~		~~Sysman Enterprise Monitoring Server~~	~~Tivoli Enterprise Managing Server (aka ITM TEMS) for SYSMAN3~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~EMS~~		~~Sysman Enterprise Managing Server~~	~~SYSMAN3 TMR for HNG-X and TCM (branch)~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~EPM~~		~~Sysman Enterprise Provisioning Server~~	~~SYSMAN3 server which provides the Data Centre provisioning and software distribution functionality to Linux and Windows Campus servers. (TPM)~~ ~~[TPM server, TCM (Campus), PXE DHCP server, TPM DB server, Rembo instance?)~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~ERP~~		~~Event Reporting Platform~~	~~Netcool OMNIBUS reporting service for event analysis and statistics.~~	~~Systems Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~EST~~		~~Estate Management Database Server~~	~~Server for Estate Management Databases~~	~~Estate Management~~	~~Windows 2003 SQL Server~~	~~New~~	~~Yes~~
~~EUI~~		~~Sysman Enterprise User Interface Server~~	~~SYSMAN3 Aka TEPS Tivoli Enterprise Portal server~~ ~~(Websphere server for Web Consoles, reporting (apache or similar), Oracles 10G?)~~	~~Systems Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~HSM~~	~~Y~~	~~Hardware Security Module~~	~~Supports key management. Also known as the Networked Attalla Device.~~	~~Security Management~~	~~Other~~	~~New~~	~~Yes~~
~~HWS~~	~~Y~~	~~Help Desk Web Server~~	~~New HNG-X Web Service for online call logging, as a virtual platform on the IOS physical platform~~	Application	~~RHEL Interstage Server (Virtual)~~	~~New~~	~~Yes~~
~~KMN~~	~~Y~~	~~KMNG Server~~	~~HNG-X Key Management server.~~	~~Security Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~KSN~~	~~Y~~	~~KMNG Workstation~~	~~HNG-X Key Management Workstation.~~	~~Security Management~~	~~Workstation~~	~~New~~	No
~~MSH~~	~~Y~~	~~Maestro Scheduler~~	~~Maestro 6.0 Legacy Server~~	~~Systems Management~~	~~Solaris 9 Server (Hydra Only)~~	~~New~~	~~Yes~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~NAP~~		~~NMS - Alarm Point~~	~~Network Management "Alien" platform, base products provisioned then managed remotely by networks team.~~	~~Network Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~NCW~~		~~NMS - CiscoWorks~~	~~Network Management "Alien" platform, base products provisioned then managed remotely by networks team.~~	~~Network Management~~	~~Solaris~~	~~New~~	~~Yes~~
~~NFM~~		~~NMS - Firewall Security Manager~~	~~Network Management "Alien" platform, base products provisioned then managed remotely by networks team.~~	~~Network Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~NMN~~		~~Network Management Server~~	~~Network Management Server (HNG-X)~~	~~Network Management~~	~~Solaris~~	~~New~~	~~Yes~~
~~NMW~~		~~Network Management Workstation~~	~~Network Management workstation (HNG-X)~~	~~Network Management~~	~~Windows XP~~	~~New~~	~~Yes~~
~~NPC~~		~~NMS - Packet Capture~~	~~Network Management diagnostics (HNG-X)~~	~~Network Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~OWS~~	~~Y~~	~~Online Training Web Server~~	~~Runs the Training Online transaction simulation Web Service Application, as a virtual platform on the EOS physical platform~~	~~Application~~	~~RHEL Interstage Server (Virtual)~~	~~New~~	~~Yes~~
~~PAN~~		~~Process Area Network Controller~~	~~Controls the processors on the BladeFrame.~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~PMN~~		~~Patch Management Server~~	~~-~~	~~Security Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~PPW~~		~~PIN Pad Proving Workstation~~	~~Workstation for testing PIN pads.~~	~~Operational Support~~	~~Workstation~~	~~New~~	No
~~RAU~~		~~Radius Core Router Management~~	~~Radius server used to authenticate users who manage network appliances~~	~~Network Management~~	~~Windows 2003 Radius Server~~	~~New~~	~~Yes~~
~~RBR~~		~~Radius Branch Router Management~~	~~Radius Server for Branch Router Management~~	~~Access and Authentication~~	~~Windows 2003 Radius Server~~	~~New~~	~~Yes~~
~~RFS~~		~~RDT File Server~~	~~File Server used by RDT for backup data and storage/ transfer of reference data files.~~	~~Application~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~RGP~~		~~EMC Secure Remote Support Gateway - Policy Manager~~	~~Provides remote access to EMC equipment for support purposes~~	~~Storage and Backup~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~RLS~~		~~RDT Linux Server~~	~~RDT Linux Server (64 bit) used for Branch Database, BAL and APOP databases, used in reference data verification.~~	~~Application~~	~~RHEL Oracle Database Server~~	~~New~~	~~Yes~~

Formatted Table

Code	Vary ~~Time~~	Name	Description	Functionality	Technology	Migration	~~DC?~~
~~ROS~~		~~Router Operational Support Server~~	~~Supports the branch router. Known as ROSS. Introduced during Horizon to support the branch router rollout.~~	~~Access and Authentication~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~RSG~~		~~EMC Secure Remote Support Gateway~~	~~Provides remote access to EMC equipment for support purposes~~	~~Storage and Backup~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~RSH~~		~~RDT Solaris Host~~	~~RDT Main Unix database server used in reference data verification.~~	~~Application~~	~~Solaris Oracle~~	~~New~~	~~Yes~~
~~SPN~~		~~Performance Management Server~~	~~Database that holds short-term performance measures from Athene, HNG-X version.~~	~~Operational Support~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~SSN~~		~~SAS Server~~	~~Provides secure access to HNG-X servers.~~	~~Security Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~STL~~		~~Security Testing Server - Red Hat Linux~~	-	~~Testing~~	~~RHEL~~	~~New~~	~~Yes~~
~~STW~~		~~Security Testing Server Windows 2003~~	-	~~Testing~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~STX~~		~~Security Testing Laptop~~	-	~~Testing~~	~~Other~~	~~New~~	~~No~~
~~SWE~~		~~Sophos Web Server Emulator~~	-	~~Testing~~	~~Other~~	~~New~~	~~No~~
~~SYS~~		~~SYSLOG Server~~	~~Network event monitoring server~~	~~Network Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~TFL~~		~~Security Test Red Hat Linux Foundation~~	~~Test only platform. Foundation only build for Security testing of Red Hat Linux.~~	~~Testing~~	~~RHEL~~	~~New~~	~~Yes~~
~~TFS~~		~~Security Test Solaris 10 Foundation~~	~~Test only platform. Foundation only build for Secuirity Testing of Solaris 10.~~	~~Testing~~	~~Solaris~~	~~New~~	~~Yes~~
~~TFW~~		~~Security Test Windows 2003~~	~~Test only platform. Foundation only build for Secuirity Testing of Windows 2003.~~	~~Testing~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~TFX~~		~~Security Test Windows XP~~	~~Test only platform. Foundation only build for Secuirity Testing of Windows XP.~~	~~Testing~~	~~Windows XP~~	~~New~~	~~No~~
~~TSH~~	~~Y~~	~~Tivoli Workload Scheduler~~	~~Server that supports the Tivoli Workload Scheduler service.~~	~~Systems Management~~	~~RHEL~~	~~New~~	~~Yes~~
~~VNS~~		~~Vulnerability Scanning Server~~	-	~~Security Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~VSH~~		~~Virtual Server Host~~	~~Platform within which other virtual servers run.~~	~~Systems Management~~	~~Windows 2003 Server~~	~~New~~	~~Yes~~
~~XCS~~	~~Y~~	~~External Client Simulators~~	~~Test only platform that provides simulators for External Client systems.~~	~~Testing~~	~~Windows 2003 Server~~	~~New~~	~~No~~
~~ACE~~		~~ACE Workstation (Horizon)~~	~~Workstation used to manage Support users and their Secure ID tokens.~~	~~Security Management~~	~~Workstation~~	~~Retiring~~	~~No~~
~~ACF~~	~~Y~~	~~Autoconfiguration Database~~	~~Runs the Autoconfiguration Database (ACDB).~~	~~Estate Management~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~ACS~~		~~Counter Package Signing Server (Horizon)~~	~~Server for signing counter packages, for Horizon counter only. Aliases include the AutoConfig signing server.~~	~~Software Distribution~~	~~Windows 2000 (Hydra only)~~	~~Retiring~~	~~No~~
~~AGE~~	~~Y~~	~~Generic Agents~~	~~Runs agents that move data between databases and correspondence servers.~~ ~~Also some stand alone online services that use the correspondence servers.~~	~~Application~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~AUD~~	~~Y~~	~~Audit Workstation (Horizon)~~	~~Workstation to access the Audit system.~~	~~Application~~	~~Workstation~~	~~Retiring~~	~~No~~
~~BAC~~	~~Y~~	~~Solaris Backup Server (Horizon)~~	~~Server to backup Solaris servers.~~	~~Storage and Backup~~	~~Solaris~~	~~Retiring~~	~~Yes~~
~~BAK~~	~~Y~~	~~Windows Backup Server (Horizon)~~	~~Server to backup retiring Windows servers.~~	~~Storage and Backup~~	~~Windows 2000 (Hydra only)~~	~~Retiring~~	~~Yes~~
~~BLS~~		~~Boot Loader~~	~~The Boot Loader platform for S92 and T60 spare installation (non Branch Router Comms)~~	~~Access and Authentication~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~BOO~~		~~Boot Server~~	~~The Boot Server platform for Horizon VSAT. This includes BBOOT/WBOOT domain controller.~~	~~Access and Authentication~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~CAW~~		~~Certificate Authority Workstation (Horizon)~~	~~Creates certificates for Horizon systems. Specialised hardware. Also known as the KMA Certificate Authority Workstation (CAW).~~	~~Security Management~~	~~Other~~	~~Retiring~~	~~No~~
~~CMS~~		~~PVCS and Dimensions Signing Server~~	~~The CM signing server~~	~~Software Distribution~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~No~~
~~CNH~~	~~Y~~	~~Horizon Counter~~	~~The current Horizon counter.~~	~~Application~~	~~Counter PC~~	~~Retiring~~	~~No~~
~~COR~~	~~Y~~	~~Correspondence Servers~~	~~Messaging server to pass messages to and from counters.~~	~~Application~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~CSM~~		~~ITU Test SYSMAN2 SMDB Server (and PDC)~~	~~ITU Test SYSMAN2 SMDB Server (and PDC)~~	~~Testing~~	~~-~~	~~Retiring~~	~~No~~
~~DCA~~	~~Y~~	~~DCS and ETS Agent Server (Horizon)~~	~~Supports on-line access to the debit card and electronic top-up services, for Horizon counters only.~~	~~Application~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~DEL~~		~~ITU Test SYSMAN2 Delivery Server (named using code SYSDEL)~~	~~ITU Test SYSMAN2 Delivery Server (named using code SYSDEL)~~	~~Testing~~	~~-~~	~~Retiring~~	~~No~~
~~DOM~~	~~Y~~	~~Domain Controller (Horizon)~~	~~Platform type for Horizon domain controllers, used for the following domains: Corporate User Domain, MIS client domain, Configuration management domain, RDMC client domain, CMS signing server domain + many others.~~	~~Access and Authentication~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~ENT~~		~~ACE Server (Horizon)~~	~~Enterprise Server (ACE) runs the RSA server software to authenticate support access via SecureID tokens (Retiring Horizon platforms only).~~	~~Access and Authentication~~	~~Solaris~~	~~Retiring~~	~~No~~
~~GST~~		~~ITU Test Ghost Server~~	~~ITU Test Ghost Server~~	~~Testing~~	-	~~Retiring~~	~~No~~
~~INV~~		~~ITU Test SYSMAN2 Inventory Server (OMDB) (named using code SYSINV)~~	~~ITU Test SYSMAN2 Inventory Server (OMDB) (named using code SYSINV)~~	~~Testing~~	-	~~Retiring~~	~~No~~
~~KAW~~	~~Y~~	~~KMA workstation (Horizon)~~	~~Key management workstation. Includes the PCI Attalla card. Used by security team to manage Horizon keys.~~	~~Security Management~~	~~Workstation~~	~~Retiring~~	~~No~~
~~KMS~~	~~Y~~	~~KMA Server (Horizon)~~	~~The server for the Horizon Key Management Application.~~	~~Security Management~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~KSA~~	~~Y~~	~~KMA Admin Workstation (Horizon)~~	~~Administrative workstation into KMA Server. Used by support rather than security team.~~	~~Security Management~~	~~Workstation~~	~~Retiring~~	~~No~~
~~LGW~~		~~Sysman Login Gateway~~	~~Sysman2 Tivoli Gateway server for initial endpoint logins before reassignment to a fixed client / campus gateway~~	~~Systems Management~~	~~Solaris~~	~~Retiring~~	~~Yes~~
~~MAN~~		~~ITU Test SYSMAN2 Tivoli Workstation (named using code SYSMAN)~~	~~ITU Test SYSMAN2 Tivoli Workstation (named using code SYSMAN)~~	~~Testing~~	-	~~Retiring~~	~~No~~
~~NRA~~	~~Y~~	~~NBX Routing Agents~~	~~Runs the Network Banking routing agent.~~	~~Application~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~OCM~~	~~Y~~	~~Operational Change Management System Server~~	~~Runs the Operational Change Management System (OCMS).~~	~~Estate Management~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~OCW~~		~~OCMS Workstation~~	~~Workstation to access the Operational Change Management System.~~	~~Estate Management~~	~~Workstation~~	~~Retiring~~	~~No~~
~~OMA~~		~~Operational Management Database Archive Server (OMDBA)~~	~~SYSMAN 2 OMDB archive server~~	~~Systems Management~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~OMD~~		~~Sysman Inventory Server (OMDB)~~	~~Server for the Operation Management Database (OMDB).~~	~~Systems Management~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~OSP~~		~~One Time Password Workstation~~	~~Provides one-time passwords for engineer access to Horizon counters.~~	~~Operational Support~~	~~Workstation~~	~~Retiring~~	~~No~~
~~PGW~~		~~ITU Test SYSMAN2 Tivoli PO Client Gateway~~	~~ITU Test SYSMAN2 Tivoli PO Client Gateway~~	~~Testing~~	-	~~Retiring~~	~~No~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~PIN~~		~~PIN Pad Proving Workstation (Horizon)~~	~~Workstation for testing PIN pads.~~	~~Operational Support~~	~~Workstation~~	~~Retiring~~	~~No~~
~~SAS~~	~~Y~~	~~SAS Server (Horizon)~~	~~The Secure Access Server that provides controlled access to other servers. Only used for the retiring Horizon servers.~~	~~Security Management~~	~~Windows 2000 (Hydra only)~~	~~Retiring~~	~~Yes~~
~~SCT~~		~~Sysman Client TEC~~	~~SYSMAN2 Tivoli Enterprise Console server to which events from counters are sent~~	~~Systems Management~~	~~Solaris 9 Server (Hydra Only)~~	~~Retiring~~	~~Yes~~
~~SDC~~		~~Sysman Domain Controller~~	~~SYSMAN2 domain controller~~	~~Systems Management~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~SDS~~		~~Sysman Delivery Server~~	~~Software repository used by SYSMAN2 software distribution~~	~~Software Distribution~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~SEC~~		~~Sysman Expedited TEC~~	~~SYSMAN2 TEC used for processing special events used for controlling systems management actions / monitoring~~	~~Systems Management~~	~~Solaris~~	~~Retiring~~	~~Yes~~
~~SMD~~		~~Sysman Management Data Base (SMDB)~~	~~SYSMAN2 - Runs the SMDB database.~~	~~Operational Support~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~No~~
~~SMK~~		~~KMS Help Desk Workstation Staging Area~~	~~KMS Help Desk Workstations~~	~~Security Management~~	-	~~Retiring~~	~~No~~
~~SMR~~		~~Sysman Master TMR~~	~~Also known as the Swing TMR for the NT branch estate.~~ ~~This is temporary – not part of SYSMAN3~~	~~Systems Management~~	~~Solaris~~	~~Retiring~~	~~Yes~~
~~SMT~~		~~Sysman Master TEC~~	~~SYSMAN2 Master Tivoli Enterprise Console server to which all other TECs send events to be displayed at operator consoles~~	~~Systems Management~~	~~Solaris~~	~~Retiring~~	~~Yes~~
~~SNT~~		~~Sysman SNMP TEC~~	~~SYSMAN2 Tivoli Enterprise Console server to which SNMP traps are sent~~	~~Systems Management~~	~~Solaris 9 Server (Hydra Only)~~	~~Retiring~~	~~Yes~~
~~SPD~~		~~Short-term Performance Database Server~~	~~The Short-Term Performance Database (STPDB) is also built as a domain controller in which it is the sole member.~~	~~Operational Support~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~No~~
~~SPG~~		~~Sysman Post Office Gateway~~	~~SYSMAN2 Horizon Tivoli gateway server for post office (counter) endpoints~~	~~Systems Management~~	~~Solaris~~	~~Retiring~~	~~Yes~~
~~SRT~~		~~Sysman Radius TEC~~	~~SYSMAN2 Tivoli Enterprise Console server to which events from Radius servers are sent~~	~~Systems Management~~	~~Solaris 9 Server (Hydra Only)~~	~~Retiring~~	~~Yes~~
~~SSG~~		~~Sysman Secure Post Office Gateway~~	~~SYSMAN2 Horizon Tivoli gateway server for client facing Campus endpoints (in DMZs)~~	~~Systems Management~~	~~Solaris~~	~~Retiring~~	~~Yes~~

Formatted Table

~~Code~~	~~Vary Time~~	~~Name~~	~~Description~~	~~Functionality~~	~~Technology~~	~~Migration~~	~~DC?~~
~~SST~~		~~Sysman S-TEC~~	~~Sysman Server TEC, also known as Campus TEC.~~	~~Systems Management~~	~~Solaris~~	~~Retiring~~	~~Yes~~
~~STG~~		~~Staging Server~~	~~Software repository on which all packaged software updates are held for use by support staff for system builds~~	~~Software Distribution~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~TGW~~		~~Sysman Campus Gateway~~	~~SYSMAN2 Horizon Tivoli gateway server for Campus endpoints~~	~~Systems Management~~	~~Solaris 9 Server (Hydra Only)~~	~~Retiring~~	~~Yes~~
~~TLM~~		~~ACSLS Server (Horizon)~~	~~Robot controller.~~	~~Storage and Backup~~	~~Solaris~~	~~Retiring~~	~~Yes~~
~~TMR~~		~~TME Management Server~~	~~Horizon SYSMAN2 Tivoli Management Region (TMR) server~~	~~Systems Management~~	~~Solaris 9 Server (Hydra Only)~~	~~Retiring~~	~~Yes~~
~~VDW~~	~~Y?~~	~~VPN Loopback Workstation~~	~~Supports virtual private network (VPN) to Horizon branches. Monitoring workstation to check availability of all VPN Routes and also availability fo Web Service endpoints.~~	~~Access and Authentication~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~VEX~~	~~Y ?~~	~~VPN Exception Server~~	~~Supports virtual private network (VPN) to Horizon branches. The exception server is used for initial key delivery (incl VPN key) from KMA to branches~~	~~Access and Authentication~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~VPM~~	~~Y?~~	~~VPN Policy Server~~	~~Supports virtual private network (VPN) to Horizon branches. The Policy Server manages delivery of VPN Policy files to the VPN Servers - linked to KMA for key revocation.~~	~~Access and Authentication~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~VPN~~	~~Y?~~	~~VPN Server~~	~~Supports virtual private network (VPN) to Horizon branches. The VPN Server provides the main routing function.~~	~~Access and Authentication~~	~~NT4 Server (Hydra only)~~	~~Retiring~~	~~Yes~~
~~APW~~	~~Y~~	~~APOP Workstation~~	~~Workstation delivered to Post Office to access the APOP & TESQA systems.~~	~~Application~~	~~Workstation~~	~~Unchanged~~	~~No~~
~~IXO~~	~~Y~~	~~POL-FS SAP Archive Server~~	~~SAP IXOS archive server~~	~~Application~~	~~Solaris SAP~~	~~Unchanged~~	~~Yes~~
~~NWC~~	~~Y~~	~~POL-FS SAP Middleware Client~~	~~SAP Netweaver Middleware client~~	~~Application~~	~~Solaris SAP~~	~~Unchanged~~	~~Yes~~
~~NWS~~	~~Y~~	~~POL-FS SAP Middleware Server~~	~~SAP Netweaver Middleware server~~	~~Application~~	~~Solaris SAP~~	~~Unchanged~~	~~Yes~~
~~PAT~~	~~Y~~	~~Patrol Workstation~~	~~Workstation for monitoring through BMC Patrol.~~	~~Operational Support~~	~~Workstation~~	~~Unchanged~~	~~No~~
~~SAP1~~	~~Y~~	~~POL-FS SAP APP Server~~	~~SAP Application Server~~	~~Application~~	~~Solaris SAP~~	~~Unchanged~~	~~Yes~~
~~SAP2~~	~~Y~~	~~POL-FS SAP Host~~	~~SAP Database Server~~	~~Application~~	~~Solaris SAP~~	~~Unchanged~~	~~Yes~~

Formatted Table