Document Title: HNG-X Technical Network Architecture

Document Type: Architecture (ARC)

Release: N/A

Abstract: This is the topic architecture for Networking under the overall solution architecture

Document Status: DRAFT

Author & Dept: Mark Jarosz

Internal Distribution: As per review details

External Distribution: As per review details

Approval Authorities:

Name	Role	Signature	Date
Giacomo Piccinelli	Solution Architect		
Phil Day	Programme Manager		

0 Document Control

0.1 Table of contents

0.2 Figures and Tables

0.3 Document History

Version No.	Date	Summary of Changes and Reason for Issue	Associated Change - CP/PEAK/PPRR Reference
0.1	30-Oct-06	Draft	
0.2	07-Nov-06	Draft for review	
0.3	12-Dec_06	Draft for Review	
0.4	08-May-07	Draft for Review	
0.5	04-Oct-07	Incorporated comments from version 0.3 Documented change of strategy for Branch Router as well as changed model for obtaining branch specific data from Estate management	

0.4 Review Details

Review Comments by :	Friday 19th October 2007
Review Comments to :	mark.jarosz GRO & RMGADocumentManagement GRO
Mandatory Review	
Role	**Name**
HNG-X Solution Design	Adam Cousins
HNG-X Solution Design	Peter Ambrose
HNG-X Solution Design	Steve Dingle
HNG-X Security	Bill Membery
HNG-X Infrastructure Design	David Sackman (or nominees)
HNG-X Service Transition	Steve Godson
HNG-X Security Architect	Jim Sweeting
HNG-X Test Design	Peter Robinson
Optional Review	
Role	**Name**
Test Design	George Zolkiewka
Development	Graham Allen
Service Network	Alex Kemp
Business Continuity	Tony Wicks
SSC	Mik Peach
SV & I Manager	Sheila Bamber
Tester	Hamish Munro
TE & VI Manager	Peter Rickson
RV Manager	James Brett (POL)
Integration	David Hinde
Data Centre Migration	Martin Brett
HNG-X Programme Manager	Phil Day
Core Services	Ed Ashford
Core Services	Pat Lywood
Core Services	Andrew Gibson
HNG-X Solution Architect	David Chapman
HNG-X Solution Architect	Ian Bowen
HNG-X Solution Architect	Jeremy Worrell
HNG-X Solution Architect	David Johns
HNG-X Solution Architect	Mario Stelzner
HNG-X Solution Architect	Dave Tanner
HNG-X Solution Architect	Dave Haywood

HNG-X Solution Architect	Ian Devereux
HNG-X Solution Architect	Rahman El-Khoulali
HNG-X Solution Architect	Ghalib Al-Kilidar
HNG-X Solution Architect	Stephen Wisedale
HNG-X Solution Architect	Temitayo Fashina
Head of Service Transition and Change	Graham Welsh
HNG-X Solution Architect	Colin Mills
HNG-X Solution Architect	Paul Tomlinson
HNG-X Solution Architect	Pat Carroll
HNG-X Solution Architect	Andrew Oram
HNG-X Solution Architect	Lee Walton
HNG-X Solution Architect	Nigel Ainge
Issued for Information – Please restrict this distribution list to a minimum	
Position/Role	**Name**
HNG-X Project Manager	Dean Parsons
HNG-X Project Manager	John Sawyer

0.5 Associated Documents (Internal & External)

Reference	Version	Date	Title	Source
PGM/DCM/TEM/0001 (DO NOT REMOVE)			Fujitsu Services Post Office Account HNG-X Document Template	Dimensions
ARC/GEN/TEM/0001 (DO NOT REMOVE)			Fujitsu Services Post Office Account HNG-X Architecture Document Template	Dimensions
ARC/SOL/ARC/0001			Horizon Next Generation - Plan X (HNG-X) Solution Architecture Outline	Dimensions
DESNETDPR002			Branch Router Design Proposal	Dimensions
ARC/SEC/ARC/0003			HNG-X Technical **Security** Architecture	Dimensions
REQ/CUS/STG/0001			HNG-X Migration Strategy - Agreed Assumptions and Constraints	Dimensions
			SG VPN for PathWay - Phase 3 Extensions	DN: Need to get this booked into PVCS if not already there
			DN: Need reference for security risk assessment of use of internet for VSAT Broadband	

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

0.6 Abbreviations

Abbreviation	Definition
ACL	Access Control List

Abbreviation	Definition
ACE	Application Control Engine
ADSL	Asynchronous Digital Subscriber Line
ASDM	Advanced Security Device Manager
BGP	Boundary Gateway Protocol
CHAP	Challenge Handshake Authentication Protocol
CLI	Calling Line Identity
CPU	Central Processing Unit
3G	Third generation of mobile phone standards and technology
A&L	Alliance & Leicester
APN	Access Point Name
BGP	Border Gateway Protocol
BT	British Telecom
C&W	Cable and Wireless
CAPO	Card Account Post Office
CE Router	Customer Edge Router
CHAP	Challenge Handshake Authentication Protocol
CNIM	Counter Network Information Monitor
CNIM2	Counter Network Information Monitor mk 2
CSM Blade	Content Switching Module
CTO	Counter Training Office
DNS	Domain Name System
DR	Disaster Recovery
DVLA	Driver Vehicle Licensing Authority
EAL4+	Evaluation Assurance Level 4+
EDGE	Enhanced Data rates for GSM Evolution
E-PAY	
FC	Fibre Channel
FE	Fast Ethernet
FSBN	Fujitsu Services Business network
FWSM	Firewall Services Module
GE	Gigabit Ethernet
GPRS	General Packet Radio Service
GRE	Generic Routing Encapsulation (specified in RFC 2784)
HNG-X	Horizon Next Generation
HSCSD	High Speed Circuit Switched Data
ICMP	Internet Control Message Protocol

Abbreviation	Definition
IDS	Intrusion Detection System
IP	Internet Protocol
IPSEC	IP Security
ISDN	Integrated Services Digital Network
L2TP	Layer 2 Tunnelling Protocol
LAN	Local Area Network
LCP	Link Control Protocol
LNS	L2TP Network Server
MAC	Media Access Control
MPLS	Multiprotocol Label Switching
MSS	Maximum Segment Size
NAT	Network Address Translation
NMS	Network Management System
NNM	Network Node Manager
NTP	Network Time Protocol
OBC	Outlet Business Change
OOB	Out of Band
P2P	Peer-to-peer
PAS	HNG-X Private Internet Address space (PAS)
PE Router	Provider Edge Router
PHU1	Portable Hardware Unit 1
PO	Post Office
PPP	Point to Point protocol
PSK	Pre Shared Keys
PSTN	public switched telephone network
RADIUS	Remote Authentication Dial In User Service
RDMS	Reference data management system
RDP	Remote Desktop Protocol
RFC	Request For Comments
RIP	Routing Information Protocol
RMGA	Royal Mail Group Account
SAN	Storage Area Network
SID	Session identity
SNTP	Simple Network Time Protocol
SP	Service Provider
SSC	Systems Support Centre

Abbreviation	Definition
SSH	Secure Shell
SSL	Secure Sockets Layer
TACACS+	Terminal Access Controller Access-Control System plus
TCP	Transmission Control Protocol
TLS	Transport Layer Security Protocol
TNS	Transaction Network Services
UDP	User Datagram Protocol
UMTS	Universal Mobile Telecommunications System
VIP	Virtual IP
VSAN	Virtual Storage Area Network
VSAT	Very Small Aperture Terminal
WAN	Wide Area Network

0.7 Glossary

Term	Definition
cblade	Control Blade
DMZ	Demilitarized Zone – A separate subnet isolated from all other networks by a Firewall
DWDM	Dense Wavelength Division Multiplexing; An optical technology to provide network bandwidth over fibre optic backbones. DWDM works by combining and transmitting multiple signals simultaneously at different wavelengths on the same fibre. In effect, one fibre is transformed into multiple virtual fibres.
IP Subnet	All components / platforms on a common subnet can communicate with use of a Router (Layer 3 switch)
IRE11	Ireland 11 Data Centre
IRE19	Ireland 19 Data Centre
IRE1X	Refers to both IRE11 and IRE19
MSS	The TCP MSS value specifies the maximum amount of TCP data in a single IP datagram that the local system can accept (reassemble). The IP datagram can be fragmented into multiple packets when sent. Theoretically, this value can be as large as 65495, but such a large value is never used. Typically, an end system uses the "outgoing interface MTU" minus 40 as its reported MSS. For example, an Ethernet MSS value is 1460 (1500 - 40 = 1460). Using such a value avoids IP fragmentation.
Network Appliance	A managed device for providing network functionality, for example a Cisco Router.
Network Appliance	Hardware and associated software for performing Network functions. Typically these are Routers, Multilayer switches, Firewalls and IDS appliances.
Platform	A Platform is something which has Enterprise management components (SYSMAN) installed
VLAN	Virtual LAN, provides a broadcast domain at layer 2
Syslog	Syslog is a standard for forwarding log messages in an IP network. It is specified in RFC 3164.

0.8 Changes Expected

Changes
Updates following review comments Resolution of all DN's Specification of OOH laptop solution

0.9 Accuracy

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

0.10 Copyright

1 Scope

This is the **topic architecture** for Networking under the overall solution architecture for HNG-X.

It covers the target Network solution for HNG-X in respect of:

1. Principles and guidelines
2. Network layout Model – structure covering components and interrelationships
3. Properties of the network

In addition to describing the Target solution this document is concerned with migrating from the Current solution to the Target solution. This Migration covers three phases;

- Preparation: Changes in existing Horizon Data Centres to support HNGX in general, for example simplifying or enabling data centre move.. This phase lasts until the Horizon Data Centres are decommissioned and therefore have left the network.
- Four Data Centre: The Period of time over which there are 4 data centres participating in the Network, these being the existing Data centres in Wigan and Bootle and the new Data centres namely IRE11 and IRE19.
- Dual running: Horizon and HNG-X in new Data Centres (IRE1x). This phase starts when Branch traffic is directed to the IRE1X data centres.

The following diagram provides an overall view of the Target Network solution for HNG-X. Note that the terms Primary and Secondary data centre are defined in section 2.3.1 . The Diagram shows a single network service into each Data centre and resilience is achieved by using the Intercampus LAN for triangulation.

Further details of the Wide Area network are provided in section 2.5.

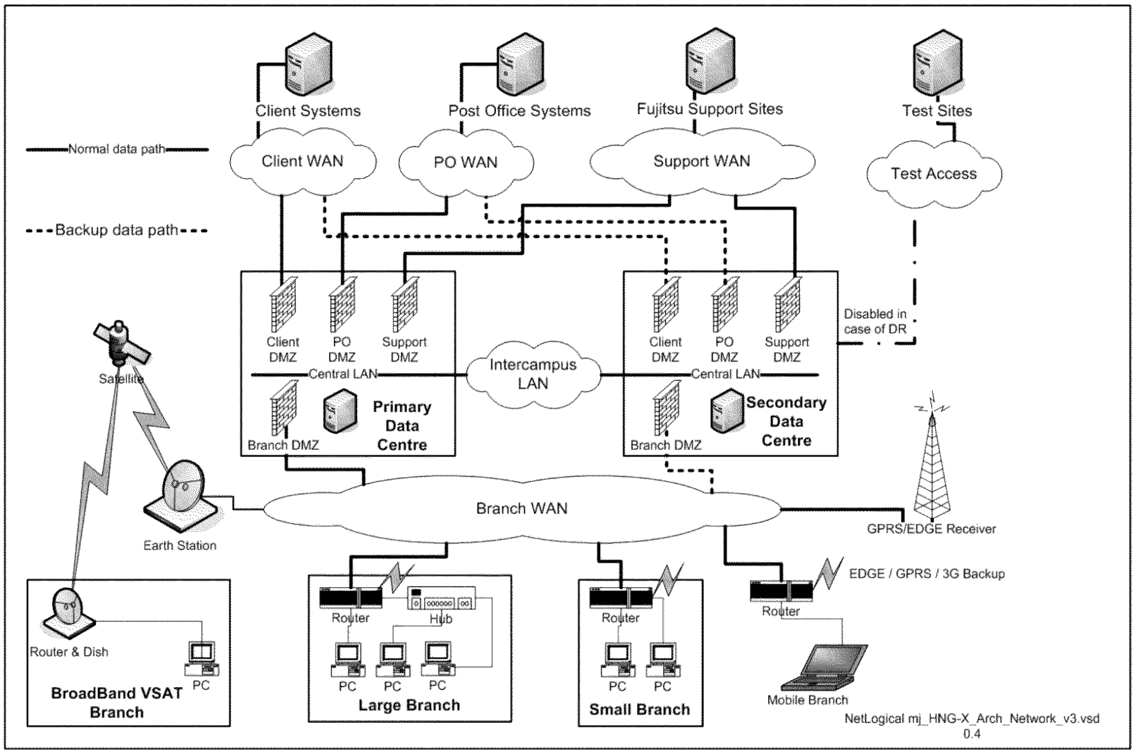

Figure 1 – Overall view of the HNG-X Target Network solution

2 Architectural description

2.1 Model

The Following diagram illustrates the model (solution space as distinct from topology) used to describe the IP and Storage Network space. The top level division of the IP Network Space is into Data Centre Network, Branch Network, Transit Network and Wide Area Network. Note that the model shows instances for clarity but these are not exhaustive.

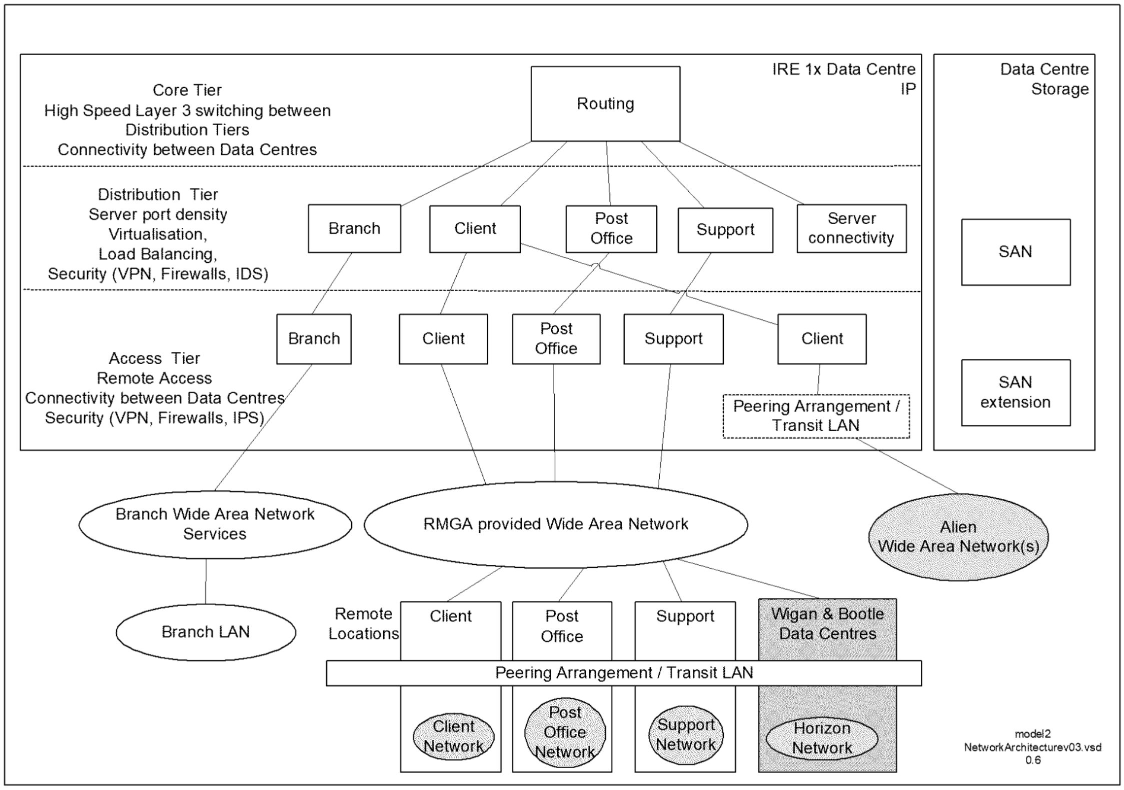

Figure 2 – Network Model

2.2 Key Principles

The following table list some key principles applied to the network.

Area	Summary
Network Protocol	The Network protocol is IP Version 4 (RFC 791)
Application Communication	All Unicast application traffic is either UDP (RFC 768) or TCP (RFC 793) The network is optimised for application communication using TCP.

Area	Summary
IP Address semantics	The semantics of an IP address may be both locality and identity. For example the data centre application interprets the source IP of an incoming TCP connection from a branch as the identity of the endpoint and assumes persistence of this identity. Conversely address pools are used in the case where the need to attribute identity does not apply, for example PPP interface addresses for Wireless WAN.
IP address space	When allocating new IP addresses for HNG-X ,Private IP addressing is used (RFC 1918 – Address Allocation for Private Internets) For those IP addresses assigned to endpoints which are eventually disappearing, for example Correspondence servers, the choice as to whether the existing IP address is retained or a new one assigned is made on the basis of simplifying the migration phases. For those IP addresses assigned to endpoints which are retained from Horizon, for example Branch Counters, the choice as to whether the existing IP address is retained or a new one assigned is made on the basis of simplifying the migration phases. The impact of the above is that some IP addresses in the target solution may not be RFC 1918 compliant as they were inherited from Horizon and retained to optimise the migration process for simplicity. Test IP addresses are allocated from an allocated from a unique collection that is not used for Live IP addresses. Specifically three is never ambiguity about whether an IP address is used for test or Live.
Peering IP address space	For each 3rd party with which HNG-X exchanges IP data grams a Peering IP address space is defined. Each such Peering IP address space will either be under administrative control of the 3rd party or under HNG-X control and shall be specified in the relevant Technical Interface document. It will also be defined / referenced in any OLA/SLA or contractual agreements for mitigation purposes and incident management. In the 3rd party case they specify the address space. In the HNG-X case the IP address space is allocated as for non-peering HNG-X IP addresses stated above.

Area	Summary
Resilience	No single points of failure exist within the Data Centre Network. The inter data centre Network is resilient with no single points of failure Wide area network resilience for main (carrying branch and MPLS traffic) high speed links is achieved by (use of two fully separate Inter data centre service). For an individual data centre, there is a single network link, terminating on a single Router (termed CE). The reason for this approach is to avoid the additional cost of providing resilient services into both Data Centres. This approach introduces the risks that; • In the event of DR, the Secondary data centre is prevented from becoming fully active as some WAN services are not available due to a failure. This is effectively a double failure • The loss of either data centre for a prolonged period will result the remaining data centre having no resilience for WAN services (Branch Traffic, Support traffic, Traffic to clients such as E-pay, DVLA and Traffic to Royal Mail group). (DN: It will be necessary to confirm that this behaviour is acceptable) Wide area network resilience for low speed links (< 8 M bits / sec) links is achieved by providing resilient services into each data centre. This is so as to avoid a single failure of such a service preventing a successful DR transition. In general Single Points of failure are avoided in the Branch Access Network. Where exceptions exist which cannot be avoided due to service constraints, for example in the ADSL network the devices terminating the ATM service are single points of failure, then these are documented in the relevant High Level Design together with a predication of the number of branches impacted and agreed contractually and with service management.
Layer 2 switches – Auto negotiation	Auto negotiation of the properties listed below is avoided, rather the ports are set explicitly with respect of these properties; Speed settings Duplex settings MDI/MDIX This applies to Network appliance switch ports and hosts Ethernet interfaces.
Optimise for the most common case	In a Branch network of 14,000+ locations, there are cases of exceptions that apply to a small number of branches. The general approach in dealing with these is optimising for the majority case.

Area	Summary
DNS	It cannot be assumed that all applications locate network endpoints using DNS.
DR transition	This needs to be achieved with no cabling change and minimal configuration change
Clustering models	Within the network Active / Standby is used to provide deterministic traffic flows. This applies to service blades which provide load balancing, firewall and virtualisations services and Appliance Firewall pairs. Similarly deterministic traffic flows are also ensured by routing, i.e., by default load share is not configured.
IP Fragmentation	IP fragmentation is the breaking up of a single IP datagram into two or more IP datagram's of smaller size. It happens because every transmission medium has a limit on the maximum size of a frame (MTU) it can transmit. As IP datagram's are encapsulated in frames, the size of IP datagram is also restricted. If the size of An IP datagram is greater than this limit, then it must be fragmented. IP Fragmentation may result in less reliable communications are Routers protect themselves from the number of fragments they handle per unit time. For this reasons explicit measures are used to avoid IP fragmentation of any TCP based traffic flows. These apply to VSAT Broadband and Wireless WAN networks...
Coordinated Changes	The need to synchronise major changes across multiple network components / services and services is avoided.
Change of Data Centre	Fundamentally this is a migration of services to be delivered out of different locations. There is no constraint on keeping the network topology and platforms the same in the new data centres in order to minimise change from existing Horizon network.

2.3 Data Centre Network

2.3.1 Data Centre state

This section provides an overview of how the network supports the transition of operational roles for the HNG-X data centres resulting from declaration of DR. An introduction to these roles and how the network supports DR is documented in section 7.3. It is important to note that from a Network perspective, the network is Active at both IRE1x data centres. Also note that the transition to DR is a manual decision.

There are two network perspectives on the operational modes of the two HNG-X data centres. The first of these is the **data plane** perspective which is concerned about where and what type of application traffic is flowing. The second is the **control plane** perspective which is concerned with traffic for management and control of Network appliances.

The HNG-X data centres have constant names according to physical location; Primary refers to the data centre where the Branch Session servers normally accept counter sessions (known as IRE11). Secondary is the other data centre (known as IRE19) .

2.3.1.1 Data plane

Network behaviour depends on state across the two data centres defined using the following notation;

Primary (DC_*state*), Secondary (DC_*state*)

DC_State is in {active, standby, inactive, indeterminate, DRblocked} and these values are explained below;

- Active means that this data centre is advertising Virtual IP's (for example the branch session servers) with the lowest cost so it is declaring itself to be the preferred data centre. It should be noted that this is a necessary but not sufficient condition for correct operation.
- Standby means this data centre has active applications / traffic flows but is not primary. For example Atalla appliances are active at both data centres and the Access network is triangulated across both data centres.In this state the data centre MAY also configured to support Test services.
- Inactive means the data centre will not have any external traffic (other than support) sent to it as a result of either the other data centre advertising routes with a lower cost or this data centre not advertising routes.
- Indeterminate means that it is not possible to assert the state of the data centre. This generally would be the case for the Primary Data centre following an unplanned DR
- DRblocked means that the High speed WAN link into the Secondary data centre has failed and that coincident with this, there is no communication possible via the path through the other data centre.

Possible states for the Primary Data Centre are; {Active, Inactive, Indeterminate} and possible states for the Secondary Data centre are {Active, Standby, Inactive, Indeterminate, DR_blocked}.

It is important to note that data centre states are not the same as service states. For example the normal Data centre state of Primary (Active), Secondary (Standby) supports the following valid service states

- NTP (Active / Active), note this also applies to Horizon services such as Riposte during migration.
- Branch Session Service (Active / Test)

Systems supporting services which are not (Active / Test) must manage the master /slave or active (preferred) state themselves and not assume any network enforcement of the service state.

The normal operation would be; Primary (active), Secondary (standby).

In this state;

- All branch sessions are with the Branch Session server at the Primary data centre
- Assuming no failures of the network services into the Primary data centre, then all Live network traffic is delivered to the Primary data centre
- Test traffic is delivered into the Secondary data centre

The transitions between states are defined by the following table;

Note these transitions only cover the first hour after the DR has been formally invoked. Longer term actions such as revoking key material etc are not covered.

Current state		Event / Action	New state		Comment
Primary	Secondary		Primary	Secondary	
Active	Active				Not valid state
Active	Standby	unplanned DR / invoke DR protocol	Inactive	Active	
Active	Standby	planned DR / invoke DR protocol	Inactive	Active	
Active	Inactive	unplanned DR	Indeterminate	Inactive	Secondary Data centre needs to be brought to Standby state.
Active	Indeterminate	unplanned DR	Indeterminate	Indeterminate	Secondary Data centre needs to be brought to Standby state.
Active	DRblocked	unplanned DR	Indeterminate	DRblocked	Secondary Data centre can be brought to Standby state once WAN communications has been restored.
Inactive	Active	Invoke Restore Primary protocol	Active	Standby	Primary data centre must be ready - so inactive here means ready to go live
Inactive	Standby	Invoke DR protocol	Inactive	Active	
Inactive	Inactive				To progress from this state either the Primary data centre needs to be restored to Active or the Secondary one needs to brought to the standby state
Inactive	Indeterminate				To progress from this state either the Primary data centre needs to be restored to Active or the Secondary one needs to brought to the standby state
Inactive	DRblocked				To progress from this state either the Primary data centre needs to be restored to Active or the Secondary one needs to brought to the standby state
Indeterminate	Active				To restore the Primary data centre it needs to be brought to the inactive and ready to go live state
Indeterminate	Standby	Invoke DR protocol	Inactive	Active	
Indeterminate	inactive				Secondary Data centre needs to be brought to Standby state.
Indeterminate	indeterminate				At least one Data Centre needs to be brought to the inactive state
Indeterminate	DRblocked				Secondary Data centre can be brought to inactive state once WAN communications has been restored.

Operational states are shown highlighted Version 0.2

Primary is Live
Secondary is Live

Figure 3 – Data Centre transitions

2.3.1.1.1 Managing state transitions

In order to support the transition of operational roles for HNGX data centres, Scripts will be provided on the NNM. The action of these scripts is defined in the following sections.

The configuration file on the Core and Access switches (refer to sections 2.3.3, 2.3.4 and 2.3.5.4) in the secondary data centre will indicate the state of the data centre (in a comment within the configuration) as follows;

State	Meaning
Standby	As per section 2.3.1.1
Active	As per section 2.3.1.1
Transition to Active	The Invoke DR protocol script (refer to Section 2.3.1.1.2) has started running. Note that towards the end of the script – that is when it has applied all changes that are significant in affecting traffic, the script will change this variable to Active
Transition to Standby	The Network restore DR protocol script (refer to Section 2.3.1.1.3) has started running Note that towards the end of the script – that is when it has applied all changes that are significant in affecting traffic, the script will change this variable to Standby

2.3.1.1.2 InvokeDR protocol

This protocol involves a series of scripts whose goal is to transition a Secondary data centre to Active from a standby state.

The actions within these scripts will result in;

- Steering all traffic (other than from the support traffic class) to the Secondary Data Centre. Note that due to WAN triangulation it MAY be the case that some traffic flows through the Primary Data centre in the case where the network service into the Secondary data centre has failed and the inter data centre WAN service and Link into the Primary Data centre are operational. This condition will prevail until the WAN service into the Secondary data centre is restored.
- Blocking all Test traffic flowing to the Secondary data centre.

The key constraint imposed on this protocol is that it cannot make any changes at the Primary Data Centre. This is because there are DR scenarios where this data centre may be in some indeterminate state with limited or no physical and network access.

Note that phases 1, 2, 3 & 4 are applied at the Secondary Data Centre. Also note that all network appliance configuration changes are committed to Non volatile memory.

Phase 1

All Test traffic flows are prevented by disabling test Ethernet ports in the Access Tier and Distribution Tier (1). This is achieved by running an SSH script on the NNM.

Note the general technique to close off areas of the network is combination of: (1) removing ports from the VLAN, (2) shutting down layer 3 interfaces, (3) reconfiguring hosts - i.e. test mode to live mode transition.

Phase 2

Change the metrics on all Routes advertised (or to be advertised when for example the Branch Session servers come up) into the Wide Area network and Branch network clouds to be less than the values used from the Live data centre. This will steer all traffic to the DR data centre since all ingress methods to the Data Centre are based on Routing. All network appliance configuration changes are committed to Non volatile memory.

Phase 3

Enable Live VLANS - these are VLANs which support platforms that are only active at no more than one data centre at a time.

Phase 4

Run checks (To be specified) and report that Secondary Data Centre is Active and ready to go

(DN Need to include storage networking aspects in next version)

2.3.1.1.3 Network restore DR protocol

This protocol involves a series of scripts whose goal is to;

- Transition a Active Secondary data centre to Standby
- After the above transition has been confirmed; Transition an Inactive Primary data centre to Active.

Details of this protocol will be confirmed during the Design phase.

2.3.1.2 Control Plane

From a network management & control perspective both data centres are active and there is a single Network management domain.

Network addressing is unique across both data centres with the following exceptions;

- Virtual IP addresses can exist at both data centres. They are unambiguous through routing and or proximity
- Subnets used to support platforms which can only be active at a single data centre are configured at both data centres. However the subnets are disabled (using mechanisms stated above under phase 1) unless the Secondary data centre is in the Active state. Therefore they are unambiguous through mutual exclusion.
- Specifically and following on from the previous point; the Network solution will support the same subnet in two data centres exclusively for the purpose of deploying HA clusters. This means that during DR Failover, the host may keep the same IP address. The subnet Data centre instance may only have hosts on it when it is at the Active Data Centre.

2.3.1.3 Properties of Secondary data centre

The mechanisms summarised in the above sections result in the following characteristics;

- When Secondary data centre is in the Active state (that is DR has been declared) – No Test Access is possible as all Ingress is shut off (By configuration of Ports).

- When Secondary Data Centre is in the Standby state it will support Test Traffic classes and Live Traffic classes. For example test counters in Bracknell may be accessing the Blade frame chassis whilst there are Active Atalla servers with live key material.

2.3.2 Data Centre layout overview

The three tiers of this hierarchical model for the Data Centre are referred to as the Core, Distribution (or Aggregation), and Access (or Peering) tiers. This section covers;

- Main functions and components of these tiers
- Traffic flows between the tiers
- Server connectivity and use of Firewalls

2.3.3 Core Tier

The functions of the Core Tier are to;

- Layer 3 switch traffic at very high speed between the Distribution Tier domains.
- Inter- Campus traffic; Layer 3 / 2 switch traffic between HNGX Data centres (IRE1x).

A Collapsed Core Model is used with the Core Functions taking place in the Distribution Tier Multilayer switches.

Note since the Distribution Tier Network is collapsed into a pair of Multilayer switches (see section 2.3.4), the optimum approach to providing layer 3 switching within the Data Centre Distribution Tier is internally within the Multilayer switches.

2.3.4 Distribution Tier

The functions of the Distribution Tier are to;

- Provide Collapsed Core Functions (refer to section 2.3.3)
- Distribute network traffic between servers
- Separate locally (within Tier) destined traffic from the network traffic destined for other Tiers through the Core tier.
- Provide port density to the server farm (*)
- ACE blade for Load Balancing and Virtualisation services
- Provide Firewall services (covered in section 6)
- QOS to separate test and live traffic plus backup traffic
- IDS services

(*) Server termination is collapsed into the Distribution Tier. Therefore the role of Distributing network traffic between Access layers which is usually assigned to this tier does not apply

The following diagram summarises the functions of Distribution Tier.

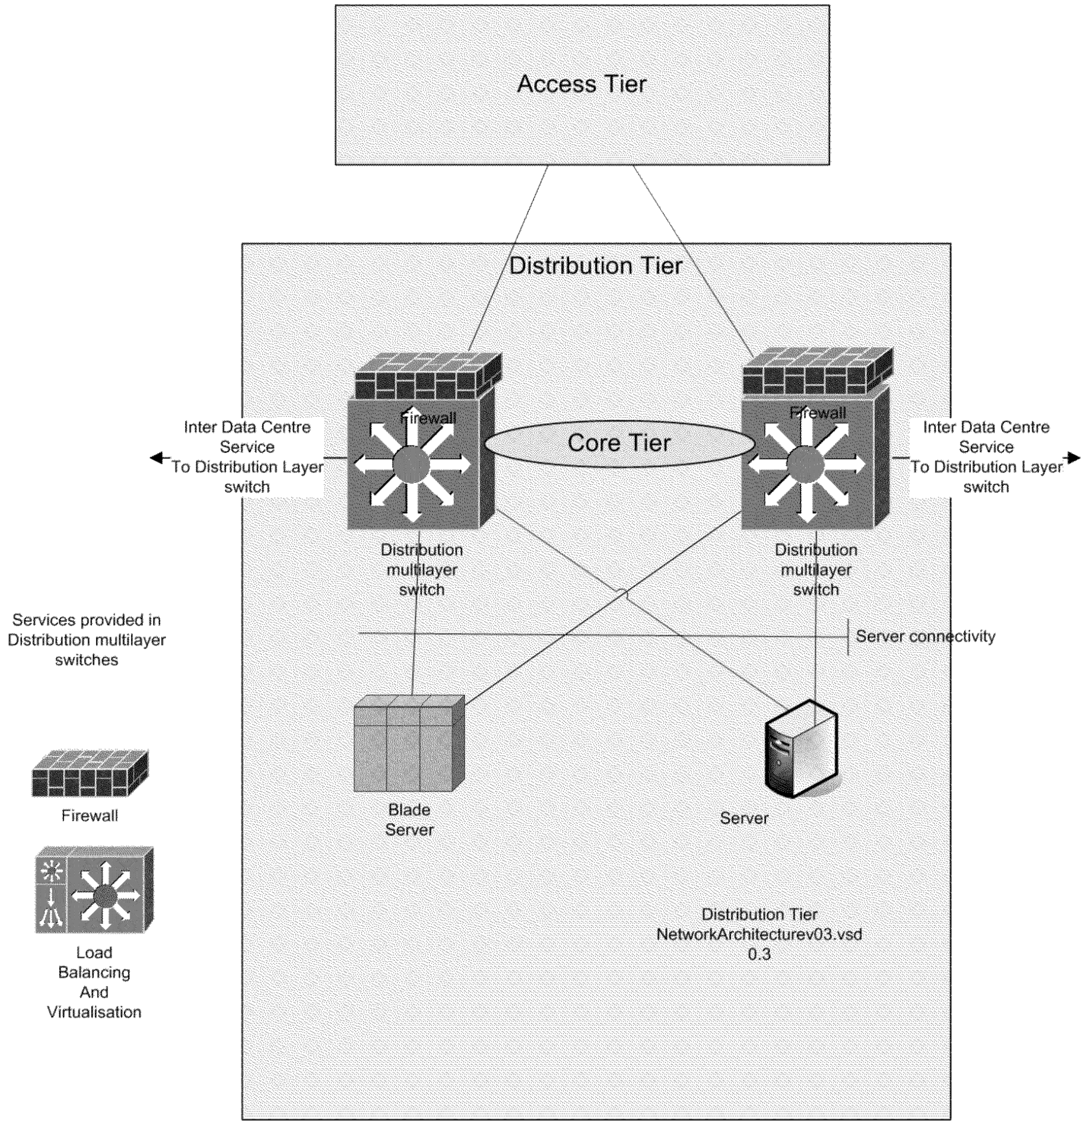
Access Tier
Distribution Tier
Firewall
Firewall
Core Tier
Inter Data Centre Service To Distribution Layer switch
Inter Data Centre Service To Distribution Layer switch
Distribution multilayer switch
Distribution multilayer switch
Services provided in Distribution multilayer switches
Server connectivity
Firewall
Load Balancing And Virtualisation
Blade Server
Server
Distribution Tier NetworkArchitecturev03.vsd 0.3

Figure 4 – Distribution Tier

2.3.4.1 Distribution multilayer switches

Enterprise class switches are chosen in order to provide the availability and stability associated with this class of hardware. The Model specified is the Cisco 6513 and these have been in use on Horizon for over 5 years. As well as proving reliable there is the advantage of design and operational experience with these devices.

These switches are chassis devices, dual power supplies connected to separate phases and with a high speed backplane into which a number of blades are fitted.

The performance characteristics of these switches and their service blades are specified in AB.1.

The blades provide for the following;

- Ethernet blades for server connectivity
- Layer 3 switching
- ACE blade for Load Balancing and Virtualisation
- FWSM blade for Firewall functions

Each switch chassis will be deployed in physically separate racks with fully resilient power.

A pair of these switches will be used to provide resilience as follows;

Servers are dual connected (refer to section 2.3.4.2 for details)

- There is one ACE blade per switch and these are in an Active / standby arrangement. Failover is state full within < 5 seconds and TCP connections survive
- There is one FWSM blade per switch and these are in an Active / standby arrangement. Failover is state full within < 5 seconds and TCP connections survive
- There is one ACE blade per switch and these are in an Active / standby arrangement. These provide stateful application redundancy with replication of TCP connection state and sticky tables

It is important to note that SSL failover is not state-full. Should an ACE blade fail then the following is likely to happen;

1) If a TCP connection exists to the Counter then it should survive the ACE blade failure as TCP failover is stateful. The newly Active ACE blade should send a TLS / SSL "Hello" in response to TCP application payload which should cause the TCP client to renegotiate the SSL session. This should be transparent to the application.
2) If no TCP connection exists to the Counter, then the next TCP connection will result in a new SSL session as opposed to the usual case of reuse of an existing session.
3) Note that the ACE blade can support 2,500 session setups / second so all 35,000 counters will be able to establish new SSL sessions within 20 seconds. Should more than 2,500 session setups be attempted per second then it is assumed that the ACE will handle this gracefully through queuing and discard. In the case of a discard then the client may be notified explicitly through a TCP reset of a failed TCP connection. The suitable recovery action is for the client to re-establish a TCP connection ideally with random exponential time back off between TCP initiating connect attempts or at least with some bound (< 5 say) on the number of such attempts. This is to avoid the case where TCP connection attempts are made repeatedly every 50 milliseconds or so (approximate ADSL network roundtrip time). Note that the behaviour of the ACE under overload will be determined during LAB testing.

The behaviour implied by (1), (2) and (3) will need to be verified during testing.

The switches are connected together with multiple GE ports on 2 different blades for inter switch traffic. The number of ports used for inter switch traffic will be determined using a capacity model (refer to section 8.2).

2.3.4.2 Server Connectivity

DN: Need to bring server terminology into line with ARC/SYM/ARC/0007 – also include as an xref.

Free standing servers

These have two network interface cards configured in an active / standby arrangement creating a single Virtual interface with an IP address. Each network interface is connected to a separate Distribution Multilayer switch Ethernet port. For those servers that have a separate interface for Management purposes, the same model applies.

Blade Frame chassis

Multi Mode

A "Multi Mode" Blade frame is defined as one which supports production workloads and testing workloads with transitions between the two to support the Data centre transition (refer to section 2.3.1).

There are 3 such Blade Frames per HNG-X data centre, each with 2 cblades. Each cblade has 8 GE ports plus 1 FE management port. All ports from 1 cblade will be connected to a single Distribution Multilayer switch with the ports on the other cblade connected to an alternative switch. Each such connection will carry traffic for multiple VLANS.

Always Active

An "Always Active" Blade frame is defined as one which supports production work only and is active at both centres continually. There is one such Blade Frame per HNG-X data centre.

All Server VLANs are spanned across the switches with per VLAN spanning tree to avoid loops.

DN: Need general overview of virtualisation and network interaction.

2.3.4.3 Application Front End Services

These are SSL offload, Load balancing and Virtualisation.

SSL offload is used to terminate SSL sessions from the counters. Refer to sections 6.2.2 and 6.3.4 for further details. Note that SSL is used in point-to-point mode between two application endpoints as distinct from tunnelled mode.

The ACE blade is used to provide Load balancing and Virtualisation. The following example illustrates both these concepts using the notation introduced in section 2.3.1.1).

1. The Counter application (HTTP client) targets a Virtual IP address (V1) which represents the HTTP server on the Branch session servers.
2. At the Primary (Active) Data Centre the ACE is probing (http) the Array of Branch session servers. Those that are functioning will respond (http response). Providing at least one session server responds, the ACE will consider that V1 exists and inject the Route unto the host multilayer switch routing table (this feature is called Route health injection). This Route will be advertised into the network using a Routing protocol.
3. At the Secondary (Standby) Data centre the ACE is performing similar probing. However because there are no live Branch session servers in existence, there is no response and the ACE does not advertise V1. Note that V1 would be advertised at lower priority from the Secondary (Standby) data centre, therefore in the event of a single miss-configuration which causes V1 to be advertised from the Secondary (Standby) data centre, branch traffic would still be routed to the Primary (Active) Data Centre.
4. In the event of DR being invoked, the Primary (Active) Data Centre may transition cleanly to Primary (Inactive). Note there may a partial shutdown in which case parts of the Primary (indeterminate) Data Centre are still functioning.

5. Shortly after DR being invoked the Secondary (Standby) data centre will transition to Secondary (Active). Once this has happened the Branch session servers at this data centre will become active and respond to probes (as in step 2).This will cause V1 to be advertised.
6. To deal with the potential problem that the Primary (Indeterminate) data centre is still advertising V1, the DR network script (refer to section 2.3.1.1) will change the priority that V1 is advertised from the Secondary (Active) data centre to be higher than that from the Primary (indeterminate) data centre.

Note that in the above example, the ACE probing is local to the data centre, there is no inter-data centre probing.

A further example of where virtualisation (Network Virtualisation as opposed to Platform) applies is the case of the Atalla security appliances. These are Active at both HNG-X data centres with different interface IP addresses. Applications will be configured to use a single virtual IP address to reach these servers. This means that there is no requirement to configure applications differently depending on which data centre they are running.

Note that there will be both Live and Test Atalla devices at the Secondary data centre (in standby mode so supporting testing), but these will have different VIP addresses.

In general Network Virtualisation abstracts both semantic roles of IP address (location) and endpoint identifier. In the case of location this can be at different data centres with the choice based on proximity. In the case of endpoint identifier this is no longer tied to the interface address of the specific platform.

2.3.4.4 Inter Data Centre Network

Each Distribution layer switch has a single GE intercampus service to a peer Distribution layer switch at the other Data Centre. For further details of this approach and the underlying services refer to section 2.3.5.6.

2.3.5 Access Tier

The functions of the Access Tier are to provide;

- Ingress and egress of network traffic to the HNG-X data centres from defined WANs. These include the C&W network, External networks from clients such as LINK and Card Account, Streamline connections and all support connectivity for appliances such as the EMC array.
- Inter data centre traffic flows
- Firewall services (using different Firewall technology to that in the Distribution Tier). Refer to section 2.3.5.5 for further details of Firewalls.
- IPSEC and SSL Termination
- ACE blade for Load Balancing, SSL Termination and Virtualisation services
- IDS / IPS services
- Radius platforms used to authenticate access from the Branch network
- Service Boundary where this is local to the data centre
- Internet Access

2.3.5.1 Access classes

The traffic flowing into and out of the HNG-X date centres falls into the types (classes) enumerated in the following table. Note that Access class refers to the method of presentation at the HNG-X data centres. The Traffic steering attribute indicates that this access class is "steered "to the appropriate HNG-X data centre using Layer 3 Routing protocols.

Access Class	Traffic steering	Overview / Examples
MPLS	Yes	MPLS VPNs carrying traffic between HNG-X data centres and the following sources. Branch traffic (ADSL, Dialled, GPRS), A&L, DVLA, E-PAY, Post Office & Support (including external Business applications such as OBC). For "Dialled" the Service boundary is local to HNG-X data centres, in all other cases the service boundary is remote to HNG-X data centres. Radius Traffic from Branch Network Service Providers Separate (from Live) MPLS VPNs (refer to section 6.2.3) are used for Test purposes.
Third party network service	Yes	The third party provides network services into HNG-X data centres. The service boundary is local to HNGX data centre. For example Link, Moneygram and Card Account These network services provide for Live and Test traffic with separation as specified in section 6.2.3. Need to cover dual running as well
Circuit based Network service	No	In this case a circuit based service is provisioned into the HNG-X data centres. Traffic steering is not applicable since all communication sessions are initiated outbound from HNG-X data centres. Separate physical circuits (refer to section 6.2.3) are provisioned for Test purposes. For example Streamline, X.25 and ISDN
Inter Data Centre	N/A	Resilience for MPLS access is provided by having one physical circuit into each data centre and triangulating between them. Also this service is used to allow active / active services to communicate (for example during dual running there are horizon active / active applications and for HNGX active / active services.. This provides the ability to run tests which require traffic to travel intercampus. This avoids having to use fibre delay simulators for either SAN extension or Data test cases. Refer to section 2.3.5.6 for further details.

Access Class	Traffic steering	Overview / Examples
Branch ISDN Dial out	No	To initiate a call to an ISDN branch. This is done in order to establish communications to the branch. The Branch router will reject the incoming call but dial back in. Traffic steering is not applicable since all communication sessions are initiated outbound from HNG-X data centres.
Support PSTN dial out	No	To provide an out-of-band management route to access the Branch Router over PSTN. Traffic steering is not applicable since all communication sessions are initiated outbound from HNG-X data centres.

Table 1 – Access class

All ingress methods to the Data Centre are based on Routing – specifically for any Access class which has attribute traffic steering set, it is possible to steer the traffic to either data centre via IP Routing. Note that test traffic will only be steered to the Secondary data centre whilst it is Standby state (refer to section 2.3.1).

2.3.5.2 Radius Services

2.3.5.2.1 Overview and Clustering

These are virtualised with a single VIP for each Radius function (e.g. ADSL authentication).

This approach is taken for the following reasons;

- Abstraction of the Radius service from number of underlying platforms and location of these platforms.
- Avoiding reliance on behaviour of multiple Radius client types to provide high availability Authentication services to the Branch Router; if the Service provider were to configure multiple Radius servers on their Client then there would need to be at least one per data centre. In the event of DR where the Primary Data Centre has failed, there may be periods during which the service provider client is trying to use Radius servers that do not exist causing degradation of service to post office branches. Note there is no out of band probing in the Radius protocol therefore any resilience protocol is likely to send auth requests to non existent servers eventually before failing over to a functioning Radius. Depending on the resilience protocol within the Radius Client, this may periodically repeat in normal operation and following restart of the Radius client. In addition with multiple Radius client types this behaviour could be significantly different for each client type resulting in multiple behaviour patterns. For example, within Horizon it was observed that the Radius client behaviour changed significantly in a version of the Cisco operating system.
- To provide determinism as to which Radius server is used for Authentication requests

The ACE blade in the Access Tier switch is used to provide Load balancing and Virtualisation. To provide deterministic behaviour, Radius Authentication and Accounting traffic from the Radius Client will be steered to the Active data centre. As far as the Service provider is concerned the HNG-X system provides one Radius server only. Therefore all the Horizon Radius servers for a given function will have the same shared secret for authentication to the service provider Radius client.

The general deployment model for creating a Radius Cluster is shown in the following diagram.

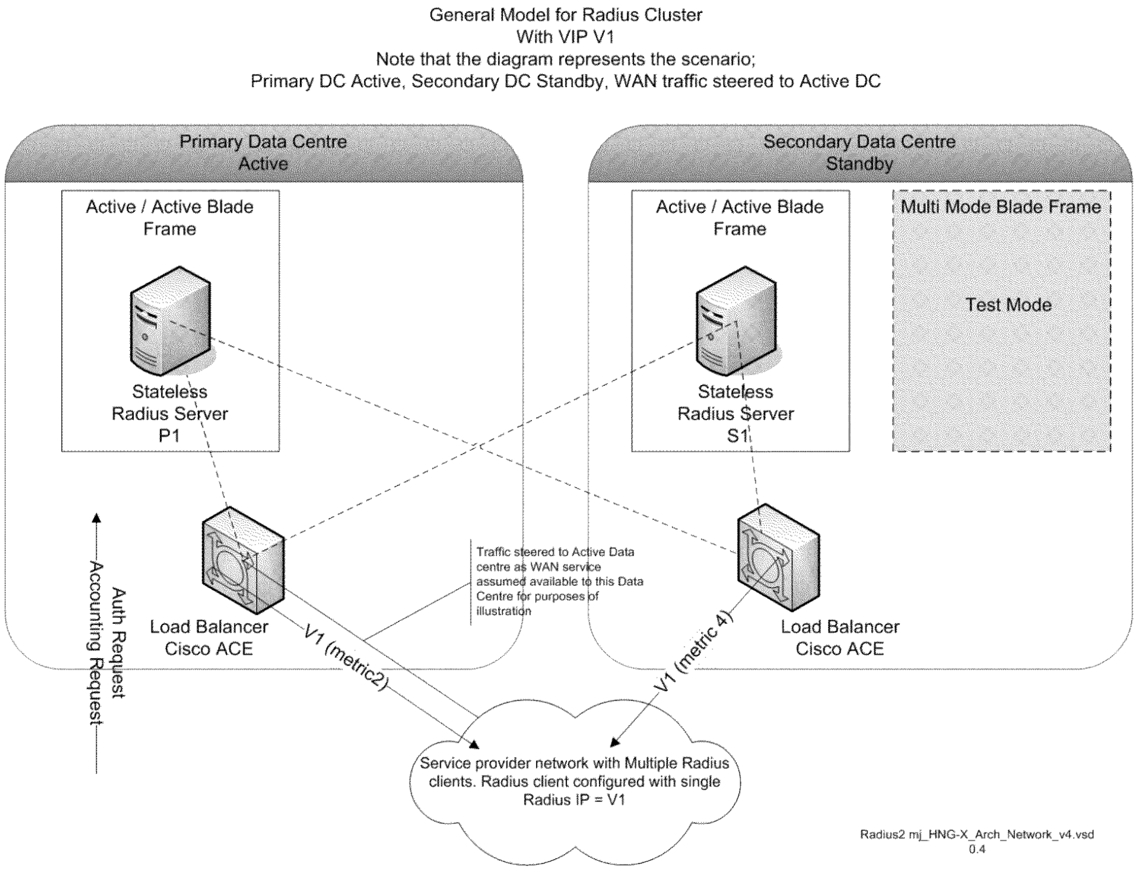
General Model for Radius Cluster
With VIP V1
Note that the diagram represents the scenario;
Primary DC Active, Secondary DC Standby, WAN traffic steered to Active DC
Primary Data Centre
Active
Secondary Data Centre
Standby
Active / Active Blade Frame
Stateless Radius Server P1
Active / Active Blade Frame
Stateless Radius Server S1
Multi Mode Blade Frame
Test Mode
Load Balancer Cisco ACE
Load Balancer Cisco ACE
Auth Request
Accounting Request
Traffic steered to Active Data centre as WAN service assumed available to this Data Centre for purposes of illustration
V1 (metric2)
V1 (metric 4)
Service provider network with Multiple Radius clients. Radius client configured with single Radius IP = V1
Radius2 mj_HNG-X_Arch_Network_v4.vsd
0.4

Figure 5 Radius Cluster

In general a Cluster consists of 2 Radius servers which are instantiated 1 per data centre as shown above.

- The Servers are ordered based on locality so in the event of WAN failure to the Primary data centre S1 will be used if available.
- In the event of a DR transition that is long term (refer to section 2.3.1), an additional Radius server may be instantiated in the Multi Mode Blade to provide two servers for resilience. This step is not part of the solution.

2.3.5.2.2 Radius Functions

The following table provides a summary of Radius functions, how these are located externally and how they are instantiated on platforms.

Radius Function	Overview	Product
Branch PPP Auth ADSL/ ISDN / PSTN Dial out profiles	Follows general Cluster deployment model, one cluster for ADSL and one for dial in / dial out.	Radiator
Boot server	For provisioning of Branch Routers	Radiator

Branch PPP Auth - WWAN	Follows general Cluster deployment model If Cisco ACS were to be used then it has been concluded that there are two clusters required (A for GGSN1 and B for GGSN2). Basically the required behaviour from the WWAN service provider is that for each GGSN nodes in their network, the Radius allocate non overlapping address blocks for the same PPP username. It is anticipated that the Flexibility within the Radiator product will mean that one Radius cluster is sufficient. Note however that the Radius servers for WWAN need to be separate from other radius servers as this is a security requirement..	Radiator
Branch PPP Accounting	Follows general Cluster deployment model – one cluster for all accounting.	Radiator
Network Device Access control and command audit	One Platform instance per data centre, these are free standing platforms. The IP address is local to the data centre. The rational for this is that services provided by this platform are fundamental to a data centre and need to be available early in the data centre commissioning process. Each network device is configured with the IP addresses of both these Radius servers.	Cisco ACS rationale is TACACS+ support
Branch Router Access Control	Follows general Cluster deployment model, one cluster	Radiator

The Rationale for using Radiator in the Boot domain is the flexibility in invoking scripts throughout the processing path. It is chosen for use in the production network due to its ability to handling real time updates and having enough flexibility to reduce the number of radius services for the Wireless WAN network from 2 per data centre to one per data centre. It is also relevant the Fujitsu services have a corporate license for Radiator which avoids RMGA having to pay any licensing costs.

*Note that the Active Blade frame is continually active in both data centres and not subject to any transition associated with DR.

2.3.5.2.3 User database maintenance

In the case Radius services which provide PPP auth functions for branches it is necessary to establish a user database. The following approach will be used;

1. Bootstrap: When a Radius server platform is booted, it detects whether it has a user database current at start of day. If not it requests via supported interfaces within Branch Estate Management a full report of all PPP users for its given technology type. It uses this report to populate the internal database using the RCAP* processing model.
2. Daily Update: Every morning at 06:00am the Radius server executes a scheduled task. Note that scheduling is performed externally (i.e. it is not a function of Radius platform). This task will use the same view as above to fetch relevant user data. RCAP will perform a delta update to the internal database.

3. Near Real Time Updates: To deal with Branch Routers that are Registered / Deregistered throughout the day, it is necessary for real time updates to be applied to the user database. The case of Registrations could be handled by simply using the Auth fail processing path to check / apply updates. However Deregistration, for example in the case of stolen Router requires that PPP credentials are revoked. For this reason the approach taken is for an update task on the Radius platform to poll for changes (say every 10 minutes). In addition an Auth Fail will cause an immediate poll for changes provided such a poll has not happened in the last 1 minute. Given that the Branch Router will retry, subsequent auth requests received after the period of time to apply the update (which should be less than 1 minute from Auth Fail) should succeed. The interface used within Branch Estate Management will have properties which enable changes since last call to be retrieved. To deal with the case where the polling process has crashed, it will request a full view and perform a delta operation. This is deemed good enough as the occurrence of the process disappearing whether through a crash or through the platform failing should not be frequent (< 10 per year).

* RCAP is utility that provides functionality for;

- Encryption (with obfuscated key) of Chap secrets allowing them to be passed around the network in clear.
- Deriving delta updates to a Radius user database based on current data base and full report of all Branches from Branch Estate management

2.3.5.2.4 State properties

This section is concerned with Radius functions for Branch PPP only.

These Radius servers are stateless in that the solution does not require a Radius packet Ri (Auth request / Accounting Request) to be targeted at particular radius server based on where previous Radius packets associated in some way with Ri had been delivered.

However the load balancer behaviour will be to prefer the Radius server in the Active Data centre. This will ensure most of the time auth requests and accounting requests will be sent to the Radius server at the Active Data centre. This simplifies diagnostics and support staff an enables offline session tracking by analysis of accounting logs.

Note that periodic accounting is provided for all access network types.

Also note that support for real time session tracking is not provided in the solution. For information purposes should this have been a requirement then it would have been provided by use of a Single Database cluster implemented across both data centres with every Radius server providing updates.

2.3.5.3 Specification of service boundary

Contractual interface specifications exist for {A&L, Card Account, LINK, DVLA, E-PAY, Post Office, Moneygram and Streamline}. These cover both live and test traffic flows.

Mitigation and incident management are covered in SLA and OLA's.

These specifications will be updated to reflect changes, amongst these being data centre location and operational mode (active / standby instead of active/active).

Design documents will cover Interface specifications for Branch and support traffic. These will cover both live and test traffic flows.

2.3.5.4 Access Tier multilayer switches

These are the same models as used in the Distribution Tier (refer to section 2.3.4.1). An ACE blade is used to provide SSL termination, load balancing and virtualisation services.

2.3.5.5 Access Tier Firewalls

These consist of appliance devices arranged in Active / standby clusters. Each cluster is local to a data centre.

The general approach is to use virtualisation so that a single appliance creates multiple virtual firewalls. A Virtual firewall has the following characteristics;

- Ingress and Egress of traffic is via VLAN's.
- Configuration of Firewall rules can be performed without regard for the rule base of other Virtual Firewalls.

Each appliance has a single dedicated management interface (Ethernet port).

For internet Access dedicated Firewalls are used, refer to (section 2.5.4).

It is not sensible to use a single cluster to provide all Access Tier Firewalls as failure of this cluster, caused by a single failure in the clustering protocol may result in all traffic flows being cut off from the date centre. To avoid this scenario where support access is cut off concurrently with any business access, a dedicated cluster is used for support traffic.

Conversely providing a firewall cluster per 3rd party organisation would result in a large number of Firewall appliances initially with no cost effective scaling.

Additionally Branch traffic is very different from other traffic due to a large number of concurrent sessions with a potentially large session churn. Therefore a dedicated cluster is used for Branch traffic.

The chosen approach is to provide four Firewall clusters; Branch, Client, Post Office & Support.

Note that any Internet traffic will traverse an External non Cisco (to HNG-X firewall) within the Fujitsu service that is used to provide this access (refer to section 6.3 for further details).

The Firewall technology for these Clusters is based on Cisco ASA 5500 Series Adaptive Security Appliances. The chosen model and its main characteristics are listed in AB.2.

The firewalls will log all events to the Logging event server – refer to section 3.1.

2.3.5.6 Inter Data Centre Network

The WAN service between the two HNG-X data centres carries IP traffic and FC SAN traffic. It is based on a DWDM service. This service needs to be highly resilient since it is used to replicate state which is required in the event of DR. There are two possible outcomes of this service failing;

a) Either, the Active data centre continues to process transactions but these are not recorded at the standby data centre.

b) Or, the Active data centre blocks transactions until the WAN service is restored.

To ensure the service is highly resilient it has the following characteristics;

c) There are two DWDM devices each data centre and the SAN extension and Network topology is such that it is sufficient for a single device to function to provide an InterCampus service.

d) Between both HNG-X data centres there are a pair of fibre optic cables. The radial distance of each of these is < 100 km and the two fibres are kept separate along their runs with no common interconnection points. The DWDM devices are used to provide data service over the fibres and in the case of Network traffic these are:

Inter Campus Services

- Access Tier: Two times 1 Gigabit Ethernet connected to Access Tier switches
- Core Tier: Two times 1 Gigabit Ethernet connected to Core Tier switches

The following diagram illustrates the approach to providing inter data centre services.

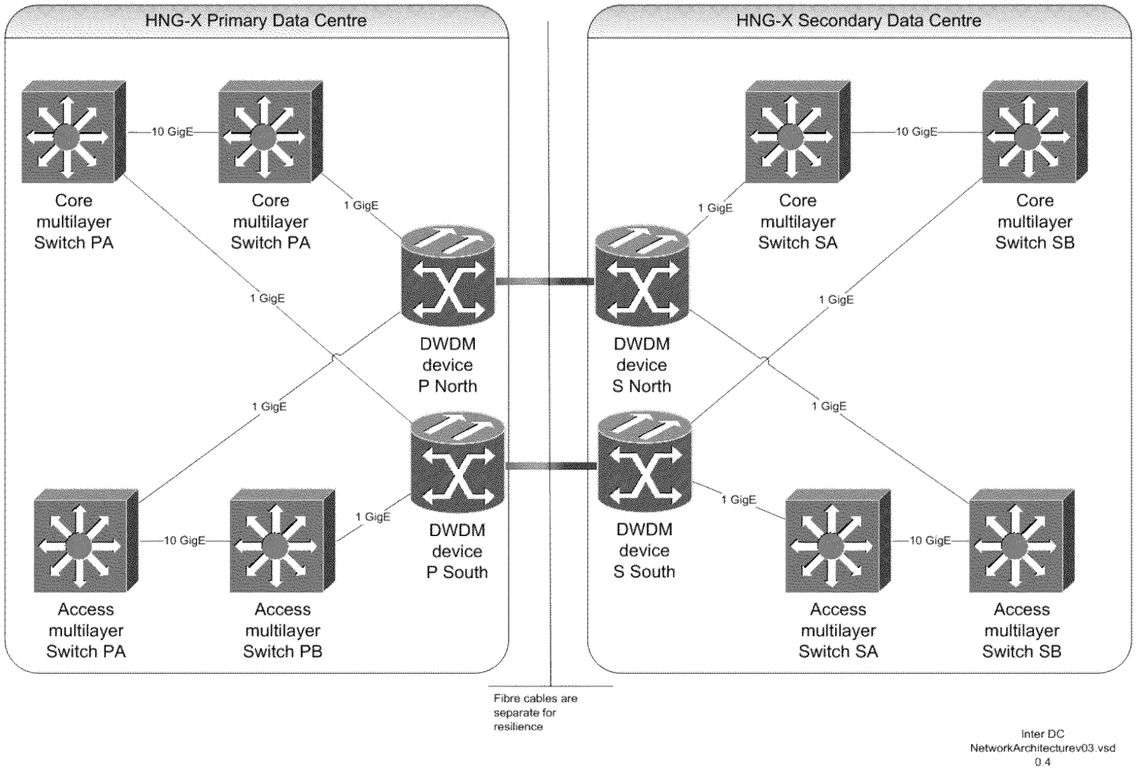
HNG-X Primary Data Centre
HNG-X Secondary Data Centre
Core multilayer Switch PA
Core multilayer Switch PA
10 GigE
1 GigE
DWDM device P North
DWDM device S North
Core multilayer Switch SA
Core multilayer Switch SB
10 GigE
1 GigE
1 GigE
1 GigE
1 GigE
1 GigE
1 GigE
1 GigE
DWDM device P South
DWDM device S South
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

Figure 6 – Inter Data Centre IP services

Note that Per VLAN spanning tree is used to create resilient intercampus VLANs providing transparent Layer 2 connectivity between both data centres.

Note that the DWDM device is provided as a managed service by FTEL – it is not a HNG-X network appliance. An impact of this from a security perspective in that it is not possible to assert (without FTEL agreement) that the inter campus traffic could not be captured..

DN: Need to obtain statement from FTEL about why packet capture is technically not possible

- Individual traffic flows across the DWDM devices are deterministic. However all Intercampus services are used in normal operation to provide active monitoring with realistic traffic flows.

Note the drawback with limited synthetic monitoring of the standby service is that the monitoring traffic may not contain sufficient traffic patterns to sufficiently test the network service {packet size distribution, fragmentation patterns and data patterns}

2.3.5.7 MPLS Termination

The main characteristics of this access class are;

- One Physical service into each data centre with triangulation between data centres for resilience to failure of physical circuit or CE Router
- IP Routing between CE Routers and MPLS Hand-off Routers for resilience
- Loss of Access multilayer switch at any data centre does not cause loss of service. In particular TCP connections from applications will typically survive such a failure.
- Production and Test Traffic classed steered to optimal data centre – triangulated route only used in failure cases..
- In the event that DR is declared then as part of the Secondary Data centre transition from Standby to Active protocol (refer to section 2.3.1 for definition of this transition) then;
    - All Test MPLS VPNS will be blocked (via ACL on MPLS handoff Router) and Access multilayer switch.
    - Preferred Traffic path for Production traffic will be to Secondary (Active) Data Centre

The following diagram illustrates the approach to MPLS termination.

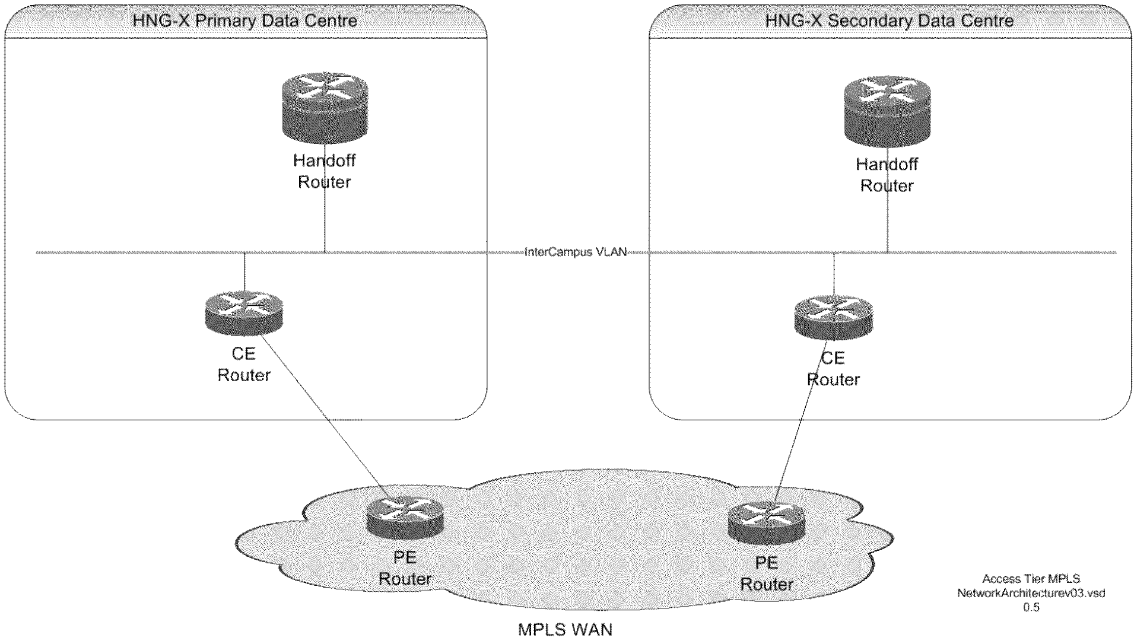

Figure 7 – MPLS access

Handoff Routers exists to terminate WAN services for particular traffic classes / instances as documented below.

Access Type / Instance	Traffic Class	Comment
Branch	MPLS	These Handoff routers participate in routing with other Routers within the Data centre – for example Horizon VPN array.
FSBN (DN:if used?)	MPLS	This is the CE Router
Client / A&L	MPLS	Contractual Interface specified by TIS
Client / CAPO	3rd Party	Contractual Interface specified by TIS
Client / DCS	Circuit based Network service	Contractual Interface specified by TIS
Client / DVLA	MPLS	Contractual Interface specified by TIS
Client / Link	3rd Party	Contractual Interface specified by TIS
Client / Moneygram	3rd Party	Contractual Interface specified by TIS
POA / Huthwaite	MPLS	Contractual Interface specified by TIS
Support / Ste04	MPLS	

2.3.5.8 Third party network service

Under this model, the 3rd party provides a network service into HNG-X data centres. The preferred model is that the 3rd party provides two Routers with diverse(*) incoming WAN services per HNG-X data centre connected onto a VLAN providing transparent Layer 2 connectivity between both data centres. At each Data Centre there is a pair of hand-off routers connected into different Access Tier switches.

Under this model;

- There are no single points of failure even with a single Active Data centre
- Service boundary is located on the Ethernet ports of the 3rd party Routers
- Routing would be used to achieve resilience both locally and across data centres
- A peering IP address space will be defined
- The HNG-X hand-off routers would be connected to Access Tier Firewalls

* Labelling the wan circuits to the Primary data centre as Pw1, Pw2 and secondary as Sw1, Sw2, then it is sufficient that diversity exist for the pairs {(Pw1,Pw2), (Sw1,Sw2),}.

2.3.5.9 Branch Access

This section provides an overview of the Branch domain within the Access Tier. Note that the general model is to have a single service provider – FJ Core services terminating all WAN network types with a common delivery for all traffic to the HNGX data centres..

Note that the Aggregation Routers in HNG-X data centre Peer with some Connect DSL routers for two-way exchange of routing information. This will be summary routes rather than per PPP session host routes and covers all network types.

Service type	Characteristics
ADSL (BT IP Stream service)	• Delivered over MPLS (refer to Figure 999108 – Wide Area Network)/ • PPP session forwarded by BT to LNS in FJ Core Service Provider cloud • Periodic Radius accounting to track user sessions • Resilient interconnect - no Single points of failure • Radius servers in HNGX data centres for PPP Authentication and PPP Accounting • Framed route to the LAN side of the Branch Router inserted into LNS by HNGX Radius servers and advertised • Session steering based on PPP CHAP username (domain part) • Test / Live separation via Session steering
Dial up (ISDN / PSTN)	• Delivered over MPLS (refer to Figure 999108 – Wide Area Network) • Service provider is C&W • Periodic Radius accounting to track user sessions • Dial up session terminated on Access Server in C&W cloud, PPP session forwarded to C&W LNS located at SDC01 / TCY02 • Framed route to the LAN side of the Branch Router inserted into LNS by HNGX Radius servers and advertised • Same interconnect as for ADSL • Radius servers as per ADSL model • Session steering via Dialled number string • Test / Live separation via session steering
Dial – out (from Data Centre) This is to "prod" a dial on demand Branch to connect in	• Gateway of last resort when no /32 available to either ADSL, ISDN or WWAN. Note that /32 routes are not required at HNGX data centres. • Dial out (for ISDN prompt) provided using Dial Out Router in FJ Core Service Provider cloud • Radius profile used for dial out – Radius server at HNGX data centre (DN: Need to confirm use of Radius profiles will provide necessary functionality)

Service type	Characteristics
GPRS / EDGE / 3G	• Delivered over MPLS (refer to Figure 999108 – Wide Area Network) • PPP session terminated in WWAN Service provider cloud • Tunnelled PPP session (see next bullet) terminated on LNS in FJ Core Service Provider cloud • L2TP from Branch router to Data Centre Routers in order to provide for Branch LAN direct access and address preservation. Note Service provider does not provide Framed Routes. • Framed route to the LAN side of the Branch Router inserted into LNS by HNGX Radius servers and advertised • GPRS hand-off Routers exchange two-way routing information (summary only) to provide resilience. • Same interconnect as for ADSL • Radius servers as per ADSL model but note there are two PPP layers (outer / inner) • Session steering based on @domain to determine which VPN and therefore VRF to terminate the connection in. • Test / Live separation based on APN's. SIM APN determines whether Live or Test access • GPRS network protection against overload in the event of widespread (. 10%) of branch outage DN: Need to provde ref to CD where the actual percentage is documented.
VSAT (Broadband)	• Uses HNGX Branch Router (replaces one supplied and managed by FJ Core Services) • Delivered over MPLS (refer to Figure 999108 – Wide Area Network) DN: Need to determine if PPP or pseudo wire model
Bootstrap Dialup - Same as production – separate MPLS VPN ADSL;- Same as production – separate MPLS VPN	• Delivered over MPLS (refer to Figure 999108 – Wide Area Network) • Dedicated MPLS VPN (one for Dial and one for ADSL) for Router provisioning purposes • Dedicated Radius (using Radiator Software)
Support dial out	• 2 * PSTN modems per Data Centre

Table 2 – Branch Network Services

2.3.5.10 Sarian VPN Concentrators

These are used to provide encryption for Management traffic between Branch Routers and the Data Centre. There are two VPN Concentrators (Sarian VC5100) per data centre in an active / standy arrangement. They each support over 14,000 Remote uses with a concurrency limit of about 200 active IPSEC sessions. IPSEC Tunnels can be initiated in either direction and will be setup on demand based on traffic flows. The configuration for the IPSEC is obtained from an SQL database (MYSQL). DN: need to find a platform to host this on.

A process similar to radius for (see section 2.3.5.2.3) will be used for Bootstrap and daily update. Since IPSEC pre shared keys are branch specific there is no near real time update requirement. The rationale for this is that the access to the data centre is already protected using PPP CHAP.

In addition the out of band PSTN management solution uses VPN concentrators (2 per data centre, active / standby). These have PSTN modems attached and are PSTN calls are always initiated from the data centre. The usage of this support route is manual with an engineer at the Branch and a support person being involed.

The VPN concentrators are configured on demand fetching details from the Branch Estate management system via using automated interfaces with the option to apply manual correction to the phone number. The Branch which is dialled may have a Personalised Router or a Gold Build Router and the configuration details will be fetched accordingly.

Note that a set of IP addresses are dedicated to PSTN dial out and the set of interface addresses for the Router can be summarised for ease of data centre configuration.

2.3.5.11 Branch Traffic patterns

Business Applications

Bespoke business application written in Java and communicating to the Data Centre systems via an http based protocol.

The Counter is the HTTP Client and data centre is HTTP server.

The business applications of data centre communications fall into tow categories

- Communications within a User Session
- Communications outside a user session context

The Session based communications comprises

- User Authentication – logon process to establish the session
- Settlement transactions (single communication exchange at the end of a customer session to commit the transactions to the branch database).
- Online transactions (Banking authorisation, DVLA etc.)
- Reporting (branch stock and accounts management, ad hoc queries etc.)
- Back office business transactions (such as Remitting in stock and cash, adjustments, cash declarations etc.)
- Audit Events within sessions (e.g. label printing event when voucher / mails label printed)
- Administration functions (create stock unit, create user, change password etc)
- Viewing Messages broadcast to the branch (plus polling for status changes)
- Training simulation for CTO branches
- Help system access

The Non Session context communications comprises

- Daily (typically overnight) Reference data changes via branch database (note only a subset is delivered through this route)

- Polling throughout the day for dynamic reference data changes (e.g. bureau spot rates data)
- Help system access outside of session (help pages limited to non session based help data)

Infrastructure Applications

These facilities are based around a standard systems management package (e.g. Tivoli) plus bespoke integration packages

Software distribution is mainly initiated from inbound TCP connections from the Counter however the ability to push software from the data centre also exists.

- Inventory management
- Software distribution
- Common reference data distribution (this is the bulk of reference data in volume terms – typically delivered overnight)
- Alert forwarding (based on both OS and business application events)
- Support access
- Diagnostic file retrieval
- Configuration setting

Some of the above actions include local communication within the branch (e.g. to cache large software download packages)

Network management communications

Under HNG-X these will typically terminate at the branch router rather than the counters

- Session maintenance (e.g. PPP and L2TP heartbeats)
- Network monitoring
- Support access
- Diagnostic data retrieval

2.3.5.12 HNG-X Bootstrap domain

In addition, there are special communications during the rolling out of a counter and branch router to establish its personality and become a managed object within the estate.
The HNG-X bootstrap domain is isolated from the production network within the data centre using Firewalls. It is also separate on the WAN using separate MPLS VPN's. Separation is maintained within the data centre VLANS. The MPLS VPN selection is based on Dialled number string or the domain name part of the HNGX username.

Architecturally the Boot loader service application components (including radiator scripts) are part of Branch Estate Management

2.3.5.13 Route Hunting

The LAN address of the Branch is constant across WAN network types and it is possible for multiple WAN interfaces to be up at once. This occurs during failover when for example during fallback to ADSL, the interface comes up and may be tested for IP flow. It also occurs during the proposed method for testing WWAN viability (refer to section 5.1.2.1). Therefore it is necessary to deal with the case where multiple potential Routes exists to the Branch. There is no routing protocol running between the Branch and the data centre. The easons for this are:

Scaling and stability challenge; Number of remote peers and stability over GPRS in particular

Optimise for the usual case; ADSL with possible WWAN backup

The approach taken is define a deterministic ordering across WAN Types both at the data centre an the Branch Router. Occasional short term asynchronous path selection should not be service affecting. The proposed ordering at the data centre is [ADSL, Dialled, WWAN, Dial out] at the branch Router it is [ADSL, Dialled, WWAN].

Note that for outbound support access a dedicated router will be used based on policy routing.

2.3.6 Topology Overview

The following diagram provides a simplified summary of the LAN topology. For details of flow constraints from security refer to section 6.

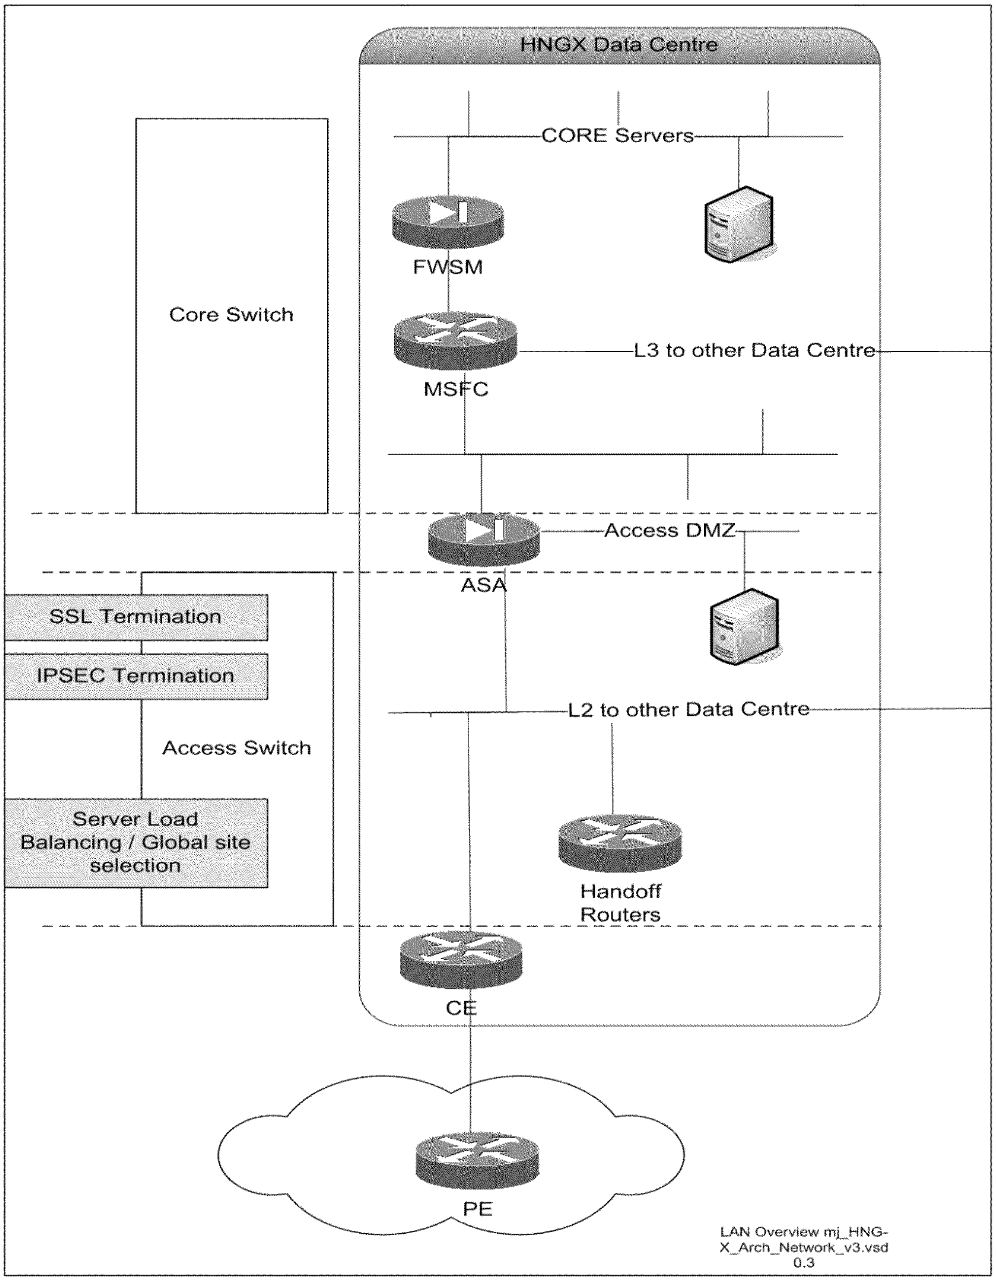

Figure 8 LAN Topology Overview

2.4 Storage Area Network

2.4.1 SAN

The SAN connectivity for HNG-x will be implemented using Cisco MDS9509 Director Class fibre channel switches in order to take advantage of the added availability and stability associated with this class of hardware.

There will be a two SAN Directors deployed at each HNG-X data centre. Each Director chassis will be deployed on physically separate racks. Both power and FC cabling will be fully resilient.

The FC fabrics deployed will use industry standard Fibre channel Protocol (FC-SW) at 1, 2 or 4 Gbps, depending on the storage arrays selected.

In order to achieve the highest degree of uptime, the SAN fabrics at each site will be deployed in pairs, each fully isolated from each other. Inter Data Centre site isolation will be achieved by means of intra fabric routing. Similarly connectivity to FC attached backup devices will be isolated by use of FC routing.

Depending on the vendor selected, FC routing is expected to be achieved through industry standard FC routing as implemented by Brocade, or by the Cisco proprietary VSAN approach.

2.4.2 Storage Area Network extension between data centres

Inter Data Centre connectivity will be provided in accordance with a service provide by FTEL (Fujitsu telecommunications) based on a resilient Dark Fibre pair provided by NTL.

Section 2.3.5.6 covers the use of these DWDM devices to provide services for IP traffic.

The number and specifications of the inter-site links for storage traffic are 4x FC (4 G bit/s) links, 2 per DWDM device deployed over 2 physically separated paths.

2.5 Wide Area Network

The functions of the Wide Area Network are to provide;

- Network Connectivity to Clients, Post Office Data Centres, Support sites and Test locations
- Remote Client LAN where a Remote service boundary exists (DVLA, A&L, E-PAY, Post Office)

The following diagram shows the scope of the WAN network and puts it into the context of Data Centres, Service providers and Access tier network services.

(DN: Need to confirm approach where ? is shown)

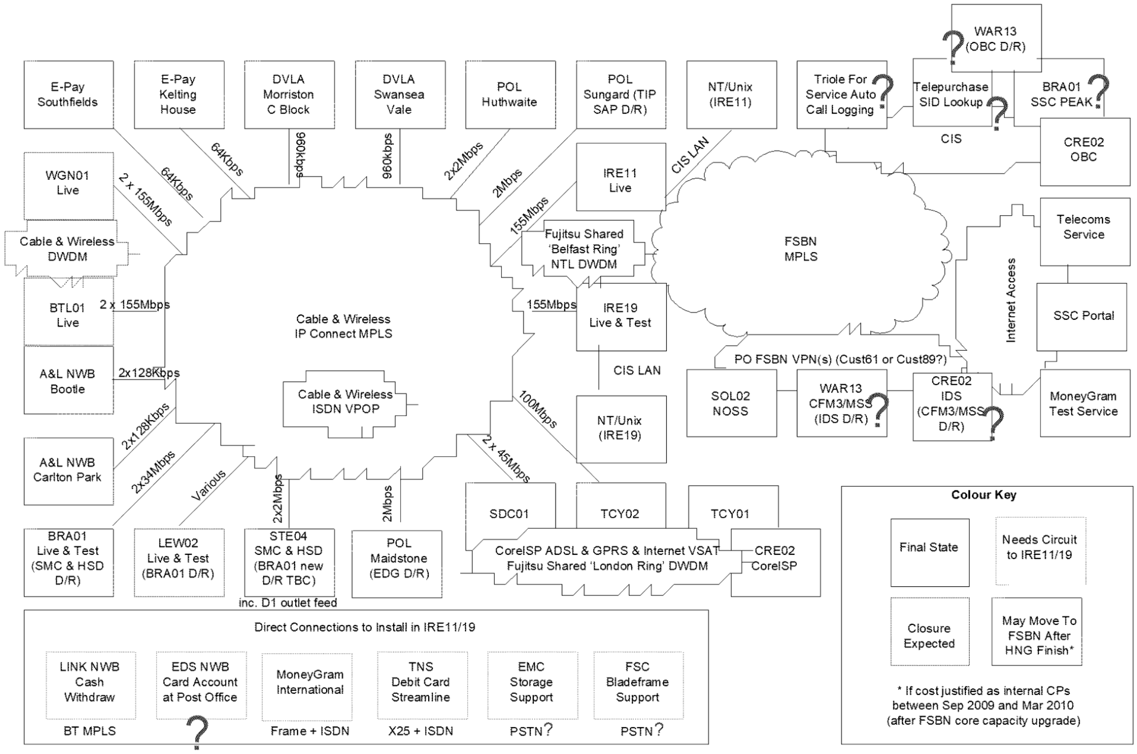

Version 0.5
October 2007
Acknowledgement to Alex Kemp for producing original

Figure 9 – Wide Area Network Target solution

2.5.1 Overview

The WAN Network delivery into the HNG-X data centres is based on MPLS terminating on CE Routers (refer to section 2.3.5.7 for further details). The principles for traffic separation across and within MPLS VPNs are specified in

Table 78 – Traffic Separation).

Remote Client LAN

The MPLS VPN is terminated at the remote site on a CE Router. The Service boundary is at this site and it exists at the Ethernet port of a HNG-X Router. The overall service is resilient to single points of failure and this is achieved by connectivity into two sites (or one site with separation between termination points) per remote client. IP routing is used to provide resilience. The MPLS VPN will carry encrypted traffic if this is specified in the relevant TIS. For example connections to DVLA, A&L, E-PAY and Post Office) follow this model with A&L, E-PAY and Post Office traffic streams being encrypted.

Point of Presence interconnect

This is similar to the remote Client LAN model but used where traffic from multiple MPLS VPNs flows across the service boundary. In this case the traffic separation is preserved by one to one mapping between VLANs over the service boundary and MPLS VPNs. This model applies to interconnection with Connect DSL for ADSL traffic and Orange for GPRS traffic.

2.5.2 Bandwidth Requirements

This section provides an initial view of the connectivity and bandwidth for the Target solution into and between the HNG-X data centres. It will be refined during the design phases.

Live Circuits:

#	Description	'A' End	'B' End	Speed	Comment
L1	Intercampus – SAN	IRE11	IRE19	n G bit/s Fibre Channel	n is 1,2 or 4 specify which
L2	Intercampus – SAN	IRE11	IRE19	n G bit/s Fibre Channel	Needs to be diverse and separate from L1
L3	Intercampus – Network	IRE11	IRE19	1 G bit/s IP	
L4	Intercampus – Network	IRE11	IRE19	1 G bit/s IP	Needs to be diverse and separate from L3
L5	3rd line support (SSC)	IRE11/IRE19	BRA01	4 M bits/s, resilient	Diversity and Separate routing required to provide high resilience
L6	2nd Line Support	IRE11/IRE19	STE09	2 M bits/s resilient	Diversity and Separate routing required to provide high resilience
L7	OBC	IRE11/IRE19	CRE02	2 M bit/s resilient	Diversity and Separate routing required to provide high resilience
L8	Ops	IRE11/IRE19	IRE11	2 M bit/s resilient	Diversity and Separate routing required to provide high resilience
L9	DR – 3rd line (SSC)	IRE11/IRE19	LEW02	2 M bit/s resilient	Diversity and Separate routing required providing high resilience. Not normally used.
L10	DR – Ops	IRE11/IRE19	IRE19	2 M bit/s resilient	Diversity and Separate routing required providing high resilience. Not normally used.
L11	3rd line (MSS)	IRE11/IRE19	???	2 M bit/s resilient	Diversity and Separate routing required to provide high resilience
L12	Branch Traffic (All Branches but need to separate by Service	Tele city/ SDC01	IRE11/IRE19	70 M bit/s resilient	Diversity and Separate routing required to

#	Description	'A' End	'B' End	Speed	Comment
	provider?)				provide high resilience
L13	IP Select via IP gateway	Tele city/ SDC01	IRE11/IRE19	10 M bit/s resilient	Diversity and Separate routing required to provide high resilience

Note that Tele city is one of the locations where the IP Gateway is delivered the other is SDC01.

Test Circuits

#	Description	'A' End	'B' End	Speed	Comment
T3	Test Access	IRE11/IRE19	BRA01	34Mbit/s with 8Mbit/s backup.	Further work needed on bandwidth. Highly unlikely to need 34Mbit/s for much of the time.
T4	Test Access – DR	IRE11/IRE19	LEW02	2Mbit/s no resilience	No resilience required.

Note that T3 is assumed to provide test access for groups not based in BRA01 (e.g. MSS) and they will use the Fujitsu Services intranet to get to BRA01. Bandwidth for this should be low.

2.5.3 Remote Access by Fujitsu services users and systems

This deals with how users and / or systems within Fujitsu services access to the HNG_X Network. In general there are three models; .

- The Red LAN model is effectively an extension / creation of a HNG-X LAN in a remote location with defined traffic flows permitted between HNGX platforms located in the HNG-X data centres and workstations on the remote LAN. Specifically under this model the end point on the remote LAN is an RMGA build workstation
- The Corporate Gateway model where all traffic between the Fujitsu Corporate network and the HNG-X network passes through a proxy server that breaks flows at layer4.
- The Secure Access server model where all incoming traffic from Corporate terminates on the SAS server – effectively acting as an Application Level firewall. The incoming traffic is always RDP (version ??).

The rationale for enabling corporate access is to enable users to use their corporate workstation to provide controlled access the HNG-X systems.

The following diagram (acknowledgements to Ian Bowen for the original version) illustrates these models.

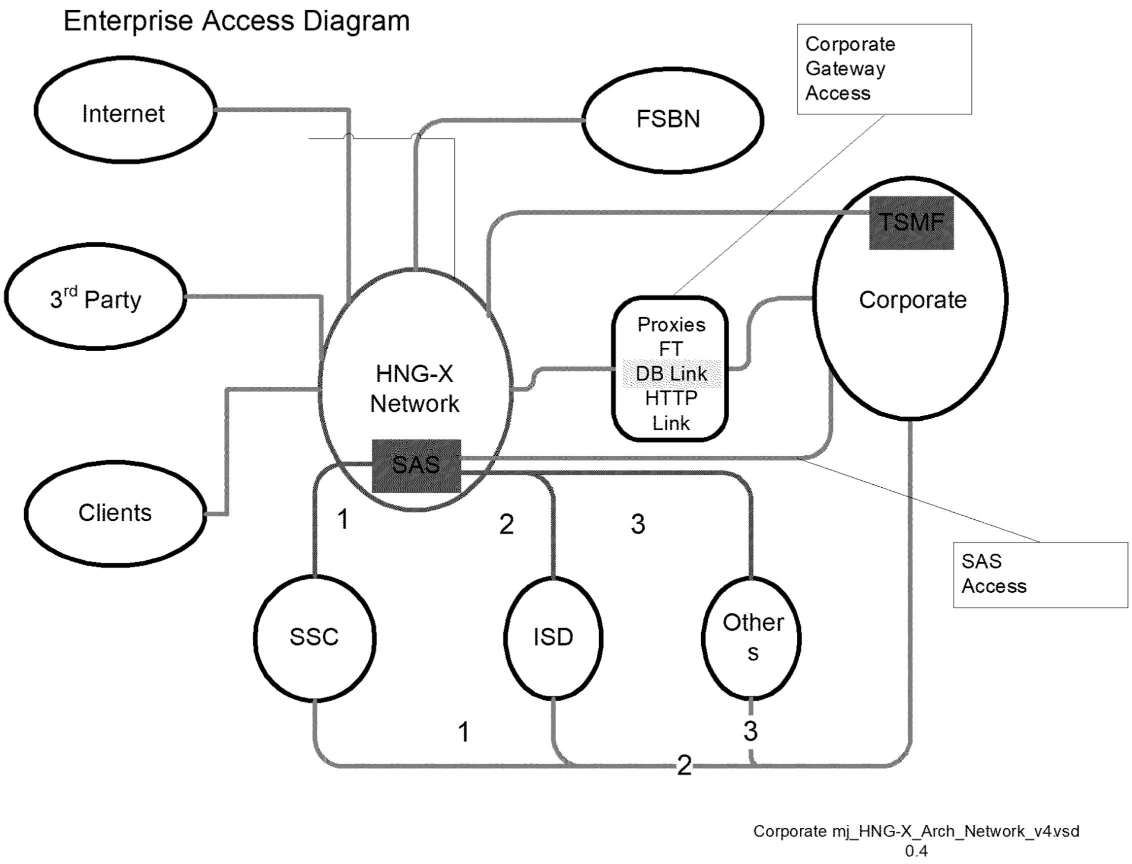

(DN: TSMF via proxy? Also remove FSBN)

The RED lines above outline the secure Campus Lan and show where it may extend beyond Data Centre bounds.

The lines marked 1, 2 and 3 show the possible routes for workstation access:

1. Direct Campus access via extended Red Lan
2. Access via Red Lan to SAS
3. Access via corporate to SAS

The definitive mapping of users / systems together with to Remote Access Models will be specified in the WAN HLD. This mapping will also cover bandwidth and resilience requirements.

For the purposes of network connectivity, Corporate are treated like any other 3rd party and are characterised as follows;

1. A number of interconnect sites, currently envisaged as Bracknell, Stevenage, Ireland and Lewes

2. Handoff model at each Interconnect – envisaged as Firewall to Firewall but specified in the Transit LAN HLD.
3. In the case of interactive users, characterised by ad hoc access and where it is acceptable to deal with prolonged network outage by moving to another site or using the OOH laptop solution then conveyance of that traffic over the CIS WAN to a Handoff point is possible.
4. In all other cases the traffic will be conveyed by CIS locally to the HNGX handoff point. It is necessary to agree with CIS that the LAN segment will be connected to a switch and from the switch to the Handoff point the network path should be non contended other the switch backplane. This is so that bandwidth available to the end user can be assured if required.

Note the reason why a single interconnect will not work is that CIS will not provide quality of service and availability service level targets.

2.5.4 Internet Access

The HNGX network has internet connectivity as a number of services requires this;- for instance Money gram and Telecoms service.

The approach to provide this is based on using the existing C&W service to provide internet access;

1. Separate MPLS VPN with dedicated Ethernet port for handoff on CE. The amount of bandwidth will be determined during the design stage and the service from C&W will need to be flexible to cover future potential changes.
2. Cross over connection from this port to Juniper Netscreen Firewalls. This could be another Firewall make that is used elsewhere within Fujitsu in an internet facing role. The key requirement is that this must be non Cisco. Rationale for cross over connection is to avoid having any HNGX device with a listening process facing the Internet dircetly.
3. The Traffic path will be CE---- x ----Juniper---1---ASA----Internet DMZ, path labelled as 1 can be switched.
4. The Internet traffic will be terminated at layer 4 within a platform located in the Internet DMZ

2.6 Branch Network and intra Branch Networking

2.6.1 WAN Technology

This section documents the approach to providing Network connectivity to the Post Office Branches.

The two main challenges in providing a network to all Post office Branches are:

- Geographic coverage of Post Office Branches resulting in not being able to cover all Branches with the same network type. In particular there are a small number of remote branches with limited communication options.
- Optimising cost and quality of service, for example using ADSL with automated backup to provide a business class service.

The Mobile Branch presents a further challenge in that a single mobile counter operates at multiple serving locations with possibly different network types at each location.

The following table compares the key characteristics of all the candidate technologies for the branch network.

Candidate Technology	Pros	Cons
ADSL (c)	1. Proportion of Branches where ADSL is viable ≈ 98% 2. Strategic technology – BT commitment 3. Always on connection avoids complexity of managing connections and is simple to manage.	1. Not available at every branch. 2. Long repair time from faults (up to 40 hours).
ISDN2e (a)	1. Fast connect time (< 3 seconds) 2. Reuse of existing infrastructure and experience in POA	1. Metered tariff makes 24 * 7 connectivity non viable. This in turn causes complexity in the solution to manage calling 2. Not available at every branch.
PSTN	1. Almost Universal availability (see Appendix A).	1. Metered tariff makes 24 * 7 connectivity non viable. This in turn causes complexity in the solution to manage calling 2. Slow connect time (>20 seconds) makes dial on demand non viable for transactions. This result in complexity where the call is held open during the online day or special action is necessary in the counter application to cause traffic when a customer session is about to start 3. Bandwidth low relative to other technologies [22 – 28k bps].
VSAT	1. Most Universal of all technologies. Sufficient to have clear view of Southern sky for service 2. Always on connection avoids complexity of managing connections and is simple to manage. 3.Broadband bandwidths available 4.Mobile solution available	1. Availability – ≈ 24 hours downtime per annum due to rain fade 2. Satellite incident can cause failure to all branches 3, Planning permission 4. Installation time
GPRS / 3G (b)	1. Cost effective packet with cost based on usage	1. Coverage is difficult to predict – both overall and for a particular branch location 2. SIM cards needs managing 3. Security is more complicated since no concept of location. Conversely uses of location services from Service provider would add significant complexity to solution. 4. GPRS bandwidth variable 5. High latency

Candidate Technology	Pros	Cons
CSD / HSCSD	Provide mobility	1. Expensive per minute rate, HSCSD more expensive that CSD 2. Slow connect time (>20 seconds) makes dial on demand non viable for transactions. This result in complexity where the call is held open during the online day or special action is necessary in the counter application to cause traffic when a customer session is about to start 3. Bandwidth low relative to other technologies [CSD 9,600 – 14,000 bps] HSCSD [.19,200 – 28,800 bps]

Table 3 – Candidate Branch Network Technologies

Notes
(a) Part of BT 21CN hence has future– confirming this with BT. Once this confirmation is obtained the document will be updated. However risk that will not be available by 2015.

(b) [DN: Have asked Orange for details of 2g License renewal beyond end 2012]

(c) Covers ADSL, ADLS2 and ADSL2+ technology

The approach taken is to define a preferred network access model for a branch as;

- HNG-X Branch Router used in all Branches including Broadband) VSAT to provide Network connectivity. The rationale for deploying this in VSAT sites is to have a common solution across all branches. Note the HNGX Branch Router will replace the Cisco 1800 provided by Core services and duplicate its functionality with respect of IPSEC tunnels.
- First choice Primary Network type is ADSL (predicted to be available at 98% of fixed branches
- Second choice Primary Network type is ISDN (Dial on demand)
- Third choice Primary network type is VSAT Broadband
- Back up Technologies enabled are GPRS / EDGE / 3G, ISDN [always on for defined period], PSTN [always on for a defined period]. Dial out "prod" is used to cause branches on dial technology [ISDN, PSTN] to connect in.
- The first choice backup technology will be Wireless WAN (GPRS, etc.). In the case where Wireless WAN is not viable then Post Office may nominate branches where ISDN or PSTN is to be used for backup.
- All Branches will have a Branch Router with Wireless WAN enabled. The viability of this as backup network will be determined as part of the commissioning process and includes a 24 network test (refer to section 5.1.2.1 for further details).

The main Characteristics of the backup network in this preferred network access model are;

- Once the Primary Network failure mode is detected by the Branch Router, failover to the backup network should occur within a few seconds (target is 5 – refer to section 7.2). In the case of PSTN as a backup the dial time needs to be added.

- Refer to section 7.2 for failover targets in other scenarios.
- The case where PPP is responding to LCP probes but data flow is not possible is termed a PPP black hole. The likely cause of this is either a software bug or double failure. This will be detected by the Branch Router observing unidirectional traffic flow for 3 minutes. Should this happen the Router will switch to the backup network and reset the ADSL PPP interface. The switch back to using ADSL will be delayed to prevent oscillation. The delay period will be specified in the Branch Access HLD.
- Once the Primary service is restored, traffic will stop using the backup network. The switch back to the primary network will be totally transparent to the application and TCP/IP sessions will survive. Since the tariff is based on bytes transferred this will minimise charges for the backup network.
- The Backup network does not have capacity to provide backup to the entire Branch estate. Rather it provides sufficient capacity for all likely failure scenarios in the primary network (10% of branches).
- The Backup Network is tested regularly to check whether it is available and any changes in availability are reported via events to the Enterprise Management System. The schedule for testing the backup network will be specified in the Branch Access HLD. The interval is likely to be every two weeks in the case where tests are successful.

In order to provide a Network service to those Branches where the preferred model is not applicable the following approach is used;

- Where GPRS backup is not available in an ADSL branch then an option will be possible to use ISDN or PSTN. In the case of ISDN and PSTN the backup network will remain nailed up for a fixed period or until service is restored to the Primary Network. Since ADSL is carried over the same service as PSTN and ISDN, this form of backup is subject to common failures. However this is mitigated by experience to date which is that most ADSL failures are service related [DN: Need to provide stats for this]

The following table summarises a prediction of the number of branches by Network technology.
(DN: prediction being reworked based on study of coverage from Orange / Vodafone. This will be supplied in next version)

Mobile Branches

A Mobile branch is a Branch, which can consist of multiple counters that operates at more than one Location. The branch Router will be configured to connect to any of {ADSL, ISDN, PSTN, GPRS / EDGE / 3G} and will travel with the Counters from the branch. There are approximately 284 luggable type mobile branches and a maximum of 40 PHU1 type.
In addition there are approximately 50 Mobile CTO branches used to provide training courses.

2.6.2 WAN Network Services

These are specified in
Table 2 – Branch Network Services

2.6.3 Within Branch Network

2.6.3.1 Logical View

The HNG-X application to use the Counter Interface IP address as a consistent identity and therefore the Branch subnet has a fixed and unique address. In addition fixed WAN IP addresses are used fro those WAN interfaces that are reachable from the data centre. The exceptions being the Wireless WAN PPP address and the out of band PSTN address.

The following diagram illustrates a Logical view of the within Branch Network and shows the various IP addresses in use. The chosen IP addressing scheme is based on making the LAN subnet visible to the data centre. This allows for outgoing connections to the branch and for additional audit of which physical counter transactions took place. The counter IP address (specifically the LAN Interface address of the Counter) will be available to the Data centre Branch Access servers on every TCP connection.

Note that Lr and L2 represent individual IP addresses on subnet L.

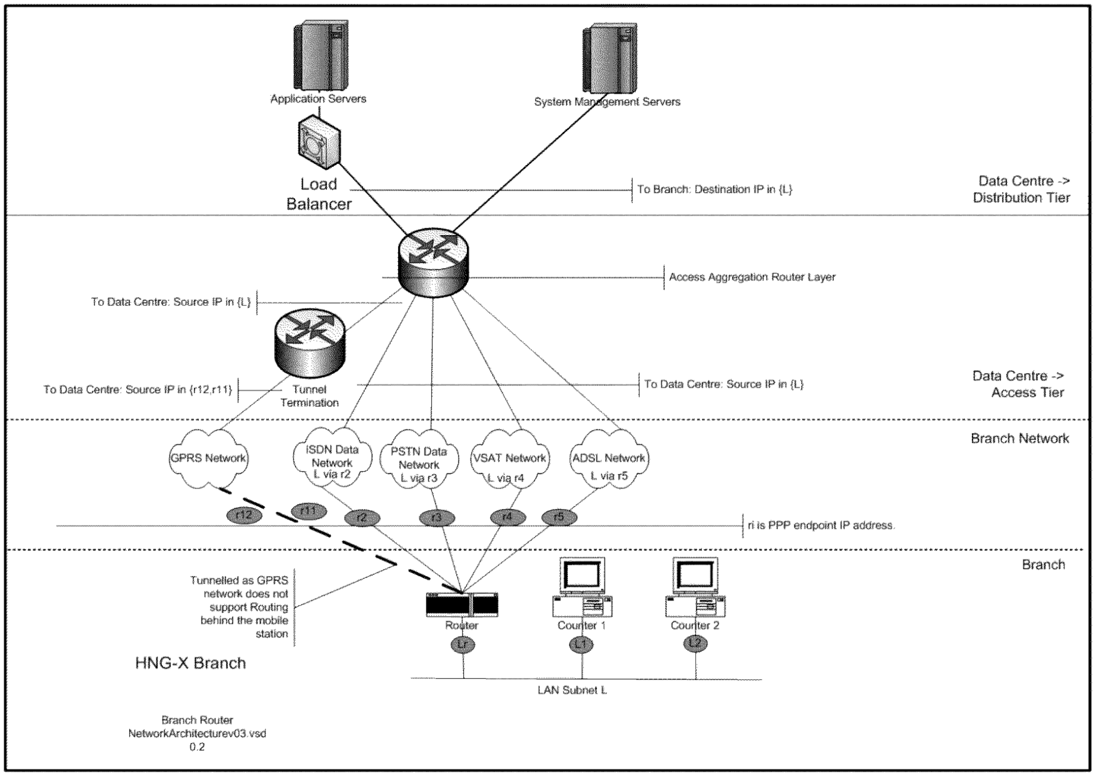
Application Servers
System Management Servers
Load Balancer
To Branch: Destination IP in {L}
Data Centre -> Distribution Tier
Access Aggregation Router Layer
To Data Centre: Source IP in {L}
To Data Centre: Source IP in {r12,r11}
Tunnel Termination
To Data Centre: Source IP in {L}
Data Centre -> Access Tier
Branch Network
GPRS Network
ISDN Data Network L via r2
PSTN Data Network L via r3
VSAT Network L via r4
ADSL Network L via r5
r12
r11
r2
r3
r4
r5
ri is PPP endpoint IP address.
Branch
Tunnelled as GPRS network does not support Routing behind the mobile station
Router
Counter 1
Counter 2
Lr
L1
L2
HNG-X Branch
LAN Subnet L
Branch Router
NetworkArchitecturev03.vsd
0.2

Figure 10 – A logical view within the Branch network

2.6.3.2 Physical View

This section summarises the Current Horizon scheme with a Gateway PC and how this is changed by the Branch Router.

Horizon scheme

Single Counter – WAN connection, No LAN

Two Counter – WAN connection in Gateway Counter, Cross over cable between Counters

Multi Counter – WAN connection in Gateway Counter, 8 port hub with additional 8 port hubs connected as required to create LAN between all Counters.

Branch Router installation

The rollout model for the Branch router is out of scope for this document. The installation of the branch Router is based on the following assertions that the Router needs to be close enough to the Gateway PC so that existing WAN (ADSL and ISDN) cables and the LAN cable can be plugged into the Router.

The proposed approach when installing the Branch Router is to connect the Gateway PC directly to the Router with a new cable and whatever network cables were coming into the Horizon Gateway PC directly into the Router. The Router will support both straight through and cross over cables. Note in a 2 Counter branch the cable between the Slave and Gateway is a cross over cable.

The Router ports will be set to 10 M bits /sec to reduce the risk of adverse effects on the Counters caused by running at 100 M bit /sec.

(An optimisation current being considered as part of the Branch Router rollout in Horizon design, is to take the opportunity to reduce the hub population from about 3,500 to about 1,600 by rewiring 3 Counter sites when the branch router is rolled out.)

As part of installing the Router, Antenna positioning process would have been performed. Outcomes are in terms of signal strength {Red = No Signal, Orange = Marginal Signal, Red = Good signal} strength. It is expected that these traffic lights are correlated with WWAN viability (refer to section 5.1.2.1)but do not assure it.

A physical view of the Branch network for 1, 2 and 3 Counter HNG-X branches is shown in the following diagram (on right).

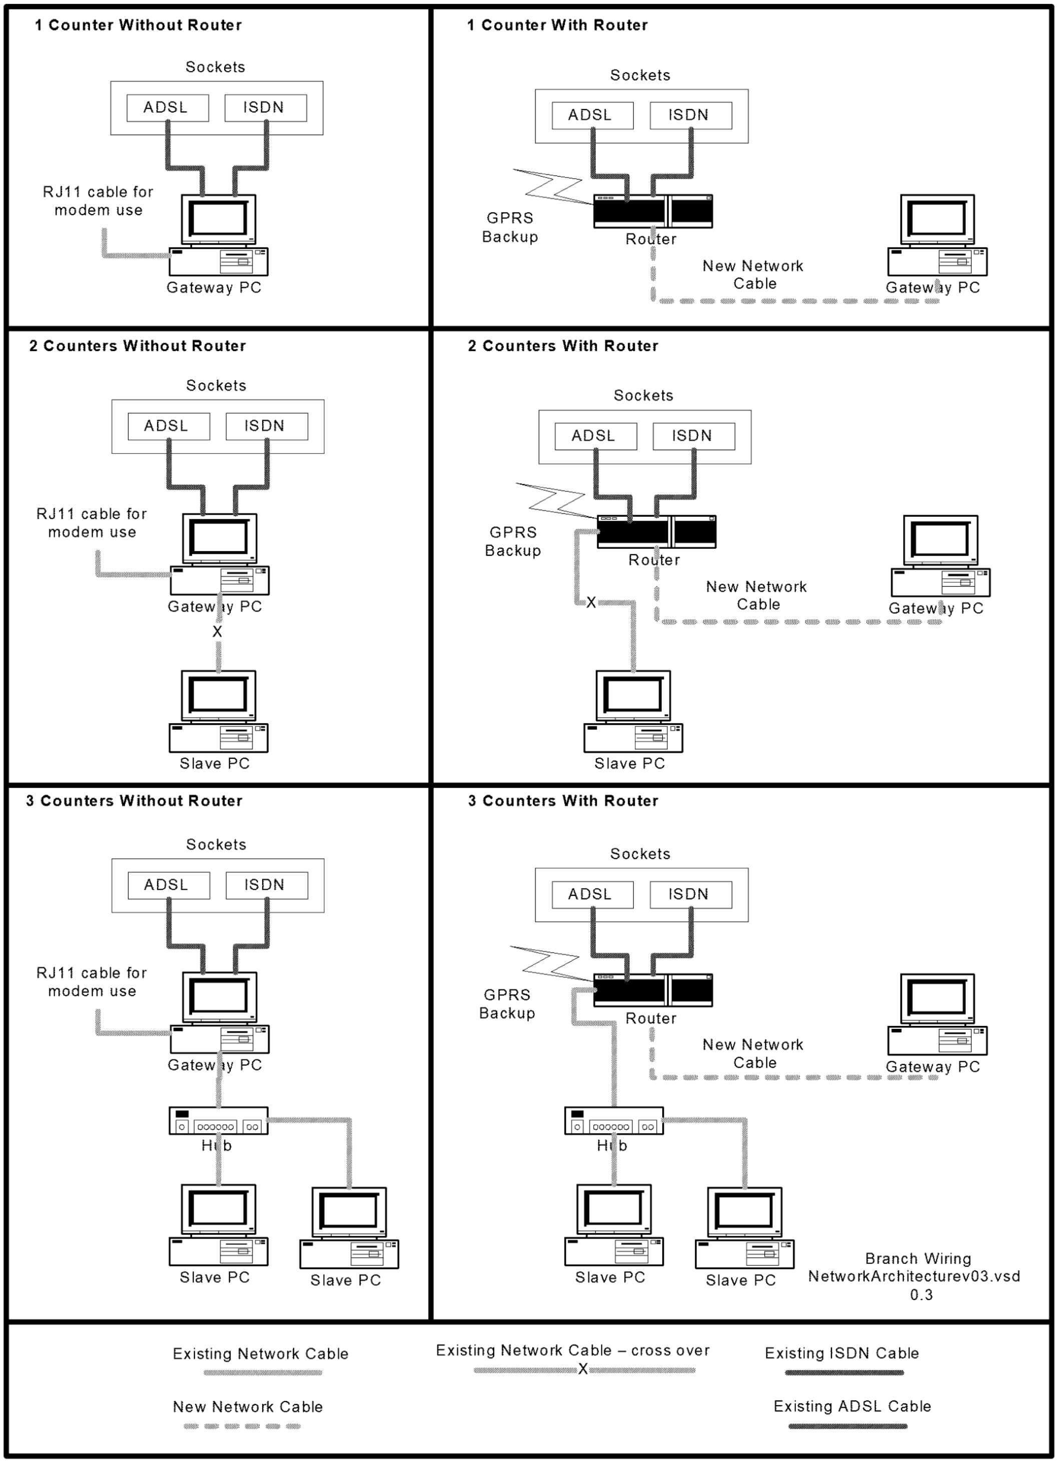

Figure 11 – A physical view of the Branch network

2.6.3.3 IP Addressing

Branch Partitioning

An enumeration of all branches is maintained with each branch having a unique assignment (P, O) where P represents the partition and O the ordinal number within partition. In general 120 partitions are used for branches to provide a grouping of branches. To represent 120 distinct partitions, 7 bits are required (2^7 = 128). If the branches are evenly divided across partitions then there are approximately, 20,000 / 120 ≈ 166 branches per partition. To allow for skew in the allocation, 9 bits are provided this gives 2^9 = 512 unique values of which 1 through 511 are usable.

(Note the figure of 20,000 branches as there are some 5,600 branches closed but still in systems).

The divisors of 120 are {1, 2, 3, 4, 5, 6, 8, 10, 12, 15, 20, 24, 30, 40, 60, 120} and each of these divisors represent the number of parts the branch estate can be divided into.

This partition scheme will be used for assigning IP addresses and the rationale for its existence is that it provides for scaling in the network should this be required. For example if a Resilient pair of Network components in the Access Tier are insufficient to handle the workload then the workload can be split across to such pairs by partition.

IP Addressing

(DN Need to ad scheme for Branch LAN if this is being changed from Horizon)

A network address scheme is introduced for the PPP WAN Interfaces used in the Branch Router. The scheme will support up to 15 I/F assignments. Note the existing Horizon Branch LAN IP address scheme is maintained whilst the Branch is running Utimaco VPN. A new LAN IP scheme is proposed for when the branch migrates to running the HNGX Java based application.

Properties of Address scheme

- RFC 1918 compliant
- Block 10.0 / 9 used for WAN PPP interfaces
- For ADSL, ISDN, PSTN, Virtual PPP – IP address assignment is dynamic via PPP, constant (same next time) and unique across all Branch Routers
- For Wireless WAN – IP address is dynamic via PPP and takes one of two values. This pair is unique across all Branch Routers.
- Address block for given Interface type can be summarised. Therefore efficient summarisation by service provider is possible.
- Maintains branch hierarchic scheme (partition, ordinal)

For Branch Router WAN interfaces the IP addressing scheme is as follows:

Octet 3								2								1								0							
7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0	7	6	5	4	3	2	1	0
Block Identifier								W	R	R B	T	Endpoint				**Partition Identifier**								WAN Ordinal							
0	*0*	*0*	*0*	*1*	*0*	*1*	*0*	*0*	*0*	*0*	*0*	*r*	*r*	*r*	*r*	*p*	*p*	*p*	*p*	*p*	*p*	*p*	*o*	*o*	*o*	*o*	*o*	*o*	*o*	*o*	*o*

Component	Values
Block Identifier	Set to 10
W	Set to 0 indicating this is WAN scheme for PPP
R	Reserved for future use set to 0
TB	0 = Live Address, 1 = Test Address0 = Live operational address.
rrrr	Endpoint (range 1-15) 1 = ADSL 2 = ISDN 3 = PSTN Modem for branch data 6 = Virtual Interface for Wireless WAN L2TP 7 = VSAT BB 8 = Virtual Interface for VSAT BB L2TP 9 = PSTN Modem for Out of Band support Note that the above Endpoints are strict one to one with Router PPP interfaces. The PPP addresses for the outer PPP interface for Wireless WAN are allocated from a pool and hence not art of this scheme. Similarly the Boot loader uses an address pool and hence its not part of this scheme.
ppppppp	partition identifier (1-123), allocated by ACDB– see TD/DES/157.
ooooooooo	WAN Host Identifier, allocated by ACDB (values 1-509254, 256 - 508 are available to live outlets, 509 is reserved for Estate Management Test Facility. 510 and 511 are restricted).

In order to achieve orthogonal address space between test environments, the collection of Partitions used for a given Test environment will be unique. The special case of VSAT sites (one partition) will be dealt with by use of unique ordinals.

A separate pool of WAN Host Identifiers from the existing ISDN and VSAT[1] pools will not need to be used for the assignments.

Some example IP addresses using this scheme are:

Live ADSL

					Address				Octet			
W	T	PPP IF	P	O					3	2	1	0
0	0	1	1	1	10	1	2	1	00001010	00000001	00000010	00000001
0	0	1	1	509	10	1	3	253	00001010	00000001	00000011	11111101
0	0	1	123	1	10	1	246	1	00001010	00000001	11110110	00000001
0	0	1	123	509	10	1	247	253	00001010	00000001	11110111	11111101

Live ISDN

0	0	2	1	1	10	2	2	1	00001010	00000010	00000010	00000001
0	0	2	1	509	10	2	3	253	00001010	00000010	00000011	11111101
0	0	2	123	1	10	2	246	1	00001010	00000010	11110110	00000001
0	0	2	123	509	10	2	247	253	00001010	00000010	11110111	11111101

Live Wireless WAN

0	0	4	1	1	10	4	2	1	00001010	00000100	00000010	00000001
0	0	4	1	509	10	4	3	253	00001010	00000100	00000011	11111101
0	0	4	123	1	10	4	246	1	00001010	00000100	11110110	00000001
0	0	4	123	509	10	4	247	253	00001010	00000100	11110111	11111101
0	0	5	1	1	10	5	2	1	00001010	00000101	00000010	00000001
0	0	5	1	509	10	5	3	253	00001010	00000101	00000011	11111101
0	0	5	123	1	10	5	246	1	00001010	00000101	11110110	00000001
0	0	5	123	509	10	5	247	253	00001010	00000101	11110111	11111101
0	0	6	1	1	10	6	2	1	00001010	00000110	00000010	00000001
0	0	6	1	509	10	6	3	253	00001010	00000110	00000011	11111101
0	0	6	123	1	10	6	246	1	00001010	00000110	11110110	00000001
0	0	6	123	509	10	6	247	253	00001010	00000110	11110111	11111101

[1] Note that VSAT sites use just one partition (123). This means that the new address scheme will result in a maximum of 508 branches on VSAT. (Currently the VSAT estate is to drop below 100 so this appears acceptable.

Test ADSL

0	1	1	1	1
0	1	1	1	509
0	1	1	123	1
0	1	1	123	509

10	17	2	1
10	17	3	253
10	17	246	1
10	17	247	253

00001010	00010001	00000010	00000001
00001010	00010001	00000011	11111101
00001010	00010001	11110110	00000001
00001010	00010001	11110111	11111101

Test ISDN

0	1	2	1	1
0	1	2	1	509
0	1	2	123	1
0	1	2	123	509

10	18	2	1
10	18	3	253
10	18	246	1
10	18	247	253

00001010	00010010	00000010	00000001
00001010	00010010	00000011	11111101
00001010	00010010	11110110	00000001
00001010	00010010	11110111	11111101

Test Wireless WAN

0	1	4	1	1
0	1	4	1	509
0	1	4	123	1
0	1	4	123	509
0	1	5	1	1
0	1	5	1	509
0	1	5	123	1
0	1	5	123	509
0	1	6	1	1
0	1	6	1	509
0	1	6	123	1
0	1	6	123	509

10	20	2	1
10	20	3	253
10	20	246	1
10	20	247	253
10	21	2	1
10	21	3	253
10	21	246	1
10	21	247	253
10	22	2	1
10	22	3	253
10	22	246	1
10	22	247	253

00001010	00010100	00000010	00000001
00001010	00010100	00000011	11111101
00001010	00010100	11110110	00000001
00001010	00010100	11110111	11111101
00001010	00010101	00000010	00000001
00001010	00010101	00000011	11111101
00001010	00010101	11110110	00000001
00001010	00010101	11110111	11111101
00001010	00010110	00000010	00000001
00001010	00010110	00000011	11111101
00001010	00010110	11110110	00000001
00001010	00010110	11110111	11111101

Table 4 – Sample IP addresses

The following table summarises the properties of the IP addresses used within the Branch.

Address Space	Assignment	Scope	Constant Next time is the same	Outgoing Connections to endpoint?
Branch LAN	Constant – preconfigured	Unique within PAS	Yes	Yes

Address Space	Assignment	Scope	Constant Next time is the same	Outgoing Connections to endpoint?
Branch WAN (ADSL, ISDN, PSTN)	Dynamic by PPP	Unique within PAS	Yes Chap user name will identify the IP address.	Limited to support
Branch WAN GPRS	Dynamic by PPP	Unique within PAS	No allocated from pool	No reachable from Data centre
Branch Tunnelled PPP	Dynamic by PPP	Unique within PAS	Yes Chap user name will identify the IP address.	Limited to support only
Branch (VSAT)	Dynamic IPSEC client	Unique within PAS	Yes	Limited to support only

Table 5 – Branch IP Address Properties

WAN IP addressing

The Router has multiple interfaces; {ADSL, ISDN, PSTN, GPRS, and VSAT}.

For {ADSL, ISDN, PSTN and VSAT} the interface IP addresses are constant and unique across all branches. This means that the IP address for the PPP interface in a given branch will always be the same.

For GPRS, there are two possible interface IP addresses depending on which GGSN is chosen by the Orange network. The choice of GGSN is made by Orange. This pair is constant and unique across all branches. Note this is different from the ISDN case where the same IP address is assigned irrespective of the destination phone number.

GPRS

The GPRS network does not support Routing behind the Mobile Station on an APN. Put simply this means that IP packets from the branch must have a source IP address which is the Branch Router PPP interface address. Similarly all packets destined to the Branch must have a destination IP address which is the Branch Router PPP interface address.

To meet the properties in 2.7.2, L2TP is used.

Without L2TP it would be necessary to use NAT to translate all LAN addresses to a single IP address – the GPRS PPP interface.

The consequences of using L2TP are;

- Additional complexity in configuration, testing and operational support of Branch Router
- Additional bandwidth usage due to L2TP header and L2TP Tunnel setup and L2TP Tunnel keep alive. (DN: need to model this).

Notes;

1. The L2TP head end will be configured to adjust the TCP MSS in order to avoid IP fragmentation.

2. In ADLS branches, WWAN (and hence the L2TP Tunnel) is only active when the Branch Router determines that ADSL has failed.

2.6.3.4 Host Components

Within the branch there is a single Counter host component relevant to the Branch Router. This component is named CNIM2 (for Branch Router) after the existing Horizon component (CNIM) which is being replaced. Note that CNIM2 carries forward the monitoring functions of CNIM but not any of the explicit control functions and so is a very different and simpler component.

CNIM2maintains information about the availability of the Branch network over time. It infers the state of the within branch LAN and WAN by listening to Syslog broadcasts and RIP broadcasts from the branch Router. The Syslog broadcasts carry information about interface transitions and firewall events. The Rip broadcasts provide a heartbeat function as they are regular (every 30 seconds) and also indicate which network interfaces are up and which interface is used for outbound traffic. CNIM2 provides an application interface to indicate the state of the network. For example this information is used is to populate the HNG_X online indicator.

Within HNG-X there is not a distinct counter role (like the Gateway Horizon) and therefore CNIM2 is installed and active on every counter. One instance of CNIM2 (Primary) will create SYSMAN events. A voting algorithm will be used to determine the Primary instance and this algorithm will run when a CNIM2 instance detects a change in the number of CNIM2 peers as well as periodically. The algorithm will be deterministic in that the winner remains fairly constant over time. For example the ordering would be based on ascending Ethernet MAC address. It would not cause any significant problems if more than one instance were active. In this case duplicate events may be received at the data centre.
To provide support staff with a mechanism for detecting which instance of CNIM2 is active, the CNIM2 process on every counter will provide a mechanism for identifying the Primary instance.

To avoid excess WAN traffic only the Primary instance (there may be more than one in the case of a LAN partition) will actively issue ICMP Pings. The result of these Pings will not update the Finite state machine for reporting the branch status; rather they will generate traffic that the Branch router uses for passive testing of network connections. The frequency of these Pings will vary depending on the current network type that is active.
Note a LAN partition exists (by definition) when there at least two sets of counters (S1 and S2) such that any counter in S1 cannot communicate in a bidirectional way with any counter in S2.

All instances of CNIM will create log files.

CNIM also causes the backup network types to be tested by use of QOS marking ICMP ECHO requests. Basically the branch Router will route on QOS to a backup interface and cause this to be made active if necessary. The return data gram to the ICMP ping will be routed via the preferred interface however as part of bringing up the backup interface the Router will have issued an ICMP ping sourced from the backup interface in which case the reply gets routed over the backup network. Therefore a successful transition to “up” for the backup network, which is reported to CNIM via SYSLOG messages indicates two way flow over the backup network.

Function ID	Summary

POL00140236

Function ID	Summary
Syslog listener	Listens on Branch Router Syslog messages and uses these to ▪ Provide persistent event storage on Counter. This is used for on demand fetching of log files for diagnostics or forensics. ▪ Maintain Finite State Machine for Branch Communications and communicate significant transitions via Windows events into Enterprise Management system. ▪ Maintain Registry information for use by other applications o IP address of Router o Communications status ▪ If Syslog message indicates incoming isdn call rejected then cause ping sequence which in turn causes dial on demand
RIP listener	Listens on Branch Router RIP and uses this to maintain Finite State Machine for Branch Communications Based on RIP contents determines whether route available to data centre and maintains relevant status. Note CNIM2 will determine from RIP payload which network type is being used to route traffic to data centre and the status of the associated interfaces. Alerts User(s) if Router reconfiguration or replacement results in a mismatch between Router LAN IP address and current counter IP address.
Regular Ping	• To generate regular traffic for the Router to use as a source of passive monitoring • To quantify network outage whilst Router is performing automated recovery (This provides diagnostic information). • To quantify whether PPP black holes are detected and recovered by the Branch router (This provides diagnostic information). • Backup Network testing
Meta Messages	CNIM monitors the status of the Branch Router based on the following sources; • SYSLOG events sent from Branch Router • Periodic RIP broadcasts sent from the Branch Router • Reachability of Slave counters as reported by Riposte CNIM provides meta states based on underlying information. These meta states are (not exhaustively); • WAN broken (which one) • Router Broken – needs replacement • LAN broken – needs cabling changes • Not determined • Router not detected

2.6.3.5 Management traffic

The management functions for then branch router are covered in section 5.1.2.

There are two classes of traffic flows;

- In clear ICMP echo request / response from the network Management system
- All other Management traffic is encrypted using IPSEC

No local incoming traffic destined for the router is permitted on either Ethernet or serial interfaces.

2.6.4 Branch Router

2.6.4.1 Introduction

Section 2.6.3 documented the Network aspects for Branch networking in terms of supporting Traffic flows over various WAN technologies. This section documents the lifecycle aspects of the Branch Router and clarifies the configuration types of the Router from a Network perspective

The key Architectural decisions are that Branch Estate Management looks after provisioning of the Branch Router and Systems Management looks after the Inventory, Remote Operation and Configuration management aspects. An important distinction between the Branch Router and other Systems Management endpoints is that the Branch Router clearly does not host and execute any Systems Management code. Rather all operations are performed by scripts which used in industry standard protocols to read and change the state of the Branch Router.

The sections titled Gold Build and Provisioning below cover how Branch Estate Management results in a Branch router receiving its relevant firmware and Configuration.

Once Branch Estate Management has provisioned a Router it will cause SYSMAN2 to register the Branch router as a managed object and then to create inventory records against the BR object. Note that the FAD code is an attribute of the Branch Router managed object.

Registering a Branch Router at a Branch causes implicit deregistration of any previous Branch Routers against the branch. This is because of the data model where a Branch (FAD) has only one current Branch Router.

The sections titled Registration and Deregistration deal with these aspects of the Branch Router.

Configuration updates, Remote operations and inventory maintenance are summarised in sections titled Configuration Updates and Remote Operations.

For completeness of description sections are included on Troubleshooting, Monitoring and Access & Security. These mainly provide references to other places in this document where these aspects are covered.

2.6.4.2 Gold Build

All Routers supplied from Sarian have a Gold Build applied. The Gold build is developed and supplied by RMGA and is a generic. Multiple versions will be allowed in the solution to;

- Provide flexibility for fixes and changes
- Allow for Firmware updates not under our control (such as changes to service provider equipment / Wireless WAN environment)

Is is necessary that a Router is in a Gold build state prior to provisioning on site.

The Management of the spares cycle for the Branch Router is specified in (need reference to Branch Estate Management TA).

2.6.4.3 Provisioning

2.6.4.3.1 Introduction

Routers deployed to Branches for installation are in a generic state with a "Gold Build "configuration. The installation process requires that the Router is connected to one or more WAN technologies that identify the location of the router via a CLI mechanism; {ADSL with SID, ISDN, PSTN}.

The Router will establish one or more PPP sessions which provide connectivity to the (Boot loader service - see section 2.3.5.12) and perform the following steps in sequence;

- Firmware catch-up if necessary
- Personality delivered to Branch Router (Branch Specific)
- Registration and Deregistration
- Cause Router to Reboot and come up on Production network

The remainder of this section covers the above steps in further detail.

2.6.4.3.2 Session establishment, authentication and Branch identification

The following table provides a chronology of events starting from session establishment up to when IP communication is achieved with the Branch Router.

Step	Result	Explanation
#1	The Branch Router initiates one or more PPP sessions, one per WAN technology that is available to it at the time of installation. The WAN technologies are {ADSL, ISDN, PSTN}.	The PPP name, configured in the Gold Build, identifies the WAN technology and includes the serial number of the Router. The ability to make the PPP name unique in a generic configuration is achieved using the Sarian capability to textually substitute the Router Serial number into PPP name. This Router Serial number is asserted by Sarian to be unique across all possible Routers. Note that there is one Boot loader service per data centre and in general all PPP session land on the same Boot loader service as this is deployed in an Active / DR model.
#2	The Boot loader will authenticate the incoming PPP session based on the CLI / SID value being known as one that is associated with some Branch. It will allocate the PPP IP address dynamically from a pool.	The CLI / SID information is available in advance from Service providers. Note that the Boot loader is running a Radius server (Radius) Refer to section 2.7.1 for details of the Branch Data model.

Step	Result	Explanation
#3	The Boot loader initiates communications with the Branch Router once it received a Radius Accounting start from the service provider. This confirms that the Service provider has set up the session. In the event of multiple WAN types being available the ordering is [ADSL, ISDN, and PSTN].	At this stage the Boot loader has learnt the following; Router serial number #S is at Branch #B Router #S is reachable via a least one WAN type For each WAN type, the IP address used to reach #S over that WAN type.

To deal with the case where ADSL SID into Authentication is not available, the Boot loader will have a configurable option to authenticate an incoming ADSL session if an associated PSTN session has recently authenticated over PSTN. This is to allow any firmware updates to be applied over ADSL. Note that there is policing applied by BT of which ADSL service endpoints can initiate sessions to the RMGA domain (refer to section 6.4 for further details). The Firmware update will be performed over ADSL and similarly the configuration update which includes sending Shared secrets and Passwords is performed over ADSL.

For the reason why it is important to use ADSL for Firmware updates refer to section 2.6.4.3.3.

Note that the advantage of using ADSL SID is that it avoids the engineer having to carry a PSTN modem and connect this to the Postmasters PSTN line when installing a Branch Router.

Note that it is necessary to provide an exception process to cover the scenario where the CLI / SID presented to the Boot loader is not known.

The process is based on the onsite Engineer calling the helpdesk and providing the Router serial number and the branch identifier. The helpdesk verifies that the installation is legitimate and initiates an automated task to cause the boot loader to associate the Router serial number with the branch for a limited period of time (about 30 minutes). This will enable the Router to be provisioned as described previously.

2.6.4.3.3 Firmware catch up

The scenarios under which Firmware catch up is necessary at installation time are;

1. Where the Branch specific configuration will not run on the version of Firmware installed on the Gold build. For example it uses features which are not supported in earlier firmware
2. Gold build firmware has known problems which must be fixed before business traffic can be reliably supported

The impact of taking the simple approach of performing a Firmware update where the Gold build is behind is predicted at;

- 13,500 Branches on ADSL (94%), update will take about 2 minutes
- If 1% of ADSL branches experience service failures during BR rollout this will affect about 135 branches and increase update time to about 35 minutes (over PSTN)
- 800 branches (or less) are on ISDN (5.6%), update will take about 15 minutes and hence increase engineer on-site time
- 62 VSAT (0.4%) branches are expected to take about the same time as ISDN

All timings based on a 4.7 M byte firmware update.

In particular where primary network service is not available in the branch the provisioning alone (firmware update) of the router **may take up to 35 minutes and violate SLA(s).** Need to validate whether this is acceptable with CS.

Assuming this is not acceptable then the following approach is proposed. This approach is based on the premise that scenarios (1&2) represents a small proportion of the instances where the current build is advanced and hence it is acceptable to defer updating the firmware from install time to a maintenance slot in the evening.

Basically the way the Boot loader application checks the Firmware version installed on the Branch Router is modified as follows;

1. Where the Router is being provisioned over its ADSL it will check the Router firmware against the scalar Target_Firmware. If a version mismatch occurs then the target HNG-X BR firmware version will be downloaded.
2. Where the Router is being provisioned over any Network type other than ADSL it will check the Router firmware against the List Allowed_Firmware. If the Router firmware version does not exist in this list then the target HNG-X BR firmware version will be downloaded.

In case 2 the Router Firmware updated will be applied as part of normal configuration maintenance (see section 2.6.4.6).

(DN: We need to decide on the approach to be taken)

The current solution for the Branch Solution has two hardware variants of the Branch Router (ISDN and non ISDN module). It is not necessary to differentiate the Firmware to be applied for the Hardware variants since Sarian provide a single image file for both variants.

Note that the strategy that Sarian are pursuing for the DR6400 series Routers is that a single image file can be applied to all hardware variants.

2.6.4.3.4 Personality delivered to BR

The Boot loader application will deliver the Personality to the Branch Router using an existing communications session. The Personality is the site specific configuration file derived from a configuration template with substitution of site specific variables. The following table documents the variables and provides an explanation of their usage.

The fundamental mechanisms for controlling the Branch Router behaviour with respect to its use of available network technologies and switching between then are the Branch specific variable Network Service Type. The Configuration Template for the Branch Router is selected based on the Network Service Type (NST) for the branch. For example in the case of Branches whose primary network is ISDN dial on demand, NST is used to differentiate between;

- Branches with normal idle timers optimised for cost reduction
- Branches with extended idle timers optimised for reliability of call connect

(DN: Need to check whether this is NST or Comms type)

A secondary selection of template is based on the hardware variant of the Branch Router being installed.

The configuration template to be used is selected based on Network Service Type (NST) and this provides mechanisms for having different branch configurations. In addition the configuration file to be deployed will in general be a function of the Router hardware type and the solution needs to allow for this. The following example illustrates this point;

Suppose that a single configuration file is used for all branches where ADSL is the Primary communications method. Then the branch routers deployed in these branches can be categorised as follows;

Enumeration	Branch Communications solution as specified by Post Office for the Branch	Router Hardware variant	Configuration requirements
#1	ADSL with no ISDN backup	ADSL with no ISDN module	The behaviour resulting from the configuration needs to avoid causing the Router to source alert / event messages indicating ISDN failures. This is so that the recipients of these diagnostic messages can avoid interpreting them in the context of the branch specific communications solution
#2	ADSL with no ISDN backup	ADSL with ISDN module	Similar to previous although different error states as no ISDN line plugged into module.
#3	ADSL with ISDN backup	ADSL with ISDN module	All alert / event messages relating to ISDN usage and failures are relevant

The decision as to whether a single configuration file is suitable for the above 3 cases will be determined during the design phase. However the solution needs to allow for Hardware specific dependencies. Sarian provides a command line to detect the hardware present in a Router – “ati5”.

The following table documents (but not exhaustively) the Branch specific variable data in the configuration.

Variable data	Comment
Branch specific LAN IP addressing	Each branch has a unique and constant subnet.
Hostname	This appears in Syslog messages and is set to include the serial number of the router. Note that the complete Syslog information includes the IP header which contains the source IP of the Syslog. That way it can be determined what branch the router was configured for.
Unit ID	This is a string of up to 20 characters that is also displayed as a command prompt when logging on remotely via telnet / SSH.
Comments embedded in configuration file	Configured for branch #b on Router #s on dd/mm/yyyy at hh:mm:ss

Variable data	Comment
Primary Radius Pre-shared Key Secondary Radius Pre-shared Key	Pre-shared Key for securing Radius authentication channel for user and/or service account access to BR device for management/maintenance/troubleshooting. In general these variables will be constant across the estate and only change as part of a planned update to the key
PPP Session Information (PPP Username/Password)	PPP CHAP Username and PPP CHAP pre-shared-key. There is one PPP CHAP Username per Router interface and this is unique. There is one PPP CHAP pre-shared-key per Router which is applied to all interfaces. For details of the PPP Username refer to section 2.7.1 The PPP CHAP pre shared key is allocated to the Branch Router at installation time and never reused. Therefore apart from chance collisions this is unique. To clarify, the next Router installed at the Branch will be allocated new pre-shared-key.
IPSEC identification and pre shared key	Each Branch Router has a unique (apart from change collision) IPSEC pre shared key
SSH Host key	This is generated on the Branch Router as part of the Provisioning process. It is to support SSH access.
Wireless WAN Viable Flag	For purposes of SLA reporting it is necessary to know for each Branch whether it has alternative backup facilities or not. This is because we can exclude from the availability calculation ANY Branch that uses ADSL as its sole method of communication – i.e. there are no alternative back up facilities available. This happens during incident cleansing when the availability SLA is calculated In the case of ADSL branches backed up with ISDN, the rationale for paying for ISDN is that this branch has backup. So from an SLA perspective such branches may be treated as being in the “having backup” category. However this is not the case for Wireless WAN. In general every Branch has Wireless WAN capability and the challenge is to categorise this branch, from an SLA perspective as being viable or not. The proposed process for doing this is documented in section [5.1.2]. This variable needs to be made available to the SLT calculation model. This flag has 3 states, Yes, No, Not determined. It is set to Not determined initially to avoid making any assumptions about Wireless WAN availability.

Notes;

The Number of hosts in Branch LAN is not required as the Branch Router does not require this information (refer to section 2.6.4.10).

It is envisaged that the following data will be in the template (so not variable data)

- Telephone access numbers
- Addresses of Radius Authentication Servers
- Addresses of L2TP Servers.

2.6.4.4 Registration

Registration has the following outcomes for the Branch Router being provisioned / replaced;

1. The Branch Router becoming known to a Management framework that provides functions such as configuration management, inventory and remote operations. The primary key for the Branch Router object is derived for its serial number. The likely choice for this framework is SYSMAN3 using tooling termed Tivoli workload manager which provides an agent less model – that is where the managed end point does not run any Tivoli code. For details of this Framework refer to ???? .
2. The Branch Router PPP details (username and pre shared keys) embedded in the configuration are added to the Radius servers in order that it can authenticate to the production network.
3. The IPSEC details embedded in its configuration are added to the VPN Concentrators to enable it to be managed from the data centre.

This registration process is dynamic since the configuration record for the Branch Router includes the serial number and the branch identifier which is not known in advance. In addition the PPP authentication credentials are constructed dynamically; the PPP name contains the Router serial number and the CHAP pre shared secret is new (refer to section 2.7.1 for further details).

The mechanisms used by the Boot loader application to cause registration are described in ????

The Radius server behaviour to support registration is described in section 2.3.5.2.

The VPN Concentrator behaviour to support registration is described in section 2.3.5.10.

2.6.4.5 Deregistration

Deregistration occurs under two events;

- Installation of a Router at a branch causing any previous Router(s) registered for that Branch to be deregistered. Note that fundamentally there is one current Router per branch as described in section 2.7.1.
- A stolen / missing branch router is deregistered using a process provided within the Management framework.

The following occurs as a result of deregistration;

4. The Branch Router object within the Management framework is updated to indicate that the Branch Router is no longer current for the Branch.
5. The Branch Router entries in the Radius servers are removed asynchronously; the key for these entries is the PPP username whose format is covered in section 2.7.1.

6. The Branch Router entries in the VPN concentrator are removed asynchronously; the key for these entries is the IPSEC username whose format is covered in section 2.7.1.
7. Approximately 10 minutes after the deregistration process started (which allows time for the asynchronous update to the Radius servers), the Management framework performs a Remote operation which consists of an attempt to logon to the Branch Router (via SSH) using all possible WAN addresses. If the Router reached is the one being deregistered (based on serial number) then the Router is instructed to reboot. The rationale for this step is to prevent a stolen Router remaining connected to the Production Network based on a session that never expires. Upon reboot the Router should not e able to authenticate as its PPP username will not be in the Radius servers.

Deregistration also prevents the scenario where a Router removed from a Branch by an Engineer, gets taken to a location for confirming it has a genuine fault before being returned to Sarian. If this location has suitable signal availability and the SIM card remains in the unit then the act of powering it on will (assuming its not faulty) cause it to connect on its Wireless WAN interface. Given that the design is likely to result in the Wireless WAN route from the data centre having higher priority for than ADSL this will cause a service denial to the Branch that the Router came from.

2.6.4.6 Configuration Updates

These are performed by the Management framework further details of which are available in External TA ref ????.

The Management Framework will have the capability of updating all Branch Routers in 3 days based on a 6 hour slot per day and sufficient bandwidth being available. The scheduling mechanism will include the function to limit concurrency.

(Dn: Need some figures here and also statement about concurrency controls)

The Triggers for Updates include the following:

- New Configuration Template
- Change in configuration generator may trigger updated configuration files
- Network Service Type (NST) change requires different configuration template
- New Firmware
- New Chap Set and/or Key Material

An example of an NST change is updating the ISDN idle timers in a Dial on Demand ISDN branch

Switching a Branch Primary Communications from ISDN to ADSL may require an NST change. The following scenario describes possible reasons for this;

Phase 1; ADSL enabled and connected; Whilst a Branch is using ISDN dial on demand as a primary communication type its configuration will be such that it attempts to use ADSL if available. This is necessary to enable the migration from ISDN to ADSL. The Branch Router will perform basic tests on the reliability of the ADSL service to make sure it is usable, for example large pings. If these fail it will switch back to using ISDN and retry ADSL after a defined period – say 2 hours.

Phase 2; ADSL observation period;

The Team managing the migration will observe that the Branch is using ADSL (via Enterprise Management events). For a period of time (at least 2 weeks) the reliability of the ADSL service will be monitored both passively (by observing events) and actively by scheduling file transfer (to and from branch) using the Remote Operations facility in the management Framework.

Should ADSL prove to be unreliable at any time in a manner which allows the branch router to switch to it but transactions are unreliable then a configuration update will be made to avoid using ADSL. This will be done as part of "business as usual" problem resolution and may involve sending an engineer to site to disconnect the ADSL service. Further attempts may be made to use ADSL as a result of discussions with BT / Core services.

The observation period has a number of outcomes;

1. ADSL is reliable and the decision is made by the ADSL migration Team to make ADSL the primary communications for the branch. The NST for the Branch will change and the configuration template may be different. In particular if ISDN is selected as a backup technology for the branch then its usage will not be dial on demand but always on during the core trading period (see section 2.6.1).
2. ADSL is not reliable but this is detected by the Branch Router thereby avoiding loss of service to the branch. Attempts will be made to resolve this with the Service Provider. Should these become exhausted then the ADSL service will be cancelled.

2.6.4.7 Remote Operations

These are provided by the Management framework and are implemented by using scripts which invoke protocols such as telnet and ftp to read / update the Branch Router. The results of these operation scan be fed back into the Inventory.

2.6.4.8 Troubleshooting

Troubleshooting will use a combination of the data available via monitoring mechanisms covered in section [5.1.2] and active mechanisms based on approach where SSC (who provide 3rd line support) create a Perl script and prototype this in LST. The script would then be "packaged up" within the Systems Management Framework which would enable the script to be executed for selected branches via 1st / 2nd line support.

2.6.4.9 Monitoring

These functions are covered in section [5.1.2].

2.6.4.10 Access & Security

This section lists the Access & Security mechanisms relevant to the Branch Router. References are provided for further detail.

Function	Summary	Reference
Branch Router security controls	Covers protection of the Branch router for traffic destined to it and controls for traffic flowing through it	Section 6.4.3.4
Management Access	Interactive / scripted Telnet, FTP over IPSEC Interactive HTTP over IPSEC SSH, SCP	Section 5.1.2

Function	Summary	Reference
Access Control & Command Accounting	Radius Server in DC contains user accounts for 3rd line Net Support, SSC and functions. Credentials are passed on from BR to DC Radius for verification, if verified access is required This solution does not plug into Active Directory and Two Factor Authentication Command Accounting via Radius	Section 2.3.5.2
Auditing	Security relevant management events are logged to CNIM and Data Centre	Section 5.1.2

2.7 Naming & Addressing

2.7.1 Branch Router data model

This section documents the fundamentals of Router naming. After introducing the main concepts some consequences are summarised..

A. A "Router" is identified by a unique and immutable serial number and represents a specific hardware instance.

B. A "Branch" consists of one or more locations where business may be transacted. It has a unique identifier – FAD code. One of these locations is termed the Storage Location. The storage location is unique for that Branch – no other Branch is associated with this Storage Location.

C. A "Storage Location" is somewhere that a valid (within Network closure – refer to 6.3.5 for definition) service termination point (*) exists. Semi-formally it is defined by the unique tuple consisting of service endpoints [PSTN (CLI), ADSL (IPStream SID), ISDN (CLI)]. At least one of the entries must be non null and any non null entries must not belong to any other tuple. So in other words a single entry in any Tuple is sufficient to identify a "Storage Location". (*) Note that at most one service termination endpoint of any technology type {ADSL, PSTN, ISDN} can exist at a Storage Location (this is a constraint to simplify provisioning)).

D. A Router has 4 possible configuration states;

- Gold Build – some version of a non branch specific configuration which attempts to "call home" when connected to any of {ISDN, PSTN, ADSL}.
- Personalised for Branch #B

- Current;- The Router has a configuration which will enable it to connect to the Production network and operate correctly at all locations for this Branch...
- Stale;- The Router has a configuration for this branch however it has been deregistered (refer to section 2.6.4.5)

E. For a given branch #B, only one Router should exist with a configuration state of current. This is enforced within the Provisioning system (refer to section 2.6.4.5).

F. The Attributes for identification of network service endpoints {CLI, PSTN, and ISDN} are required to be known in advance for Storage Locations for the purposes of provisioning. These attributes are not required to be known within the network because there is no checking of these attributes as part of PPP Authentication process (refer to section 6.4).Note however that the CLI attributes are recorded in Radius logs for purposes of audit / diagnostics.

G. The Wireless WAN network service endpoint is the SIM which has associated phone number MSISDN. This is presented as part of Radius Authentication. This phone number is recorded in inventory as associated with the Branch Router.

H. The PPP usernames and IPSEC usernames are associated with the Router. For PPP usernames this is a fixed prefix per interface followed by the serial number. This is fundamental as it is the Router which is authenticated to join the network and not the Branch. To illustrate this point consider the case where 3 stale Routers exist for a branch. None of these should be able to authenticate to the network.

Consequences of this model are:

1. A Router can only be provisioned automatically at a Storage location. An exception process is in place to deal with a mismatch between the known CLI and actual CLI presented. This exception process is based on associating the Router serial number with the presented CLI for a period of time having validated this is legitimate via a manual process.
2. The Network design is based on one Current Router per Branch. Exceptions will result in potential denial of service type disruption. Deregistration makes this unlikely to happen.
3. The DNS name for the Branch Router is defined as part of the DNS & Naming High level design activity. The name is specific to the Branch and not the Router.

Radius Session tracking is not used since IP stream provider does not support periodic Radius accounting

From a Support perspective the Branch Router is identified using its LAN side IP address (LR) and this is reached via any functioning WAN interface. The decision to which WAN interface to use for an IP datagram is made by the Access layer as covered in section 2.3.5.9.

*** Service termination points**

ADSL service termination points are identified by SID as specified in (BT STIN 456). The SID is known in advance and can be presented to the user radius depending on the top level domain presented in the username.

For PSTN and ISDN service termination points are identified by CLI. These are known in advance,

2.7.2 Application IP Endpoints

The following diagram illustrates the IP endpoints for applications.

IP address properties are provided in
Table 56 – Branch IP Address Properties.

The key properties are

1. The Data Centre application has visibility of the Branch Counter LAN IP address for all interactions with the Branch Counter. Therefore it can infer which Counter is initiating traffic flows. Note Counter 1 always has the same address etc.
2. Outbound connections can be initiated from the date centre to applications listening on well known ports on the counters. For example use of SSH by SSC for support.

In the diagram below;

- Lr and L2 represent individual IP addresses on subnet L
- {A} is the collection of interface addresses for all servers which are endpoints for branch application traffic
- {L} is the collection of all possible Counter Interface addresses
- {S} is the collection of interface addresses for all servers which are endpoints for Sysman traffic
- {V} is the collection of all possible Virtual IP addresses used by the Counter Application

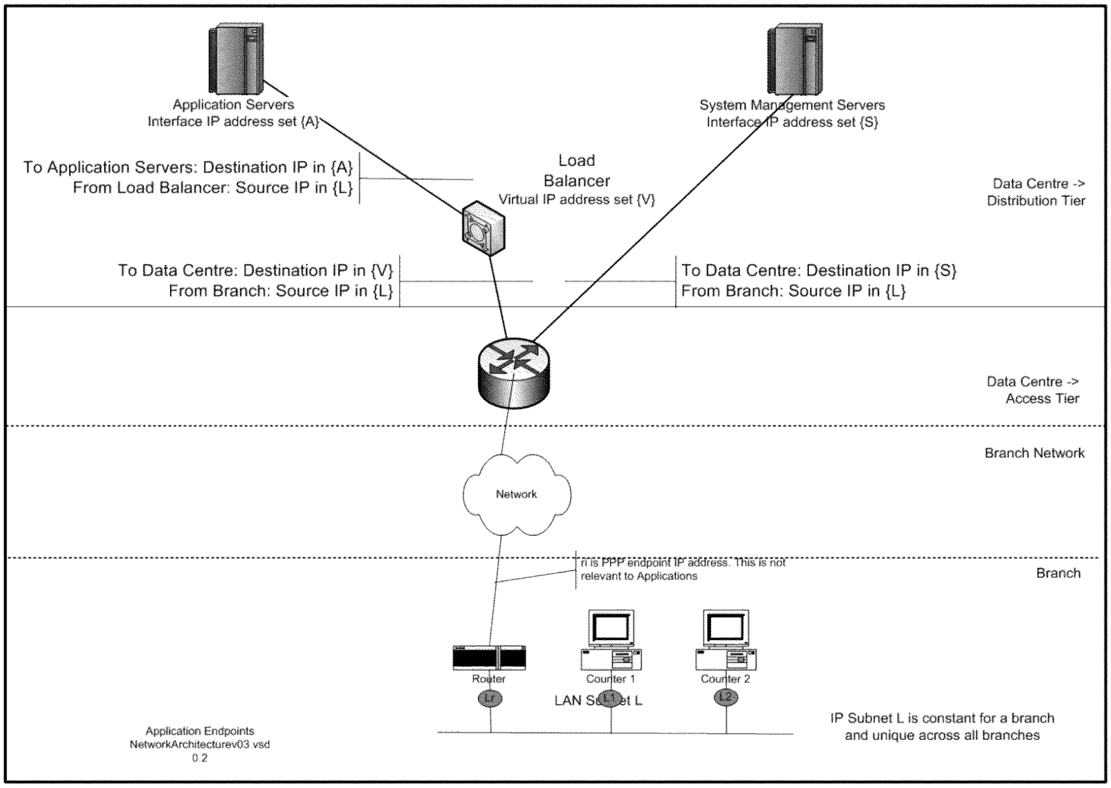
Application Servers
Interface IP address set {A}
System Management Servers
Interface IP address set {S}
To Application Servers: Destination IP in {A}
From Load Balancer: Source IP in {L}
Load
Balancer
Virtual IP address set {V}
Data Centre ->
Distribution Tier
To Data Centre: Destination IP in {V}
From Branch: Source IP in {L}
To Data Centre: Destination IP in {S}
From Branch: Source IP in {L}
Data Centre ->
Access Tier
Branch Network
Network
ri is PPP endpoint IP address. This is not
relevant to Applications
Branch
Router
Counter 1
Counter 2
Lr
L1
L2
LAN Subnet L
IP Subnet L is constant for a branch
and unique across all branches
Application Endpoints
NetworkArchitecturev03.vsd
0.2

Figure 12 – Application IP Endpoints

2.7.3 DNS

A Resilient DNS service will be provided for name resolution. Further details will be provided in the DNS High Level Design and this will cover use of Canonical names, the server hierarchy, domain structure, naming scheme, interaction with Active Directory and location of DNS platforms. In addition this document will cover security considerations, specifically controls.

The DNS will be used as follows;

- All data centre platforms will used domain names to locate service endpoints
- The DNS service will be virtualised using functions in the multilayer switches. A single constant IP address is configured for each server. The goal is to eliminate / significantly limit the need editing of hosts files on any platforms.
- Support staff can refer to all counters, network appliances and the Branch Router by name.

Note that the DNS is not used for Counter interactions.

2.7.4 HNG-X IP Addressing

2.7.4.1 Principles

1. All values according to RFC1918, therefore in {10.0.0.0 - 10.255.255.255 (10/8 prefix) , 172.16.0.0 -172.31.255.255 (172.16/12 prefix), 192.168.0.0 - 192.168.255.255 (192.168/16 prefix)}. Note that the VSAT addresses will be incorporated into the 10. range.
2. Distinct ranges for following domains;
   2.1. Data Centre LAN (domain is data centre)
   2.2. Inter Data Centre
   2.3. Virtual IP addresses
   2.4. WAN (POA owned remote site LAN Core VPN WAN)
   2.5. Inter-WAN (data centre to POA owned remote site (p2p) ,Core VPN WAN
   2.6. AS peer address space (NAT etc)
   2.7. Loop back addresses
   2.8. Branch WAN
   2.9. Branch LAN
3. It is possible for support staff to readily identify where an address belongs with respect to listed domains
4. Addresses can be summarised to respected domain
5. Branch traffic can be partitioned for scaling
6. Classless Inter Domain Routing (CIDR) is the method used for assigning IP addresses
7. Masks on octet boundaries are preferred as these make it simple to interpret IP addresses in preference to efficient usage of address space

8. Orthogonal to existing Horizon IP address space therefore allowing coexistence. . Note that Horizon IP addresses will need to exist in the HNGX data centres whilst Horizon remains operational.

2.7.4.2 Application

#	Domain	Address topology
2.1	Data Centre LAN	256 subnets per data centre of size 254 hosts IRE11 Summary 172.16.0.0 / 16 Subnets 172.16.0.0 /24 through 172.16.255.0 /24 IRE19 Summary 172.17.0.0 / 16 Subnets 172.17.0.0 /24 through 172.17.255.0 /24
2.2	Inter Data Centre	Summary 172.18.0.0 /16 Subnets 172.18.0.0 /24 through 172.18.40.0 /24 Total 41 subnets of size 254 hosts
2.3	Virtual IP addresses	Summary 172.19.0.0 /16 Subnets 172.19.41.0 /24 through 172.19.140.0 /24 Total 100 subnets of size 254 hosts

#	Domain	Address topology
2.4	WAN (POA administered remote LAN)	Summary 172.20.0.0 / 16 For each traffic class, such as DVLA, BRA01 support, E-pay etc index the traffic class by i. Up to 255 traffic classes are allowed. Summarise traffic class by 172.20.i.0 / 24 Within each traffic class there are 16 subnets with 14 hosts per subnet, Subnet mask is / 28 For example for DVLA suppose the traffic class is 1 then we have the subnets shown in Appendix A. These subnets are use to cover Remote LAN and WAN subnets
2.5	Inter-WAN (data centre to POA owned remote site (p2p) and Core VPN WAN	As for 2.4 but using 172.21.0.0 / 16 *Identify data centre, if applicable, by overloading last octet of subnet, e.g. < 128 is IRE11.*
2.6	AS peer address space (NAT etc)	As for 2.4 but using 172.22.0.0 / 16
2.7	Loop back addresses	Totals hosts =256 * 2 (assuming subnet mask is /32) Allocated from{192.168.20.0/24, 192.168.21.0/24}
2.8	Branch WAN (Hub & Spoke model)	10.x.y.z as specified in Section 2.6.3.3.
2.9	Branch LAN	10.x.y.z as specified in Section 2.6.3.3.

2.7.4.3 DVLA subnets example

Subnet	Mask	Subnet Size	Host Range	Broadcast
172.20.1.0	255.255.255.240	14	172.20.1.1 to 172.20.1.14	172.20.1.15
172.20.1.16	255.255.255.240	14	172.20.1.17 to 172.20.1.30	172.20.1.31
172.20.1.32	255.255.255.240	14	172.20.1.33 to 172.20.1.46	172.20.1.47
172.20.1.48	255.255.255.240	14	172.20.1.49 to 172.20.1.62	172.20.1.63
172.20.1.64	255.255.255.240	14	172.20.1.65 to 172.20.1.78	172.20.1.79
172.20.1.80	255.255.255.240	14	172.20.1.81 to 172.20.1.94	172.20.1.95
172.20.1.96	255.255.255.240	14	172.20.1.97 to 172.20.1.110	172.20.1.111

172.20.1.112	255.255.255.240	14	172.20.1.113 to 172.20.1.126	172.20.1.127
172.20.1.128	255.255.255.240	14	172.20.1.129 to 172.20.1.142	172.20.1.143
172.20.1.144	255.255.255.240	14	172.20.1.145 to 172.20.1.158	172.20.1.159
172.20.1.160	255.255.255.240	14	172.20.1.161 to 172.20.1.174	172.20.1.175
172.20.1.176	255.255.255.240	14	172.20.1.177 to 172.20.1.200	172.20.1.201
172.20.1.202	255.255.255.240	14	172.20.1.203 to 172.20.1.206	172.20.1.207
172.20.1.208	255.255.255.240	14	172.20.1.209 to 172.20.1.222	172.20.1.223
172.20.1.224	255.255.255.240	14	172.20.1.225 to 172.20.1.238	172.20.1.239
172.20.1.240	255.255.255.240	14	172.20.1.241 to 172.20.1.254	172.20.1.255

3 Platforms

DN: Need to use common definition for Platform.

By definition a platform is something which has Enterprise management components (SYSMAN2) installed. This section addresses the requirements of the architecture in terms of dedicated platforms and/or providing components that will coexist with other software applications on generic platforms. Support of the components on Operating system genres is identified.

In addition any overall architectural requirements for storage including back up and disaster recovery are stated.

3.1 Dedicated platforms

This table lists those platforms which are required for the Network Architecture.

DN: Need to include Test Platforms. – those in IRE11/19 and those elsewhere (if applicable)

Function	Number	Type	Storage	Backup	Rebuild
ADSL Radius	1 per Data Centre in Active blade frame.	Radiator Windows 2003 Server	73 GB	Yes – logs	Required -
GPRS Radius	1 per Data Centre in Active blade frame.	Radiator Windows 2003 Server	73 GB	Yes - logs	Required
Dialled Radius Including dial out profiles	1 per Data Centre in Active blade frame.	Radiator Windows 2003 Server	73 GB	Yes - logs	Required
Accounting Radius	1 per Data Centre in Active blade frame	Radiator Windows 2003 Server	73 GB	Yes - logs	Required
Radius for Network Appliance login	One Platform instances per data centre as free standing platforms.	Cisco ACS 4.0 Windows 2003 Server	73 GB	Yes - logs	Required
Branch Router Access Control	1 per Data Centre in Active blade.	Cisco ACS 4.0 Windows 2003 Server	73 GB	Yes - logs	Required
DNS Refer to section 2.7.3. This entry will be filled in once design is complete.	Subject to model				
Network Sniffer	One free standing platform per data centre	Wireshark			

3.2 Platforms / appliances provided as part of managed service

NMS; NNM HP Open view Node manager	1	SunFire V890 (Sun4U), Solaris 5.9	DAS – 220GB	Yes – logs	Part of managed service
Alarm Point	1	Fujitsu RX300S2, Windows 2003 Server	DAS - 36GB	Yes – logs	Part of managed service
Cisco works	1	SunFire V280R, Solaris 8.	DAS – 73 GB	Yes – configuration & logs	Part of managed service
Logging server	2	Syslog NG		TBD	Part of managed service

Notes

1) IDS/IPS appliances have been omitted from the above table since they are treated as network devices.

2) DN: The Boot Radius is assumed to be an Estate Management platform and hence not introduced with the network solution space.

4 Network section

This section is included for numbering consistency and has no relevant content.

5 Manageability

This section documents how the network fulfils its responsibilities to support the overall manageability of the solution.

5.1 Network Management

There is a single Network management domain covering network components (including the fibre channel switches) at both HNG-X data centres. At each data centre there is a management platform (NNM) and these platforms operate in an Active / Active mode. Each NNM instance manages all HNG-X network appliances at its local data centre plus all Branches, WAN services and Remote locations.

Therefore, for example the HNG-X Routers at DVLA will be managed by both NNM instances.

5.1.1 Non Branch HNG-X network components

The following table provides a summary of Network management functions for all network devices other than the Branch Router;

Management Function	Approach
Logging & Events & Alerting	Logging: All network devices are configured to send SYSLOG messages to the logging system at the local data centre. The Logging system will forward these to both NNM instances. SNMP traps & Informs These are sent to each NNM instance Alerts These are generated to the NNM and forwarded to the Enterprise Management System via an SNMP gateway on the NNM.
Remote Operations	Support access, inband this is either SSH or Telnet over IPSEC tunnel or HTTPS (refer to 6.3.2).
Time Sync	All data centre components use a GPS receiver as the NTP time source. Non data centre components used Access tier switches as NTP time source.
Configuration / Inventory	Configurations are managed by Cisco works
Traffic paths	OOB access to Aurora via OOB firewall All network management over device LAN interfaces is in-band except for security appliances such as Firewall and IDS. There is no dedicated management network reaching all devices. All console ports connected to out of band network with audited access as for Remote Operations (Aurora).
Monitoring	ICMP Ping and SNMP polling from NNM

5.1.2 Branch Router

This section provides a summary of management functions for the Branch Router (Sarian DR6410)..

Management Function	Approach
Logging & Events & Alerting	The Branch Router uses a local application - CNIM2 which runs on every Counter for these functions. Refer to section 2.6.3.4 for details of CNIM2 functionality provided. In the case where CNIM2 received events from the Branch Router then it will use the Enterprise management system to forward events to the logging system in the Data Centre. Rationale for the approach of using a local application is that; • Alerts need to be preserved when no Network connection exists to the data centre. • Event storms need to be avoided and CNIM2 will randomise intervals between sending events and throttle events being sent • Multiple sources (events and content of Routing updates) can be correlated to create single composite events / alerts for more efficient diagnosis of problems. The logic for performing this correlations runs outside the Branch Router. To deal with the scenario where in a single counter office, the LAN connection between the Branch Router and the Counter is down, the Branch Router will send any SYSLOG messages relating to LAN connection transitions to the Data Centre. To provide some degree of protection against event storms, the number of such events will be limited to 10 per hour. *(DN: Need to ensure that the IPSEC head end service will support this load – scaling horizontally using VRRP+). Security related events will also be forwarded to the data centre.* In this case the event destination will be the Logging system. A virtual IP will be used as the destination so that the branch Router does not need to be aware of which data centre to target. CNIM2 will also detect when a Counter looses communication with the Branch Router since CNIM2 runs on every Counter. It will log this occurrence and use SYSMAN to alert when this has happened. The SYSLOG messages from the Branch Router will travel over IPSEC even though they do not contain passwords. This is because they may contain other information which is of value to profiling the betwork such as IP addresses, phone numbers. The path from the VPN concentrator to the SYSLOG server is not encrypted.

Management Function	Approach
Remote Operations	These can be performed interactively or via scripts controlled by SYSMAN3 components. The Network paths are; 1. In band – same network session as used for Branch traffic 2. Out of band -PSTN dial up In both instances the support traffic must be encrypted as passwords are carried in the payload; The normal approach will be to use an IPSEC tunnel between the data centre VPN concentrator and the Branch Router. This enables protocols such as ftp, telnet and http to be supported. The within data centre path from the source endpoint to the VPN concentrator will also be encrypted using IPSEC to avoid passwords in clear. The Data centre platform which is the endpoint for support traffic is protected by a Firewall. To meet the requirement that no Encrypted traffic passes through a Firewall, the Firewall will End Point Platform-----FW-----VPN Concentrator------(WAN)-----BR BR = Branch Router End Point Platform-----FW always up IPSEC tunnel To deal with the case where the IPSEC is broken on the Branch Router, it will be possible to establish an SSH session to the branch Router from an SSH client on the SAS server. Interactive traffic can be either Http or Telnet (both over IPSEC) from the Router Operational support system (ROSS). Access to this platform is via the SAS server to provide audit and common access. (DN: Note the reliance on the SAS for RDP protocol version 6.0) Batch scripts on the ROSS are scheduled via SYSMAN and use Telnet and ftp (both over IPSEC) access the branch Router. SSC create perl script and prototype in live This would be a workflow in TPM which can be called by 1st / 2nd line

Management Function	Approach
Time Sync	Router connects to service on Branch Router Operational Support System using SNTP once every 24 hours Randomised over [20,28] hours.
Configuration, Inventory and provisioning	This are described in section 2.6.4
Back Network Testing and viability reporting	Periodic testing of ISDN backup networks is required this will be performed by CNIM using DSCP routing to bring up ISDN. At least one ICMP echo request / reply will be required for success. The outcome is reported via Enterprise management events.
Monitoring	ICMP Ping of Branch Router on LAN IP address from NNM. This will be a plain unencrypted ping. Note that NNM has no explicit information as to which network type is being used for a branch. However the Logging platform will have events detailing transitions. There is no plan to use SNMP access to the Branch Router since the Router does not support version 3. The Router will block SNMP access. Refer to section 6.3.2 for details of SNMP usage.

5.1.2.1 Wireless WAN viability monitoring

Within Branch estate management a variable is maintained for each Branch indicating whether this branch is defined to have viable WWAN backup for the purposes of SLT availability calculations. Put another way RMGA accept SLT's on availability for such branches. Note that this flag has no influence on the Router behaviour. For example the Router will be configured to always try and use Wireless WAN. The rationale is that this will "do no harm" as the Primary network is down as far as the Router is concerned. The Branch Router configuration design needs to ensure that this does not preclude failback to the Primary network by more than a few minutes.

This section proposes how this variable will be maintained.

Constraints / assumptions;

1. WWAN is not always up so cannot assume reachable from data centre without intervention.
2. Selection of outgoing network – priority is [ADSL, GPRS]

Every time a Branch Router is registered for service in a branch #B, if that branch has not yet had its WWAN test performed, a task will be performed to do this.

The task will perform the steps as outlined in the following table;

#	Summary
1	Scripted via Remote operations; Over inband management network Change Router configuration (in memory only) to nail up WWAN interface. Confirm Router is time synched to Data centre by using the NTP client on Router to check and if necessary perform an update. This covers both the case where ADSL is the Primary network and where the

#	Summary
	Router has failed over to using GPRS From the data centre ping the L2TP endpoint on Router – this is an in clear ping following. The test is that every 10 minutes; a sequence of four 100 byte, 4 1000 byte and 2 1480 byte ICMP echo requests. The timeout is set to 15 seconds. The test runs for 24 hours and the success criteria will be based on Packet loss, average and 95pctile response time (by packet size). In addition these values will be applied to both core hours and non core hours time periods. The test is deemed to have completed if there was connectivity measured at least once in every 24 hour period – so a response in each bucket. In the case of a completed test the flag will be set to "YES" or "NO" based on the defined criteria to indicate whether this branch is defined to have Wireless WAN backup for the purposes of SLT availability calculations. For each test a log will be produced containing a time stamped sequence of all relevant events – ICMP echo replies and outcomes, signal strength and Wireless network transitions the latter two being available from the branch router and CNIM event logs. All time stamps will be at source meaning that they will be from 2 different sources but should be synchronised to within a few seconds. The branches where the signal strength exceeds a defined value (for example -86dBm but to be agreed with service provider and Sarian) where the test failed will be investigated further with the Wireless WAN provider as this result is unexpected. This may result in a further test being scheduled. Non complete tests will be rescheduled manually.
2	Test time period complete; The Router changes will be undone to revert it to normal operation. This may require a reboot; - This decision is subject to Design work.

On an ongoing basis it will be necessary to monitor whether there are any changes to the WWAN viability. The methods proposed for this are detailed below;

Note that each Counter within the Branch can determine which network type {ADSL, PSTN, ISDN, WWAN} is being used (subject to recent transitions) and transitions are reported in the event log.

As part of "Business as usual", a Postmaster will phone in as a consequence of network down (this means no communication for at last 10 minutes with suitable alert on screen). As part of the diagnostic process it may be determined from the UI on the Counter that the current network type is Wireless WAN and this will also include the signal strength and network type. The helpdesk will know from the inventory for the branch whether Wireless WAN is defined as viable. If yes and this is not consistent with the diagnostic outcome then an issue should be raised for investigation. The outcome may be a manual correction to the Wireless WAN viable flag or rescheduling of the Wireless WAN test as a result of changes by the service provider or at the branch – Antenna repositioned.

Note to provide suitable diagnostics for issue investigation It is assumed (but this needs to be confirmed) that transactions successes and failures within a branch are being logged locally and retained for a period of time – ideally 1 month detail, 3 years in hourly summary form. This will provide a history of behaviour of the Wireless WAN network.

Building on the above assumption, if transaction failures are being counted then Events should be raised on criteria when a significant number of transaction shave been attempted over Wireless WAN and a significant proportion have failed. For example if every other transaction failed and at least 100 were attempted. The events would be used to trigger investigations similar to the Help scenario above.

5.1.2.2 Branch Router Diagnostic sources

Details on each of the diagnostic sources discussed in the above diagram are covered in the following table.

#	Summary
1 - G/W UI	The User Interface is available to the Post Master and Engineer to interrogate CNIM and initiate a PING test to the data centre.
2 - Router Lights	The Router has multiple lights which indicate the state of operations. For example the DSL light could be used in conjunction with (6 POM reboot) to determine if the ADSL service is available at the branch.
3 – Tivoli Health checks	These exist for the current Gateway and need to be amended to include the Branch Router
4 – Historic data	This covers data maintained at the Data centre for the branch. It includes; • Radius logs (Authentication / Accounting / interim accounting - not ADSL) • Result of NNM periodic polling of branch • Quality of service records produced by CNIM2 • Events created by CNIM2 for the Branch Router and forwarded to the data centre
5 G/W UI	As per (1) but use of CNIM2 on the slave. This enables the case to be determined where for example the Gateway cannot see the Branch Router but the slave can.
6 POM reboot	Whether reboot of the branch router (if determined appropriate by support unit) solves the communication problem. If not it may provide further evidence as to where the problem lies – for example flashing DSL light indicating ADSL training problem.
7 Engineer on site	An engineer is sent to site to connect out of band support modem. The router event log is pulled back and analysed by script to check for a number of predefined problem cases.

5.1.3 Network probes

The approach to providing network traces is;

1. LAN traffic – specific platform exists with Wireshark software. Use of this technology will be coved by process agreed with Security manager as full packet capture is possible..
2. WAN Traffic – X.25 only by agreement with POA security Manager. (DN: Need to ensure technology is in place for this)
3. Branch Router – tracing on demand – note full packet is captured but application traffic is SSL protected therefore is permissible (*Dn: Need to validate with security*).

6 Security

This section documents the mechanisms by which the Network Architecture meets requirements for providing a range of security controls.

It covers;

1. Network domains model for network segmentation and enforcing rules on traffic flows (section 6.1). The purpose of this model is to provide a framework for specifying the security enforcing functionality in the Network.
2. Network controls (section 6.3).
3. Network closure (section 6.3.5)

6.1 Network Domains Model

Need to explain what DMZ means and what security zones mean update this section based on work with Jim

The term Network domain is defined to mean a collection of platforms and network components which can communicate together without any network based securing enforcing controls applied to this communication path. For example two servers in a DMZ are defined to be within the same network domain. Any traffic which crosses network domain boundaries passes through an Enforcement point which is a Firewall.

The following diagram illustrates how Network Domains fit within the Network tier model (refer to Figure 2 – Network Model).

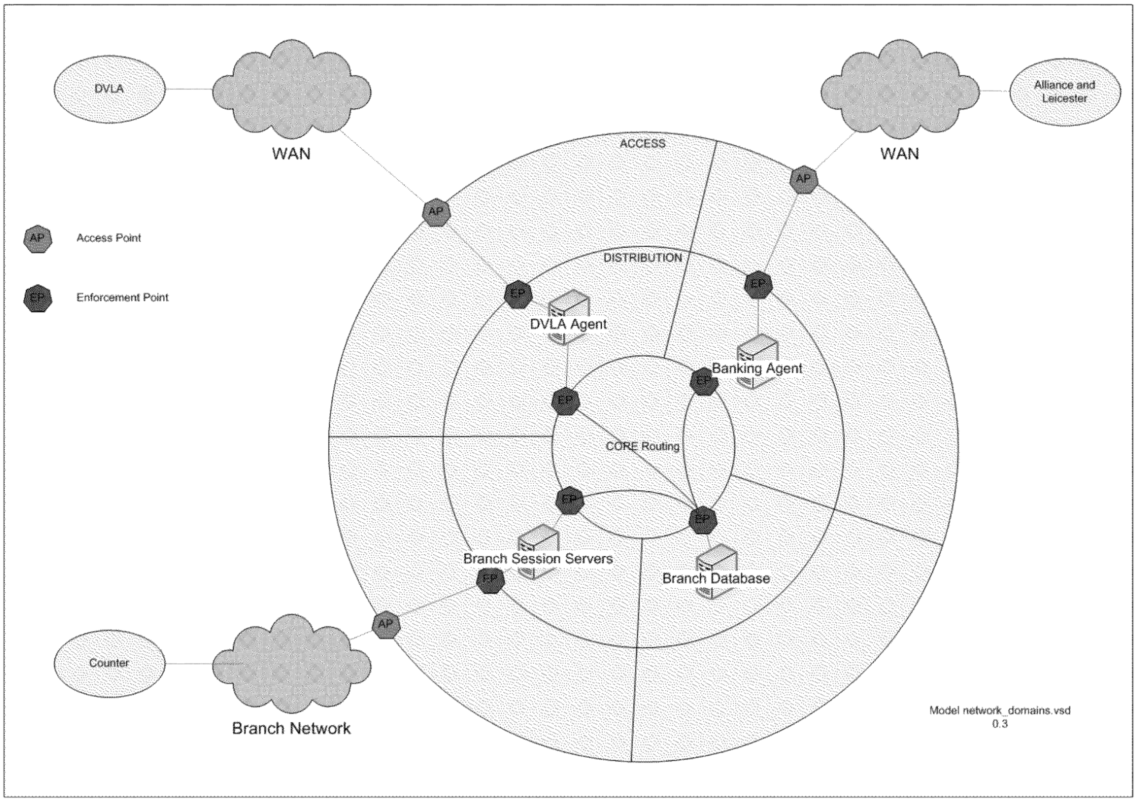

Figure 13 – Network Domains

Enumerate Network domains based on latest list from Jim

Network Domains are the basic building blocks for enforcing security in the Network. The following rules are applied to use of Network Domains.

Property	Description
Communication flow restriction between network domains	Communication flows between platforms (* see comment below on Virtualisation) and Network components are constrained by the hierarchy shown in Figure 2 – Network Model. Communication is only possible upwards (towards the core) and downwards (away from the core). This means that for example, that any platforms in the Branch Session Server network domain cannot communicate directly with any platforms in the Branch Database network domain. Rather they are constrained to communicate via the Core Routing Tier and associated Enforcement points. * Virtualisation creates multiple logical platforms within a managed Virtualisation server. Each such virtual platform will be in a single network domain. The Virtualisation server applies security controls to ensure separation of Virtual platforms and VLANS.
VLAN network domain membership	A VLAN cannot be in more that one Network domain

Property	Description
Platform network domain membership	Each platform belongs to one Network domain. The term platform applies to virtual platforms within the Virtualisation server and real platforms.
Firewall network domain membership	Firewalls may be present in multiple Network domains. Sub interfaces (so VLANs) within the firewall will be dedicated to Network domains.
Enforcement points and traffic forcing	Enforcement points are Firewalls and all traffic passes through these when traversing Network domains
Physical separation	The only allowable Network components on the boundary between the Distribution and Access Network tiers are Firewalls and in this case, a single physical interface on such a Firewall can only be present within one Network tier
Firewall types	Firewalls in the Access / Distribution boundary will be a different models to those used in the Distribution / Core boundaries.
Firewalls	All Firewalls used in the network shall be certified at EAL4+

6.2 Traffic classes, flow constraints and separation

6.2.1 Traffic classification

Traffic classes are use to enumerate all possible traffic flows into and out of HNG-X data centres for purposes of describing security properties. All such flows ingress and egress the HNG-X data centres via the Access Tier.

The Traffic classes are listed in the following table

Traffic Class	
Live	Test
CAPO	Test_CAPO
LINK	Test_LINK
E-PAY	Test_E-PAY
DVLA,	Test_DVLA,
A&L,	Test_A&L,
Streamline online	Test_Streamline online
Streamline batch	Test_Streamline batch
Moneygram	Test_Moneygram
Post Office	Test_Post Office
Support	Test_Support
ADSL Branch	Test_ADSL Branch
Dialled Branch	Test_Dialled Branch
GPRS Branch	Test_GPRS Branch
VSAT Branch	Test_VSAT Branch
Branch Router Support Traffic (IPSEC)	Test Branch Router Support Traffic (IPSEC)
Branch Router Support Traffic (ICMP)	Test Branch Router Support Traffic (ICMP)
Network Management	

Table 6 – Traffic classes

For example the DVLA traffic class covers the network flow between HNG-X endpoints located in the Distribution Tier and DVLA sites for Application traffic;

6.2.2 Flows constraints

Each Traffic class is treated as follows in the Access Tier

Ingress

Network Termination-> IPSEC Head End (if applicable) -> Firewall in Access Tier ->SSL Head End (if applicable) - > IDS sensor

Egress

This is the reverse of the above path

Network Termination includes basic security controls (Access control lists). Specifically if a Traffic class contains a traffic flow which should be IPSEC then this will be policed.

In each case the Firewall will limit the Traffic class to specific IP address ranges, protocols and ports ranges. State full rule will control the direction in which TCP connections can be initiated. All blocked traffic will cause alerts, subject to an upper bound at which alerts can be generated.

Similarly the IDS sensor will alert exceptions and violations.

The following diagram provides a Logical overview of Traffic from the Access Tier through to network domains. It shows the position of the Firewalls, IDS sensor, IPSEC termination and SSL termination.

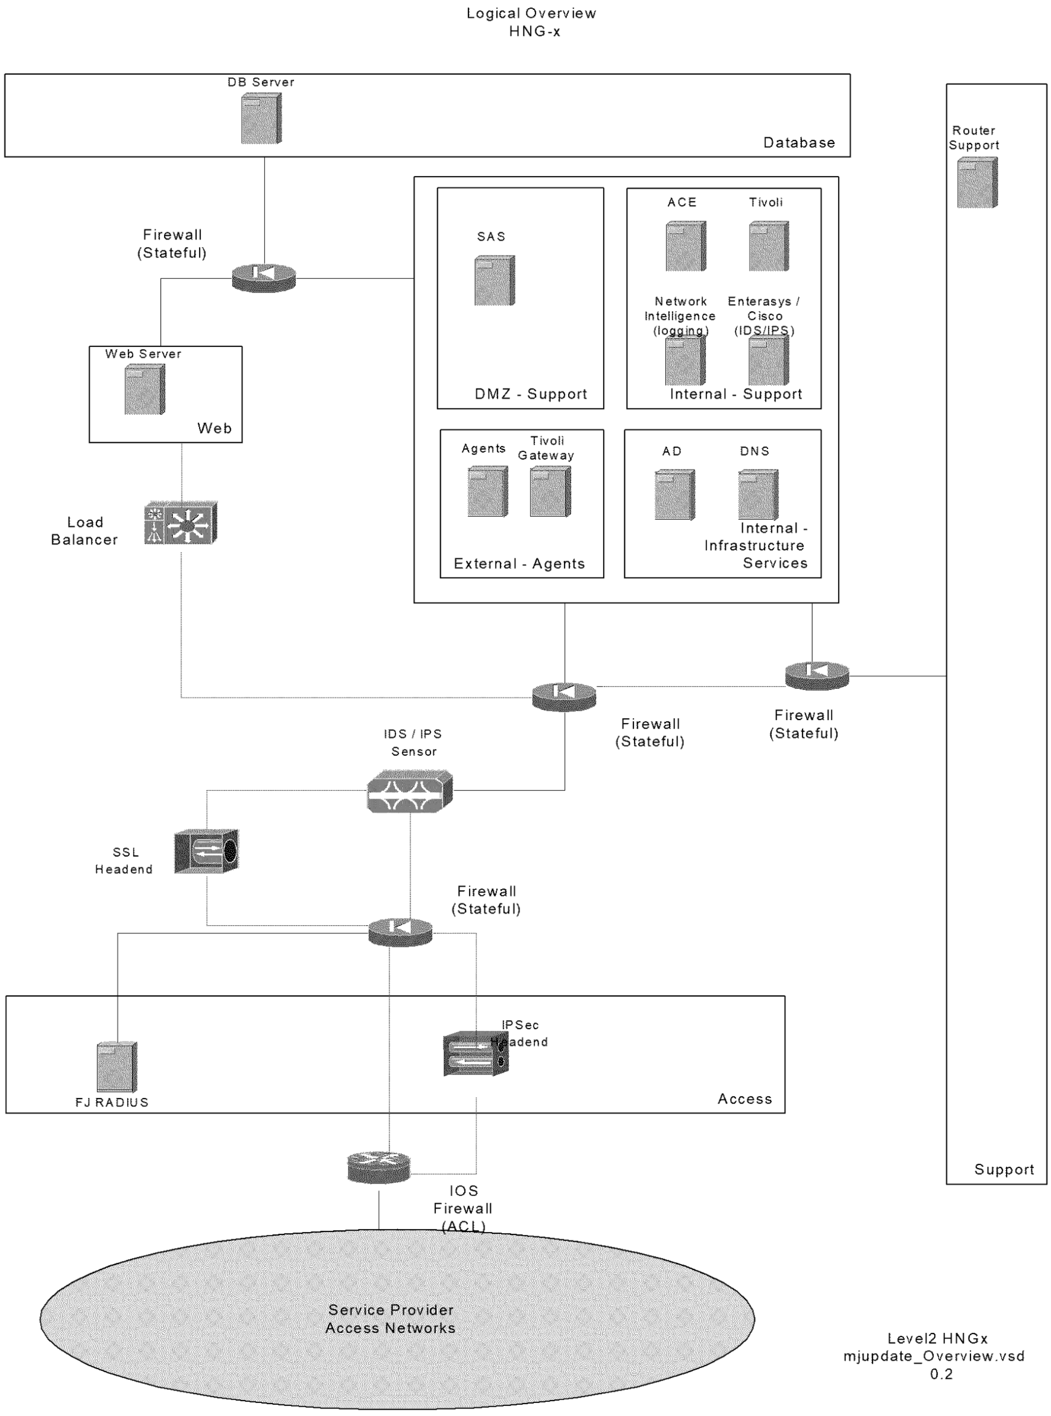

Figure 14 – Security overview

6.2.3 Traffic Separation

The following table specifies the approach for separation of Traffic classes (listed in Table 67 – Traffic classes).

Traffic separation	Traffic classes are separated as follows; Distribution Tier – Traffic class endpoint in dedicated Network domain. Access Tier – LAN side One of the following approaches; ▪ VLAN separation of traffic so that a given VLAN only carries one traffic class ▪ Dedicated Ethernet ports Access Tier – Wide Area Network side For each traffic class one of the following approaches; ▪ Dedicated MPLS VPN ▪ Dedicated IPSEC tunnel over Shared MPLS VPN. ▪ Dedicated circuits (leased or point to point dialup) – for example Streamline Note that VRF's (Virtual Route forwarding) in data centre environment are not required since the IP address space is administered from a single domain. Note that Network domains are not the same as traffic classes. For example both the DVLA an E-PAY network domains will have traffic flows into the branch database network domain.

Service boundary	Every traffic class has is a service boundary which represents the demarcation between HNG-X infrastructure and 3rd party infrastructure. Depending of the 3rd party, the service boundary will be located at one of the following locations; a) HNG-X data centres b) 3rd party data centres c) Fujitsu points of presence with onward carrying of traffic over the FSBN In each case the traffic from these locations will pass through a Firewall in the Access Tier. For example, in the case of CAPO, LINK and Moneygram, the service boundary is local to the Data Centre. The Connection will be HNG-X Firewall -> Layer 2 switch -> 3rd party Ethernet port.
Separation from other networks – for example Internet	Such traffic must pass through an external (to HNG-X) non Cisco firewall

Table 7 – Traffic Separation

Note that since any Production class is be definition separate from any Test class the mechansism abobe will be used to separate the traffic.

(DN: Need to add some points about isolation and include control plane)

6.2.4 Traffic Policing

The Multilayer switches used in the Distribution and Access Tier are configured to apply traffic policing;

- Protection of Control plane by controlling type and rate of packets that consume router CPU resource – for example fragmented ICMP echo requests
- Traffic Storm Control
- Rate limiting of test traffic classes on ingress / egress, layer 3 switching and control plane

6.3 Network Controls

6.3.1 Identity and Audit

Summary	Description

Summary	Description
Authentication - Logon to Network Components* via LAN interface is via (SSH) * Not Branch Router	Interactive traffic always forced via SAS Protocol; TACACS+ Device -> Cisco Secure Access Control Server (ACS) User names in Windows Active Directory (AD) Two factor authentication
Branch Router – logon via WAN interface using http / telnet carried over IPSEC	Device -> Cisco Secure Access Control Server (ACS) Protocol; RADIUS User names in Windows Active Directory (AD) Single factor authentication - password
Authentication - Logon to Network Components via Serial Port	Always via Aurora at data centre Will try ACS (as LAN) but fall back to local accounts. Principles; • Generic username, separate last resort password per device for console access only • When password used then it is changed • Key under lock and key DN: Explain local account in more detail Under this scenario the Network component being accessed may be unlikely to be able to access the ACS. Should this be the case then local accounts will be used as backup where the passwords are device specific by role and issued under a defined process.
Command Audit	All commands issued by the users when they are logged in to network components are logged. The logging is performed by the command audit feauture:

6.3.2 Network Management

Management Traffic Command line	All Command line Management traffic supporting command access is encrypted – either via SSH, IPSEC or HTTP over SSL.
Monitoring Traffic	This comprises ▪ SNMP version 3 ▪ ICMP Echo request / reply Note that SNMP community strings are treated as passwords and therefore only SNMP version 3 is supported. (Otherwise it would be necessary to encrypt all earlier versions of SNMP from the NNM which adds significant complexity). Note that need to get a let from security on the ASA devices as they do not support SNMPv3. ASA; The adaptive security appliance provides support for network monitoring using SNMP V1 and V2c. The adaptive security appliance supports traps and SNMP read access, but does not support SNMP write access.
Audit / Logging / Alerts	All network components will send logging information (Syslog) to a single logging endpoint at the local data centre. Behind this logging endpoint there will be a resilient platform for persistence storage and analysis of logs. *DN; The likely software used is Syslog NG* Note that the Branch Router will send logging information locally as well.
Non HNG-X access	Network components and supporting platforms (such as Cisco works) will not retrieve data from external web services. *(DN: A process will be defined to cover application of firmware and operating system updates).*

6.3.3 IPSEC

The use of IPSEC for traffic classes was stated in section 6.1.

Summary	Description
Encryption	Traffic classes will be encrypted over the Wide Area network if specified by the relevant TIS or HNG-X Security Architecture. The IPSEC tunnel will terminate on devices within the HNG-X service boundary. For example this is the case with A&L where HNG-X Router are at A&L Data Centres. The encryption algorithm will be AES 256.

Summary	Description
Authentication	Based on Certificates (except for branch Router) Branch Router will use PSK. The rationale for this is resuse of the Branch Router CHAP solution for key management. The IPSEC keys will be 15 characters in length from alphabet {A-Z, a-z, 0-9}. Entropy source will be same as used for CHAP – Hardware Random Number Generator.
Protection against single configuration error	IPSEC devices are deployed in the following topology. IPSEC Router -> Downstream Router - > WAN The downstream Router will apply an ACL to ensure that the traffic from the IPSEC router is IPSEC traffic only. Also the IPSEC Routers will be configured not to be able to negotiate an NULL encrypted stream. Therefore a single configuration error will not compromise WAN encryption.

6.3.4 SSL

Summary	Description
Encryption	SSL is used in transport mode between the client application on each Counter and the SSL Head end in the HNG-X data centre (refer to Figure 14151614 – Security overview). All transaction data travels over SSL secured TCP connections. The Cipher suite to be supported are; RSA_WITH_AES_256_CBC_SHA RSA_WITH_3DES_EDE_CBC_SHA (DN: Need to confirm which one is to be used)
Authentication	This is based on certificates issued by a certificate server to the SSL head end. The Network administrator will generate the private/public key pair for the router. The private key cannot be seen by anyone, including the administrator. The Certificate server will be a central authority for deciding whom to issue, revoke, etc. the certificates. This server will be managed by the Crypto team providing a clear separation of management.

6.3.5 Other Network controls

These will be defined as part of the Network Security HLD and cover amongst others; - Anti spoofing, no proxy ARP, no source IP routing.

6.3.6 Blade frame controls

DN: Need to specify why it is valid for single chassis to span security zones.

6.4 Branch Network Closure

This section documents controls within the Branch network to limit connectivity.

6.4.1 Network Service properties

The following table summarises what controls exist for service end points connecting to the HNGX data centres. This section only considers restrictions which prevent a session request being sent to the Data centre. Session level security controls are covered in section 6.4.3.1.

Network Type	Connection Incoming to Data Centre Initiated by Branch	Connection Incoming to Branch Initiated by Data Centre
ADSL These controls are implemented by network supplier. The service is IPStream provided by BT	Only known service endpoints (PSTN service with ADSL delivered over this PSTN service) can initiate sessions to the Data Centre. These service endpoints will be restricted to branches which are live. This restriction is a feature of the chosen ADSL service;- BT IP Stream Note that the BT ADSL SID service is not yet launched. Should it then it will be enabled so in Radius but not used	Session cannot be initiated from the data centre. This is a property of the underlying network service.
Dialled (ISDN)	No CLI checking	Incoming calls to Branch Router are rejected. This will cause the Router to send a SYSLOG message which in turn causes CNIM to send data targeted at the data centre resulting in the Router attempting to dial the data centre and establish a network connection. (This above mechanism is termed dial back it should not be confused with PPP dial back, It is used to provide manageability in ISDN branches as these are not permanently connected)

Network Type	Connection Incoming to Data Centre Initiated by Branch	Connection Incoming to Branch Initiated by Data Centre
Dialled (PSTN)	No CLI checking	Incoming calls to Branch Router are for "out of band support". This "out of band" route is for support staff to access a non communicating Router whilst an Engineer is on site. This requires PSTN line to be attached by Engineer. Note that this mechanism can be used both during provisioning and normal running Branch Router modes.
GPRS / UMTS Network supplier implemented through use of SIM	Valid SIM – registered to APNs belonging to Horizon. Note the key distinction where the SIM grants access rather than a service termination point	Session cannot be initiated from the data centre. This is a property of the underlying network service.
VSAT Broadband	Valid encryption key (PSK) in site Router enables IPSEC tunnel to be initiated	Valid encryption key (PSK) in VPN concentrator enables IPSEC tunnel to be initiated to be branch

6.4.2 Provisioning security

The term provisioning is defined to mean a site visit with Router for the purposes of installing that Router at the site.

A Branch Router can only be provisioned over a Network Technology which identifies the branch to the Provisioning system (via Radius attributes) at the HNG-X Data Centre. The technologies that provide this identification this are PSTN and ISDN which present the CLI. The ADSL IPStream service MAY have similar capability in the near future where a session identifier (SID) will be presented. Should this be the case then provisioning over ADSL will be the preferred method. Both the CLI and SID have the properties that they are known in advance, unique and constant; hence they can be used to identify the location where the Router is initiating a session from. This location determines which configuration the Router should have. The use of CLI and SID provide a level of assurance that the session is being initiated from a valid location. Mention exception process The section has been updated to allow for an exception process (Enginr phone help desk to let branch router with serial number S provision without CLI type attribute from network)

The Generic Router has a "Call Home" configuration which includes PPP CHAP Secrets and IPSEC secrets. These will be used to authenticate Router to the provisioning system and establish an encrypted Tunnel in order to protect passwords.

The Provisioning system will download a personalised configuration to the Branch Router which includes specific key material for IPSEC and CHAP. The provisioning system will also ensure that any previous Routers which have been provisioned at this location are not able to initiate session to the data centre. This is achieved by near real time updates* to the HNG-X Data Centre Radius servers. Therefore when engineer swaps a Branch Router, the one being replaced is longer usable.

Potential Denial of service where large numbers of Routers become unusable is mitigated by limiting;

- The number of Locations that can provision per day to 200 (configurable)

- The number of provisions per day per Personality to 5 (configurable)

Refer to section 2.7.1 for a definition of the terms Location and Personality.

* External authentication may be a better way of achieving this. The method will be determined as part of a design study.

6.4.3 Branch Router controls

6.4.3.1 Session establishment

This section covers session authentication during normal operation as distinct from provisioning which is covered in section 6.4.2.

The Branch Router and Data Centre mutually authenticate using a shared secret. The protocol is PPP CHAP and at the HNG-X Data Centre the CHAP secrets are stored protocol in Radius servers. To ensure separation of the network types, there are separate Radius servers for:

- ADSL
- Dialled (ISDN/PSTN)
- GPRS/UMTS

The security of the password used for the CHAP authentication is important and a robust solution that handles the passwords is required. The key requirement is that the risk of compromise of an individual password is low and that the risk of compromise of multiple sites must be almost non-existent.

The way that this is handled is that passwords are always transported and stored in encrypted form and that the full list of (encrypted) passwords is never allowed outside of the secure data centres.

For every branch there will be a unique CHAP secret which is the same across all network types {ADSL, Dialled, and GPRS}. The CHAP secrets are unique (up to randomness) across all Branches. This means that for 15,000 branches there will be a 45,000 CHAP secrets and they will almost all be different*.

Associated with each CHAP secret is a PPP username. This username includes the identity of the location and has a suffix which is used to provide a mechanism for non synchronised change of CHAP secret. For example if the current username in a branch is H12345A, a new secret is created at the data centre and the user suffix is set to B. Whilst changes are being rolled out to Branch Routers the Branch Router can authenticate with either H12345A and the old Chap secret or H12345B with the new Chap secret.

* The secrets are generated using a random source therefore with 11 characters, a few collisions may occur by chance.

6.4.3.2 Management traffic

All Network Management traffic between the HNG-X data centre and the Branch Router travels through IPSEC tunnels with AES 128 encryption. IPSEC shared secrets are used and these are unique (subject to collisions) per branch. IPSEC secrets are subject to the same handling as CHAP secrets; - always transported and stored in encrypted form and that the full list of (encrypted) secrets is never allowed outside of the secure data centres. A similar method to that used for CHAP secrets (see section 6.4.3.1) of applying a suffix on the user name provides for non-synchronised change of IPSEC secrets.

6.4.3.3 Out of Band support session

The security controls for Out of Band support session are summarised in this section.

For Incoming PSTN calls to the Branch Router;-

1. PPP CHAP is used to authenticate the incoming session as being from the HNG-X data centre
2. IPSEC used to authenticate the HNG-X Data Centre and provide an encrypted tunnel for Management traffic.
3. User logon is authenticated an audited via Data Centre Radius. There are no local accounts on the Branch Router.
4. The Branch Router Firewall is configured to limit traffic to only that necessary for support – incoming Telnet and incoming HTTP.

6.4.3.4 Router Controls

1. Hard router resets are handled – the Router goes back to a default configuration that is security safe.
2. Configuration of the router firewall to limit the WAN traffic to that explicitly allowed and using state full connections.
3. Serial port is blocked.
4. Remote command sessions: No local user accounts, all users authenticated via Radius and this is audited.
5. Router is configured as switch. Therefore, most of the time, an individual device does not see all Unicast traffic. The exception to this occurs for a few seconds when the router is learning addresses and broadcasts all traffic.
6. Limit range of Counter LAN IP addresses that can source / receive packets, the range is a subnet and takes no account of the number counters configured for the Branch. Therefore adding or removing counters will not require a change to the Branch Router configuration.
7. There is no restriction applied to the MAC addresses that can source packets from the LAN .
8. Protocols and Ports limited to valid ranges
9. TCP and UDP traffic will be limited by using state full connections rules.
10. Anti spoofing, no proxy ARP, no source IP routing
11. Traffic destined for the Router

- Support traffic (ad hoc plus scripted sessions from Router Operational system); All this traffic travels over IPSEC tunnels to and from the Router
- Provisioning traffic; All this traffic travels through IPSEC tunnels to and from the Router
- ICMP Echo request from Network Management station

Note that SNMP traffic is not allowed. The use of SNMPv2 for monitoring* was considered but rejected on the basis that use of Syslog for logging and alerting combined with automated a scripts for fetching information from the Router was sufficient.

* SNMPv2 can be used for monitoring only by not allowing set operations. Specifically the following operations are allowed;

SNMP Inform, Response, and Trap PDUs from Router to NMS Network Management System

SNMP Get, GetNext, and Response PDUs from NMS to Router

DN: Need to update as SNMPv3 is available

DN: Need to specify mechanism to access Router back at base

6.4.3.5 Counter WAN Interface

Since the ex Gateway PC has WAN interface cards fitted (ADSL, ISDN and serial), these will be disabled as follows;

- No drivers associated with ADSL or ISDN cards
- No Modems defined
- Windows RAS service not started

6.4.4 Other Branch Controls

Note that the Branch Hub is unmanaged and therefore no susceptible to security attacks as there is no FTP, Telnet, SSH or HTTP listener.

7 Recovery and resilience

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

7.1 Outline Approach

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

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

Component	Resilience mechanism
Access Tier Firewall	These are configured as an Active / Standby pair with TCP connection state replication between the pair. When the Standby Firewall determines that the Active Firewall has failed then it will (by design) take over with no loss of Application connectivity due to loss of state.
ACE used as layer 4 Load Balancer	At each HNGX Data Centre there are 2 Catalyst 6513 switches each with one ACE blade. The ACE blades are configured in an active standby pair with state replication between them. When the Standby ACE determines that the Active ACE has failed then it will (by design) take over with no loss of Application connectivity due to loss of state. Note that the 6513 are on separate power phases.
Distribution Tier Firewall - FWSM	At each HNGX Data Centre there are 2 Catalyst 6513 switches each with one FWSM blade. The FWSM blades are configured in an active standby pair with state replication between them. When the Standby FWSM determines that the Active FWSM has failed then it will (by design) take over with no loss of Application connectivity due to loss of state.
LAN Connectivity	**Application servers** These are configured with two physical network interface cards that are "teamed" to create a single logical interface. Each such card is connected to a separate Catalyst 6513 switch port. This arrangement means that the failure of any one of {network interface card, Catalyst port, Catalyst switch} does not result in loss of LAN connectivity. Failover takes less than 2 seconds. **Next hop for application server** Application servers do not participate in Routing protocols. The next hop gateway is determined by the having two possible Gateways running VRRP (Virtual Router redundancy protocol). In the event that the Active Gateway fails, the standby will take over this role with no loss of Application connectivity.
SAN Extension	There are at two DWDM provided Fibre Channel services using separate components and over two fibre circuits. There is at least 5m between the fibres at all times. The storage array determines which path to use.
Inter Campus IP traffic	There are two DWDM provided 1 Gigabit Ethernet services using separate components and over two fibre circuits. There is at least 5m between the fibres at all times. A Routing protocol is used to determine a working path.
SSL offload	This is performed within the ACE blade. When a standby module takes over the functionality of the active module, the existing SSL sessions are lost. New SSL sessions are established on the standby (now active) module using the same configuration available on the active module.

Component	Resilience mechanism
Catalyst 6500	At each HNGX Data Centre there are 2 Catalyst 6513 switches in an Active / Active arrangement within the Distribution Tier. The partial or full failure of one of these results in the other Catalyst being used as described elsewhere in this table – for example LAN Connectivity. Each Catalyst functions as an independent layer 3 switch and other layer 3 devices will select the functioning Catalyst based on interior Routing protocols. Note there are also 2 Catalyst 6500 switches in the Access Tier. The same considerations as stated above apply.

7.2 Component Targets

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

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

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

DN: Need to specify that the key configuration items (and which files they belong in) which make the above times possible are listed in the relevant High Level Designs (so need a table enumerating the HLD’s). Also need to state that a “two man check & change approach is applied to these parameters.

7.3 DR

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

    - The Primary site traffic is mainly Production traffic with some Test traffic
    - The Secondary site traffic is mainly test traffic with some Production traffic.
    - Separation of the Production and Test traffic classes is maintained through a variety of mechanisms (refer to section 6.2.3).
    - The resources in shared platforms such as Catalyst 6513 available to Test traffic classes will be limited to avoid the very low risk of Test impacting Production (refer to section 6.2.4).
- Branches use Virtual IP addresses (VIP) to target services. Since the network is Active / Active, each branch will simply select the Data Centre proving the service and detect this based on IP Routing since the VIP is advertised into the Access network. This is how Horizon supports web services today.

Consequences of this approach are that:

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

8 Performance

8.1 Introduction

This section addresses the following two facets of performance management;

- Defining how the architecture addresses the capacity requirements
- Defining how the performance of the components within this architecture can be monitored and measured

8.2 Capacity Requirements

The Architecture uses the following approaches to addressing Capacity Requirements;

1) Creation of a performance model to quantify component targets and predict bandwidth usage throughout the network. This will be carried out using OPNET modelling tools. Fujitsu Services have expertise in this toolset.
2) Ensuring the Architecture is scalable
3) Selection of suitable components – this is especially important in the case of the Post Office branch network due to the large number of endpoints and therefore potential churn of sessions

8.3 Performance and scalability

The following table summarises how performance and scalability is achieved for key areas of the Network Architecture.

Component	Performance and scalability
Branch Access Network	This network is currently in place for Horizon and is supporting throughputs of over to 70 M bits / sec which is more than the Design target for HNGX. The same Network component types and approach will be used to extend this network for HNGX. In general scaling of the branch network is achieved either by replacing components with higher specified ones or by partitioning by branch. Since each HNGX branch, has a deterministic and constant IP address subnet, the network can simply be partitioned such that a subset of branches is supported by a subset of Access network components.
Application Tier Firewalls	The specified component significantly exceeds Design targets. Since the Firewall model is highly specified, scaling would be achieved by partitioning the Branch Access network as stated above and thereby splitting the traffic / session load over multiple resilient pairs of Firewalls.
ACE (Load Balancing and Virtualisation)	The specified component significantly exceeds Design targets. Scaling, if required, would be achieved by use of multiple ACE blades each supporting a separate virtual server farms. The branch application would be configured such that across the estate multiple Virtual IP address targets are used. Refer to comment about blade scaling under Catalyst 6500 entry.

Component	Performance and scalability
Distribution Tier Firewall – FWSM blade	The specified component significantly exceeds Design targets. Scaling, if required, would be achieved by use of multiple FWSM blades (up to 4) in a single Catalyst 6500 chassis. Workload partitioning would be achieved by allocating VLAN's to separate FWSM blades. Refer to comment about blade scaling under Catalyst 6500 entry.
LAN Connectivity	Servers have a resilient connection into the Catalyst 6500 Core switch pair. This connection can be 100 M bit/sec or 1 Gigabit / sec depending on throughput requirements. Blade servers have 8 * 1 Gigabit / sec resilient connections. Refer to comment about blade scaling under Catalyst 6500 entry.
SAN Extension	This solution is in place today for Horizon. Performance is achieved by • Ensuring fibre distance between data centres is < 100km. • Using DWDM service to provide 1,2 or 4 Gigabit /sec Fibre channel connection between SAN switches Scaling achieved by firstly increasing the speed of individual links (up to 4 Gigabit/sec) and use of use of extra FC links between Fibre Channel switches.
Inter Campus IP traffic	This solution is in place today for Horizon. Performance is achieved by • Ensuring fibre distance between data centres is < 100km. • Using DWDM service to provide 1 Gigabit /sec Ethernet connection between Catalyst 6500 switches Scaling achieved by use of extra GE links between Catalyst switches.
SSL offload	The specified component significantly exceeds Design targets. Scaling, if required, would be achieved by use of multiple blades in a single Catalyst 6500 chassis. The multiple blades would be load balanced by the ACE. Refer to comment about blade scaling under Catalyst 6500 entry.
Access Tier and Distribution Tier Multilayer switches-Catalyst 6500	The same top end-switching platform as used in Horizon is specified for HNGX. In the event of insufficient slots for blades, Access switches will be connected to a 1 GE port and provide 100 M bit/sec ports. No over subscription will occur.
Mcafee Intrashield IDS / IPS	The specified component significantly exceeds Design targets. Scaling, if required, would be achieved by use of additional sensors with traffic split across them By the Access Tier switches.

Component	Performance and scalability
RSA/ NI EnVision logging server	10,000events per second license RSA claim “From 500 EPS to 3000,000 EPS and from 320 gigabytes to 3 terabytes per appliance” DN: Need to validate claim

Table 8 - Performance and scalability

8.4 Component Targets and Risk Assessment

This section documents the results of a simple model to predict component targets for the network and asses the risk of these targets not being met by the components.

Derived Network Measures

	Peak		
	Transaction	Reporting	Context Loss (1)
TCP Connections per Second	828	558	583
Max Concurrent Active TCP Connections	34,954	34,954	
Stale connection timeout in ACE (seconds)	300	300	
Maximum Stale TCP Connections	248,484	167,442	
Packets per second from branches	8,283	5,581	
Packets per second to branches	8,283	5,581	
SSL Session establishments per second	7	0	583
Concurrent SSL sessions	34,954	34,954	
Encryption Load (Mbytes / second)	4.3	2.7	

Component targets Blades	ACE	SSL
Concurrent active connections	34,954	
Concurrent active sessions		34,954
Connections per second	1386	
Session establishments per second		583
Packets per second from branches	8,283	
Packets per second to branches	8,283	
Encryption Load (Mbytes / second) (3)	tbs	tbs

Component targets Access network	Firewall	Routers
Concurrent active sessions	34,954	
Connections per second	1386	
Packets per second from branches	8,283	8,283
Packets per second to branches	8,283	8,283
Throughput (Mbytes / second)	4.3	4.3
Throughput (Mbits / second)	35	35
Derived Network Measures - Context loss		

TCP Connection

Scenario is all Counters try to establish a TCP connection within one minute

TCP Connection establishments per Second 583

SSL Sessions

Scenario is all Counters try to logon over within 1 minute

SSL Session establishments per Second 583

Notes

(1) Context loss refers to the case where SSL or TCP sessions need to be re-established. In the case of SSL this will occur when an ACE blade failover occurs. By Design TCP connection should survive hwoever this depends on the interaction of the SSL protocol between the Server (ACE) and the Client (Counter). A possible scenario would be that the Counter doesnot support renegotiation of SSL parameters mid TCP connection and resets the TCP connection.

(2) SSL session establishment will be on PC loading, unless sessions are timed out at a synchronised time. The Counter business apps should ensure that they apply randomisation to SSL establishment following SSL failure. The range should be [2..60] seconds.

(3) This must be less than 2 * Network bandwidth, so < 200 M bits /sec.

8.4.1 Risk Assessment

Access Network

The current Branch Access network in Horizon supports throughput in excess of 70 M bits / sec (Design target for this network under HNGX is 35 M bits /sec).

The main changes in moving to HNGX are;

- Replacing the proprietary transport used in Horizon by Riposte (based on UDP) with TCP. Since Riposte has no congestion control, this move will remove significant complexity and risk of overload caused by congestion from the network.
- Introduction of a Branch Router – this is being done in Horizon as part of the migration to HNG-X.
- Introduction of Wireless WAN (GPRS / UMTS) for backup purposes.

There are many measures within Horizon to avoid network overload caused by traffic storms. These involve;

- Randomisation of connection attempts and back off within a branch
- Scheduling of outgoing traffic from the data centre
- Avoiding use of Backup network for transient network problems (< 1 minutes)

These measures will be carried forward into HNGX.

Conclusion – Existing technologies (ADSL / Dialled) are low risk since solution already in use.

GPRS network is new and subject to but performance risks can be mitigated by use of techniques in place for Horizon. *(DN: Need to specify mechanisms for limiting active sessions)*

ACE Blade

Connections

1 million concurrent TCP connections (Design target is 34,954)

165,000 connection setups per second—Layer 4 (Design target is 828)

Throughput

Total combined throughput of 4 Gbps (client to server and server to client) (Design limit is = 200 Mbps)

1.25 million packets per second (Design target is <17,000)

16,384 Real servers (Design Target < 40)

Conclusions

- Risk of not meeting design target is very low.
- Important to note that these are already in use on Horizon hence there is design and operational experience.

SSL Termination

Connections

3000 new connections per second (Design Target is 583)

60,000 simultaneous connections (Design Target is 34,954)

Throughput

The ACE can support up to 300 Mbytes /sec of bulk encryption (design limit = 200 M bps)

Conclusions

- Risk of not meeting design target is very low.
- No experience within POA of these. This will be mitigated by prototyping in lab environment.

FWSM

The Cisco FWSM provides;

100,000 connections per second (design target is 828)

5 G bps throughput (design target is 35 M bps)

1 million concurrent connections (design target is 34,954)

Conclusions

- Risk of not meeting design target is very low.
- Important to note that these are already in use on Horizon hence there is design and operational experience.

Access Tier Firewall

Clear text throughput: Up to 650 M bits /sec (Design limit =200 M bits / sec – for branch firewall)

Concurrent connections: 400,000 (Design target is 34,954)

20,000 connections per second (design target is 828)

Conclusions

- Risk of not meeting design target is very low.
- No experience within POA of these. This will be mitigated by prototyping in lab environment.

Catalyst 6500 switch

The model specified has a 720 Gbps backplane and can switch 400 Mpps at Layer 2 / Layer 3.

In terms of configuration the 6513 is specified which is a 13 slot chassis.

Slot	Usage
7,8	Slot 7,8 Fabric card – 720 Gbit /sec backplane
9 – 13	Line cards in slots 9, 10, 11, 12, 13, each Line card 40 Gbits /sec onto backplane; Slots 9 through 13 - 48 port 10/100/1000 Total ports = 288
1	ACE
2	FWSM
3	SSL
4,5,6	96 port 10/100 cards, shared 32 G bps bus Total FE only ports = 288

Max available GE Ports 288 (Design target is 100, leaving 188)

FE only Ports 288 + GE left over ports (188) = 476 (Design Target is 380)

Conclusions

- Important to note that these are already in use on Horizon hence there is design and operational experience.
- Need to update model for port requirements since chassis port utilisation is about 80% = (480/(288+288)). May need to use GE ports to create 10 FE ports using L2 access switches. In particular some of the ports will be from the Access Tier switches.

IDS sensor

Each appliance is specified to support throughput of over 1 Gigabit /sec (design target is < 200 M bit / sec) across all traffic streams. This is the total throughput in + out of the HNG-X Data Centres.

Conclusions

- Risk of not meeting design target is low.
- No experience within POA of these devices. However a managed service will be used from Core services through the device lifecycle (design / deploy / maintain) and this team has plenty of experience of these devices. .

8.5 Capacity Monitoring and Measurement

The measures in place for Capacity Monitoring will include;

- Branch network QOS / bandwidth measured from Branch and analysed locally. Exceptions reported via events.
- Within the Capacity Management framework there are reports produced showing trends in defined performance measures, for example free memory on each multilayer switch. The Network Management system collects raw data performance data using SNMP polling. This data is delivered, via SYSMAN to the Capacity Management workstation for production of the reports.
- Service provider capacity reports showing trends plus detailed days in reporting period ; Connect DSL / C&W MPLS / GPRS / FSBN.

Near real time alerting mechanisms in place will include:

- NNM will raise alarms on defined thresholds, for example Firewall CPU utilisation > 80% for last 5 minutes.
- Multilayer switches will raise alerts to the NNM when traffic storms are mitigated or when traffic policing is necessary

8.6 Traffic Management

8.6.1 Traffic shaping

For HNGX branches (so non Riposte), traffic shaping will be applied to SYSMAN traffic during the day to prevent Reference Data delivery from impacting normal branch operations. This covers the case where there is a large backlog of reference data deliveries in progress at 09:00am

Quality of Service Mechanisms (Marking and queuing) will be applied to backup traffic in order to limit its impact on other traffic classes.

9 Migration

This section addresses migration from the existing Horizon solution.

9.1 Branch Network Strategy

The diagram below outlines the proposed branch network strategy for HNG-X. DN: This needs updating to cover change in barcnh Router strategy and outcome of decision on whether Horizon branch services move to IRE1x.

Branch Network Strategy
Version 0.8

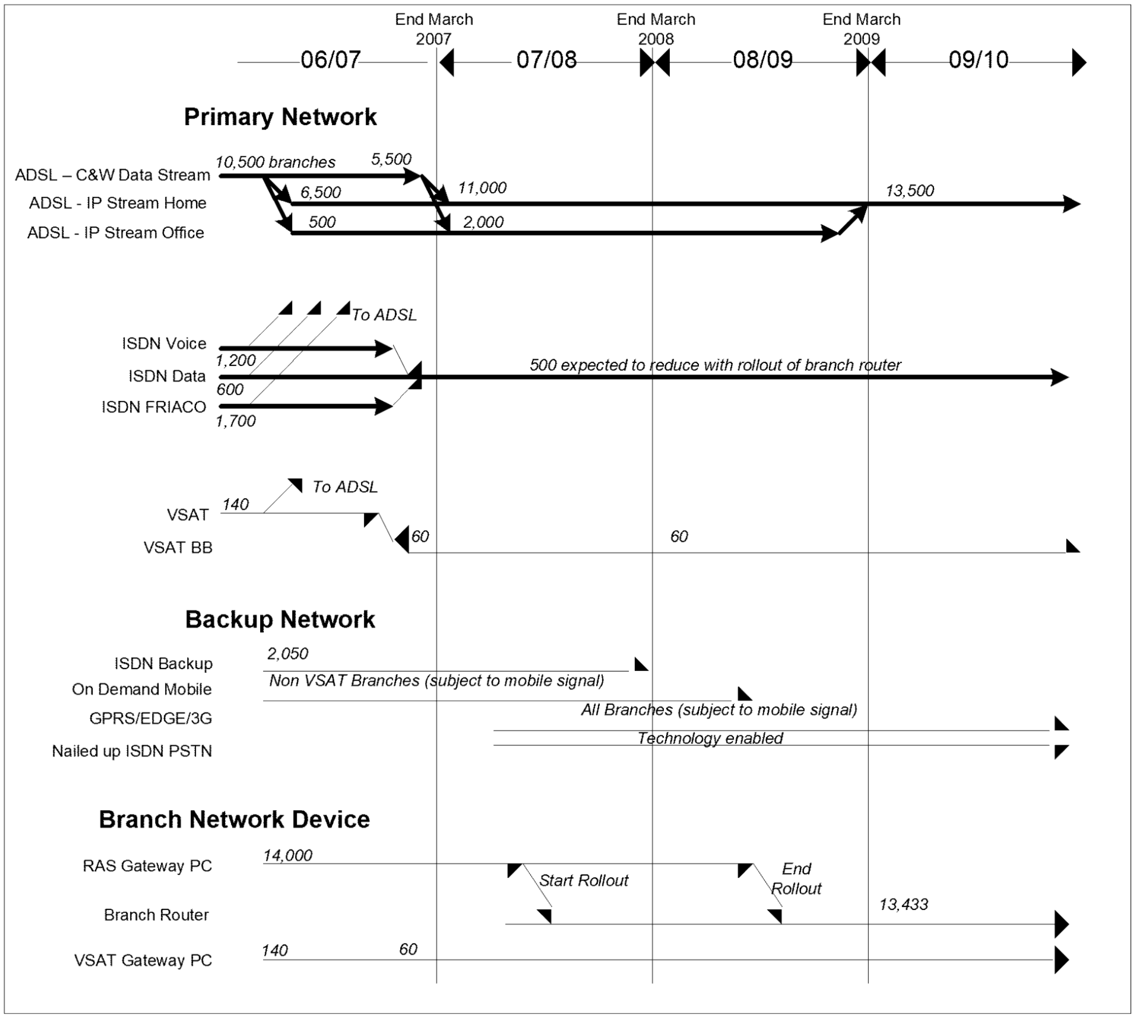

mj_Router_Rollout_v3.vsd
Branch Strategy

Figure 15 – HNG-X branch network strategy

The key characteristics are:

1. Replacement of VSAT with Broadband VSAT service. This removes the bandwidth issues associated with the satellite connection. It is suggested the PSTN line is provided to allow for out of band management if required (although not in current network pricing).
2. Migration of as many branches as possible to ADSL IP Stream – initially a mixture of Office (for larger sites – assume 15%) and Home (for smaller sites – assume 85%). This technology is the most cost effective for an estate the shape of Post Office with large number branches (estimated at >70%) in BT exchanges where Local Loop unbundling is not taking place.
3. Once the new HNG-X application is installed, the use of IP Stream Office will be ceased to reduce costs further. This approach is seen as too large a risk for the current solution due to Riposte's use of UDP and its intolerance to network congestion.
4. ISDN will consolidate to one type (ISDN Data Only – i.e. remove combined current Data and Voice usage of the current Bronze type) by March 2007.
5. By March 2007, the ISDN Voice circuits in the data centre can be consolidated (as they are no longer being used for inbound traffic) to reduce network and router costs.
6. The gateway PC communications function in the branch will be replaced by a router with rollout expected to start in June 07.
7. ISDN backup and on demand mobile (using HSCSD) will be replaced – as the router is rolled out by GPRS/EDGE/ 3G.
8. Other backup approaches (essentially any combination of supported technologies) are also available in the solution, but no such branches have been included in the HNG-X price.

9.2 WAN Network Strategy

The diagram below outlines the proposed WAN network strategy for HNG-X. DN: Need to update in next version.

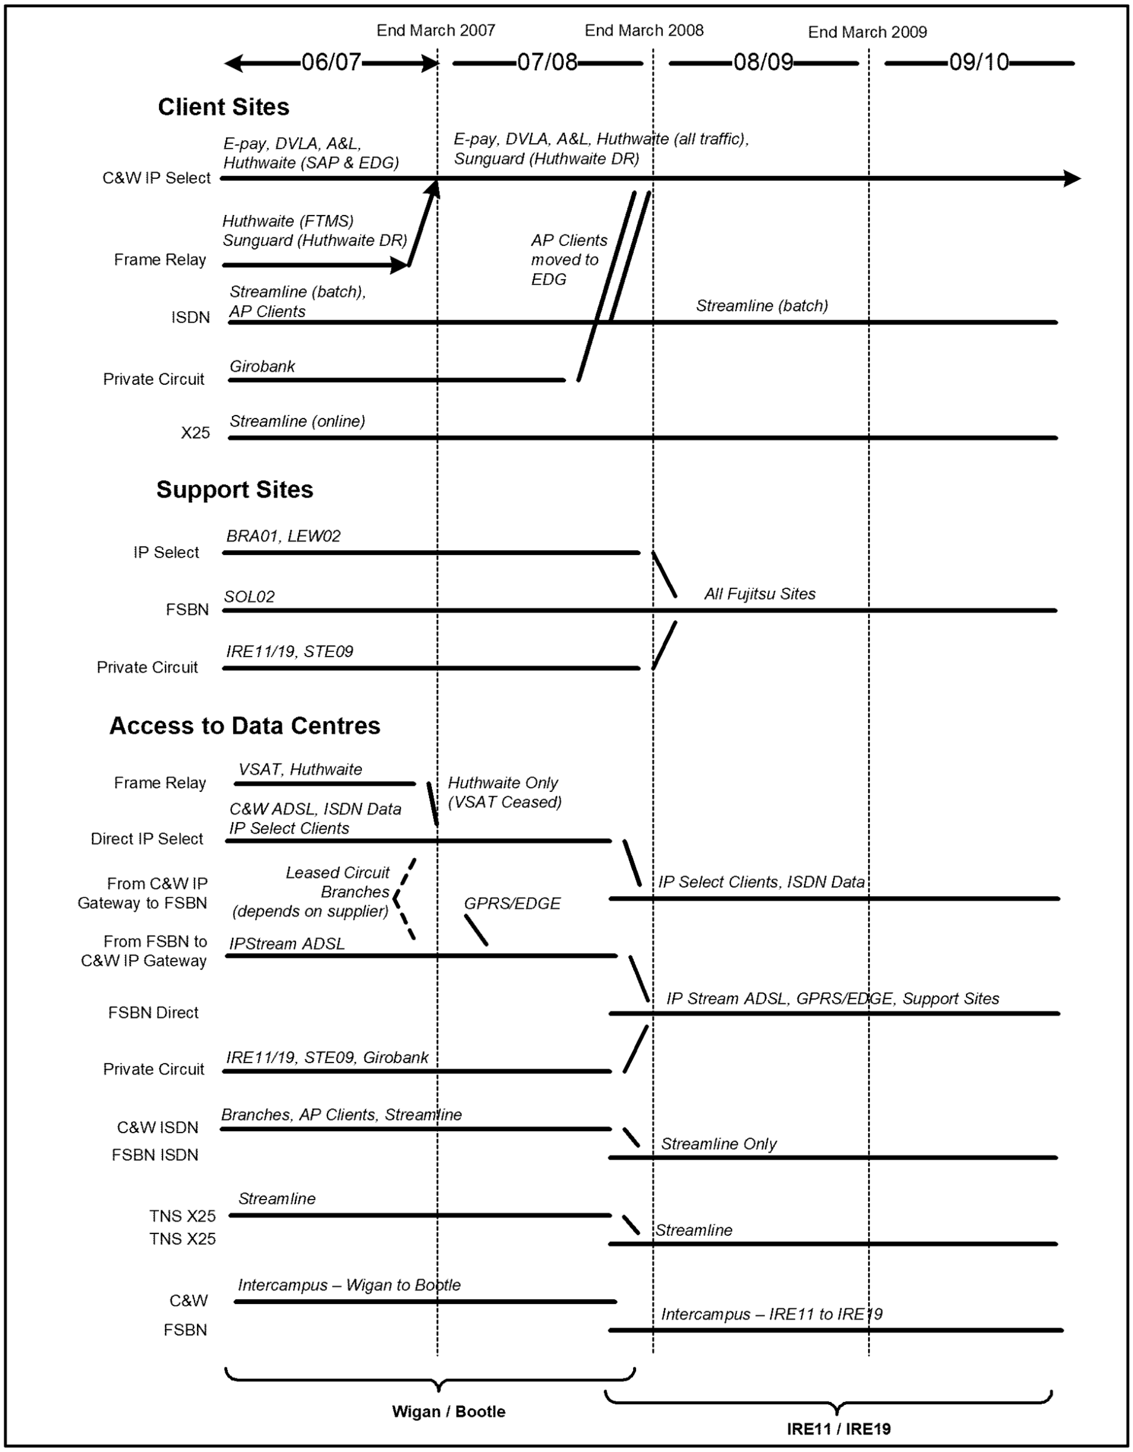

WAN network strategy
NetworkArchitecturev03.vsd
0.4

Figure 16 – WAN network strategy

DN: Need to include the following; CRE02 for OBC plus IDS, Ste09 should be ste04, Network support access – Warrington.

DN also need to take account ELS01 move and Lewes

Most changes are driven by the move of data centres planned to be around March 2008. Most changes before this are to make the migration easier. The impact on test circuits is not shown, but will typically follow the same approach as for live with earlier dates required.

For the client connections the key characteristics are:

1. The C&W IP Select services to client sites (e-pay, DVLA, A&L and Huthwaite continue).
2. The frame relay connections to Huthwaite and Sunguard are ceased once the FTMS traffic is moved to the IP Connect circuit. This is expected by the end of 2006/7.
3. ISDN connected AP Clients and Girobank are moved to EDG connections before the move of data centres.
4. ISDN and X25 connections to Streamline continue.

For the support sites the change is to replace the many different connection methods, by a standard approach though the FSBN.

For access to the data centres, the key characteristics are:

1. Frame relay connections supporting VSAT Branches and Huthwaite will ceased by the end of 2006/7.
2. The IP Select connections that continue in the new data centre will be provided via the C&W IP gateway in SDC01 and Telecity.
3. IPStream ADSL will link to Wigan and Bootle via the C&W IP Gateway. GPRS/EDGE traffic will be added to this as part of the Branch router rollout.
4. In the new data centres (IRE11/19), IPStream ADSL, GPRS/EDGE and Fujitsu Support sites will connect via the FSBN directly.
5. The service providing access for the VSAT Broadband will be carried over the FSBN (using the C&W IP Gateway while in Wigan/Bootle).
6. ISDN PRI in Wigan and Bootle support Branches, AP Clients and ISDN BRI for Streamline. All AP Clients will have moved to EDG before the migration to the new data centres. Therefore the ISDN services being directly supported in the new data centres are {Streamline batch files, out of band support and dial out to isdn dial on demand branches}.
7. X25 for Streamline is unchanged apart from the move to the new data centres.
8. The Intercampus connections at the new data centre will be provided by DWDM services.
9. The bandwidth of the C&W IP Gateway is assumed to be 90Mbits/s until the new data centres are used when it drops to 10Mbit/s. This is due to the branch IP Stream ADSL traffic not having to use the IP Gateway once the new data centre are in operation.

During the migration from Wigan/Bootle to the new data centres in IRE11/19, the data centres need to connect to each other to allow for the transfer of data. This is via the C&W IP Gateway at Bra01 (refer to section 2.5) with the C&W IP Select Service and FSBN providing connectivity into Wigan/Bootle and IRE11/19 respectively.

This migration has a number of phases and there are different bandwidth requirements through the IP Gateway for each phase. This is estimated in the table below. These estimates will need to be validated as part of the detailed design.

#	Phase	Daytime Usage	Overnight Usage
0	Start Position	IP Stream ADSL + GPRS/EDGE: 50Mbit/s FJCore <-> IPGW <-> Bootle/Wigan	IP Stream ADSL + GPRS/EDGE : 20Mbit/s
1	POL_FS Moved (Weekend A)	IP Stream ADSL + GPRS/EDGE: 50Mbit/s Huthwaite IP Select (POL-FS Only): 2Mbit/s Total – 52Mbit/s	IP Stream ADSL + GPRS/EDGE: 20Mbit/s Huthwaite IP Select (POL-FS Only): 2Mbit/s Total – 22Mbit/s
2	+ Batch Services Moved (APS, TPS, DRS etc)	IP Stream ADSL + GPRS/EDGE: 50Mbit/s Huthwaite IP Select (POL-FS + FTMS): 2Mbit/s Total – 52Mbit/s	IP Stream ADSL + GPRS/EDGE: 20Mbit/s Huthwaite IP Select (POL-FS + FTMS): 2Mbit/s Harvesting of data across link (SQL*NET Traffic): 50Mbit/s (more if available) Total – 72Mbit/s
3	+ Online Services Moved (NBS, DCS etc)	IP Stream ADSL + GPRS/EDGE: 50Mbit/s Huthwaite IP Select (POL-FS + FTMS + Track&Trace): 2Mbit/s Branch Online Transactions: 20Mbit/s E-pay, DVLA, A&L IP Select: 2Mbit/s Total – 74Mbit/s	IP Stream ADSL + GPRS/EDGE: 20Mbit/s Huthwaite IP Select (POL-FS + FTMS + Track&Trace): 2Mbit/s Branch Online Transactions: 2Mbit/s Harvesting of data across link: 50Mbit/s (more if available) E-pay, DVLA, A&L IP Select: 2Mbit/s

#	Phase	Daytime Usage	Overnight Usage
			Total – 76Mbit/s
4	+ Branch Services Moved – VPN, Cor Servers, Generic Agents (End state)	ISDN Branch Data: up to 6 Mbit/s Huthwaite IP Select (POL-FS + FTMS + Track&Trace): 2Mbit/s E-pay, DVLA, A&L IP Select: 2Mbit/s Total – 10Mbit/s	ISDN Branch Data: up to 6 Mbit/s Huthwaite IP Select (POL-FS + FTMS + Track&Trace): 2Mbit/s E-pay, DVLA, A&L IP Select: 2Mbit/s Total – 10Mbit/s

9.3 Four Data Centre operation

Key characteristics during this phase are:

a) All 4 Data Centres are active from a network perspective and fall within a single Network Management domain and single IP addressing domain.

b) The IP address as seen from branch will not change because doing so would be complex. In practice this means that the Correspondence servers and VPN servers do not change their interface IP addresses. Also it should be noted that the Riposte traffic carries IP addresses and therefore likely to cause problems with use of NAT.

c) Due to (b) the principle has been adopted that most Horizon IP addresses will not change during the movement of Horizon components between data centres. The exception (remaining 20%) are those which are long life (survive in HNGX) – these get a new IP address according to RFC1918.

d) The further principle is that IP addressing must be unambiguous which in general means that duplicate interface addresses cannot exist concurrently – for example at more than one data centre. Note that Virtual IP addresses can exist concurrently – routing is used to sort out how these are used.

e) There is no provision for subnets existing at the Horizon data centres concurrently with the HNGX data centres. Specifically there is no use of L2TPv3 pseudo wire to enable spanning of subnets between Horizon and HNG-X data centres.

f) Due to (d) it means that any advance provisioning at the new data centre must use “temporary” IP addresses. These would be reassigned to the correct values during the migration epoch.

g) The traffic between Wigan / Bootle and IRE11 / IRE19 during the period of 4 data centre working will consist of a small subset which is delay sensitive with delays being apparent at the branch. For example between Weekends “C” and “D” the Debit card agent is located in Ire1x whilst the Correspondence servers are in Wigan / Bootle. The protocol between them is RPC which delivers’ at most on message per network round rip. The network will treat this traffic as higher priority in the sense that should there be any contention on the CE to PE link then packets in this traffic class will be given priority. The simplest way of achieving this is to dedicate an MPLS VPN and route traffic to this alternatively QOS marking can be used.

9.4 Branch Network migration

This section summarises the main changes in the network during Migration phases with respect of Branch traffic.

- Support for Framed Routes in IRE1X located Radius to enable HNGX branches on NT4 Gateway. This is an interim measure to support the NT4 HNGX gateway without Utimaco VPN.
- Re location of dial out "prod" service in SDC01 / TCY02 to provide gateway of last resort prior. This enables HNGX branch Router and is part of the target solution.
- Provision of LNS servers for C&W dialled service in SDC01 / TCY02 in preparation for closing Wigan and Bootle. This is performed prior to moving Horizon branch services to IRE1x so that no changes are needed during weekend "D". Specifically Routing will be determined which data centre traffic is steered to. The fundamental constraint is that the LNS are selected by dialled number string. If this change is not carried out prior to weekend D then all ISDN branch traffic Horizon and HNGX will go via Wigan and Bootle. This is part of the target solution.

9.5 Branch Migration to HNG-X

9.5.1 Introduction

This section provides an outline of the Migration steps in getting a Branch from its current state to Target state.

The following key dependencies exist in the current solution;

1. A Branch will not be migrated from its current Horizon state until all Horizon Branches connect to endpoints in IRE1x. In practice this means that Branch Migration follows following weekend "D" as described in [HNG-X Migration Strategy]
2. The Branch Router solution is targeted at HNGX Application branches only. This means that a Branch Router can be deployed at or after the epoch when a branch migrates to the HNGX application.
3. The previous point implies that the solution must support a HNGX Application branch working with a NT 4 Gateway PC providing the communications.

It is assumed that the migration of the operating system to XP in a branch follows the introduction of Branch Router.

Note that a HNGX application branch is defined as a Branch that is which is running the HNGX application. An agreed constraint is that all counters in a HHGX application branch must run the HNGX application there is no dual working allowed.

9.5.2 Phase 1 Data centre preparation

For Horizon Gateway PPP names; Framed Routes to Branch Subnet for PPP sessions and VPN branches in preparation for switching off VPN. Note the name is changed to provide for a phased introduction of framed routes.

9.5.3 Phase 2 HNGX application migration

The first Migration step is to deploy the HNGX application. The key changes from a network perspective are described below;

Branch Change	Network implications
Riposte removed / switched off	No more UDP Messaging traffic No more TCP traffic for web services
POLO removed / switched off	None – this has no impact since the CHAP secrets are protected under a scheme which does not rely on the Post masters Token KEK. This statement is included to make the point of no impact explicitly.
Utimaco VPN encryption switched off on Slave Counters. It is also switched of on all interfaces within the Gateway – so LAN and all WAN.	This is dynamic (no reboot required) and causes the VPN filter driver to pass traffic transparently. The default behaviour on reboot will be set to not encrypt. There is no need to unbind the VPN filter driver from the protocol stack. Note that MTU correction on interfaces is not removed. This avoids the need to reboot the platform to bring this into effect. Since the application traffic flows are TCP based, the MSS negotiation during connection setup will settle on the smaller MSS asserted by the Counter and thereby avoiding fragmentation of IP datagram's for this TCP traffic. The results of this behaviour are largely benign; in streaming mode TCP will use approximately 1400 byte IP datagram's instead of 1500 byte ones. The impact on throughput is predicted to be negligible. Also note that to enable peers to communicate within the branch regardless of their encrypt / no encrypt setting, a protocol based on broadcast destination IP addresses should be used.
Apply Security enforcing components	DN: These are still to be defined. But these are to compensate for removal of the Utimaco VPN and leaving the counters with NT4. It is expected to include a combination of a installing a Firewall, security patches and possibly changes to use of shares.
Gateway Counter has IP address .254 added to its interface.	This requires a reboot (to be confirmed). At this stage it cannot be asserted that this change can be applied in advance of the Branch migration as the Utimaco VPN does not support multiple IP addresses on an interface. Further investigation is required to determine whether this can be achieved.
Slave Default Route changed	Slaves Counters are changed to use .254 on the Branch LAN subnet as the default gateway. This can be achieved with no reboot, dynamic change via routing table update with the change made persistent with a registry change to the default route value. This is done so that no change is required to the Counters when the Branch Router is introduced. In the interim period the Gateway counter will continue to be the IP Gateway.
Gateway Counter PPP name	This is changed to allow for a phased introduction of Framed Routes, All other aspects remain the same
Change ISDN idle timers	For those branches which use ISDN dial on demand as the Primary Network type it is necessary to lengthen the idle timer. This is because the current value (5 seconds) is less than the transaction timeout for TCP based transactions. Within Horizon Riposte is used to keep the line up, in

	the case of web services a Riposte priority messages is created specifically for this purpose. The impact of this change is that additional dial ports are likely to be required from C&W.

The Consequence of the above from an endpoint point view is that;

- The HNGX Branch Application Layer servers have the interface IP address available to them in IP communications. For the slaves they will observe the source IP the Ethernet interface address, for the Gateway they will observe the source IP as the WAN IP address. Note that this WAN IP will change to the Gateway Ethernet IP address when the Branch Router is introduced so the BAL needs to allow for two IP addresses for the Gateway.
- SYSMAN2 sees no change to the Branch from an addressing point of view. It is perceived as a significant risk to cause a disruption to the communications for the enterprise management service whilst this same service is causing and orchestrating changes.

The problem that needs to be solved is how to robustly and with no rendezvous between the Branch and the Data centre VPN layers ensure that Data Communications with Branch following Migration (with possible regression) continues to be possible.

In the case of the new HNGX application servers at the data centre (BAL) This problem is trivially solved since traffic originates in the data centre on new platforms. Since this traffic only flows to branches which are HNGX application branches, there is no decision to be made as to whether this traffic should be encrypted by the Utimaco VPN layer

However the case of traffic originating at legacy platforms such as SYSMAN2 and the NNM presents a more difficult problem. This is because the applications on these platforms continue to communicate with the Branch. Prior to the Branch Migration this traffic needs to be encrypted, post Migration it needs to be in clear with the possibility of regression back to encrypting the traffic.

The proposed solution is to use the features in the Utimaco VPN product, namely "Encrypt when You Can Policy", that were specifically introduced to avoid synchronisation of crypto / plain settings during rollout of the VPN product. Specifically this legacy traffic continues to flow through the VPN layer. The VPN layer will make a decision as to whether to encrypt or not as detailed below;

The remainder of this section provides more detail to substantiate the viability of this solution.

Migration event	Summary
Branch enabled for HNGX migration.	A given branch it is enabled for migration in advance of when the Postmasters button to trigger Migration is enabled. Branch Estate Management makes available a list Branches with the following categories; 1. Those which are HNGX and have passed the point of no return. 2. Horizon branches – subcategory (A) those that eligible to migrate in the next n days (tbs but probably 4), subcategory (B) the rest. An update is performed based to the VPN servers based on this list and existing process; All branches in category (1) are set to plain communication so that no attempt is ever made to establish a session key. All branches in category (2A) are set to "Initiate Crypt". The full semantics of this are covered in [SG VPN for PathWay - Phase 3 Extensions]. In summary this causes no change to existing Horizon working. All branches in category (2B) are set to "Use Crypt". This is what happens

	today as normal process. This approach limits the number of branches in the indeterminate state "Initiate Crypt".
SYSMAN 2 commits the HNGX – Software at a Branch	The Branch Migration process will abandon if the WAN is not available. During the process a central decision is made by SYSMAN whether to complete the Migration – "Commit".
Branch Migrates and may Regresses	Some time during the night the branch may switch to HNGX mode in which case communication between the branch and legacy Horizon endpoints in the data centre will fail as the branch is no longer encrypting traffic. This failure will cause one or more VPN servers in the Data centre to log events in their event log "Plain packet received from IP address 223.64.0.1". To enable communications, it is necessary to clear the session cache on the VPN server for the branch. By design, the VPN server has a policy setting for any legitimate branch in this category will have an Initiate Crypt. Clearing the session cache for the branch will enable either plain or encrypted communication. To cause this there will be a VPN watcher process on every Server. (This to a certain extent reuses the approach when VPN was introduced). The basic logic of this process will be; For every "Plain event " on the VPN server; (If the IP peer = branch is in the "Initiate Crypt" state) and (Session not used in 15 minutes for inbound decrypt) and (the session cache for this IP peer has not been cleared for 30 minutes) then Clear Session cache for IP peer = branch Note that regression simply causes the Branch to revert to encrypted communication. The check on the session being used is to prevent a spurious reset. This may not be necessary – to be confirmed during design phase. No check on Branch Migration state is included for simplicity. Superficial analysis suggest this done not weaken the security assertion (see IDS point below) – further work required.
Traffic flows	Plain traffic outbound originating from Legacy platforms to Horizon branches– via VPN layer but does not get encrypted All Plain traffic in bound via IDS, VPN always bypassed.

9.5.4 Phase 3 Introduction of Branch Router

The Physicals aspects of this are in the Section 2.6.3.2.

The BNR solution is retired which implies changes to the Gateway.

A Gateway reboot is required to remove the .254 interface IP address and set a default route pointing at the Branch Router.

At this stage the Branch Router solution is as per Target solution apart from that the Branch LAN subnet is not compliant with RFC 1918.

9.6 Radius Services Migration

Prior to moving the Horizon Branch services to IRE1x (Weekend "D"), the Radius servers will be established in IRE1x. Once they are fully operational which means they are pulling and applying updates from Branch Estate management, the various service providers for Horizon WAN services will migrate to using the new Radius servers. This will be done in advance of Weekend "D" as there is no need for this to be synchronised.

9.7 External Interfaces

In general there are two types of changes to these interfaces;

- The network connectivity needs to include IRE1x
- The number of TCP connections may change

The general approach taken to Migration is that all changes in respect of network connectivity and where necessary layer TCP connection models are made and tested (at least to establishing layer 4 connectivity) in advance of the Migration weekends. This is to avoid the significant risk of problems during the migration weekend.

Some specific cases are covered in the remainder of this section.

NBS

In the case of NBS, the Agents will handle both Horizon and HNG-X transactions so the TCP connectivity with the FIs is unchanged in respect of the number of connections. In advance of network migration, changes will be applied such that Agents can initiate connections (EDS and A&L) from any of the 4 data centres albeit subject to the constraint about connectivity being unchanged. It is likely to be necessary to extend the range of IP addresses that can source connections in the RMGA domain in which case this will be specified in an update to the relevant TIS to obtain agreement.

LINK his has a different connection model as they initiate TCP connections with Horizon / HNGX acting as the TCP server albeit virtualised across two data centres Wigan and Bootle. This virtualisation will be extended to include IRE1x.

Streamline and E-PAY

There are new DCS and ETS Agents that handle HNG-X transactions; these Agents are modelled on the NBS Agents. The existing Horizon DCS and ETS Agents will continue unchanged and will run in parallel with the HNG-X Agents. For DCS the period of parallel running is up to the end of rollout of PCI changes to the Horizon Counters. For ETS the period of parallel running is up to the end of Branch Migration.

For DCS, each of the four Horizon Agents connects to a single TCP endpoint which is mapped to an X.25 SVC. Therefore adding the HNG-X Agents should be a simple extension of this model. The maximum number of SVC's will need to be checked and if necessary increased in agreement with TNS the X.25 network provider.

For ETS there is currently a pair of e-pay TCP ports for each agent/cluster. Additional ports are required for the additional HNG-X agent. Therefore the expected changes are additional TCP connections and extending the range of IP addresses that can source connections in the RMGA domain

For DVLA the number of TCP connections is not constrained as they are made on demand per transaction. The range of IP addresses that can source connections in the RMGA domain will be extended.

9.8 Horizon PCI changes

The Horizon PCI change is being rolled out to the branch estate gradually and will not start until after Horizon has moved to the new Belfast Data Centres.

On day 1 of Belfast the Horizon OSR component within the BAL will not be used.

- The Correspondence Server -> Routing Agent traffic goes directly to the Banking Authorisation Agents (LINK, CAPO + A&L)
- The Horizon style Debit/Credit Card Agents (not shown in your diagrams) accesses the Correspondence Server directly.

After activation of the Horizon PCI change there will be a pilot period where some of the Banking and Debit /Credit Card traffic goes via the BAL. This is expected to start soon (probably within a week) after the data centre move – but may be later if accreditation testing has not completed.

Horizon PCI requests for Banking and Debit / Credit cards will come directly from the Utimaco VPN layer to the BAL - they will not go through the Correspondence Servers and Routing Agents.

There is a security requirement that only traffic emerging from the VPN Servers can access the Horizon OSR component within the BAL. Specifically, it must not be possible to use the Horizon OSR component from HMG-X counters.

Once pilot is complete, the change will roll out gradually across the estate over a period of a couple of months. After this point the Routing Agents will not handle any Banking traffic.

10 Testing and validation

The purpose of this section is to document any architectural issues that affect the viability and cost effectiveness of validating the solution.

Main challenges identified are;

1) Inception of new products requiring familiarity and significant prototyping, these are;
    a) ACE blade
    b) Cisco ASA Firewall including failover characteristics under load
    c) IDS appliance
    d) Branch Router and associated IPSEC head end
    e) New functionality in network – SSL Termination.
    f) GRE tunnel Head end appliances for GPRS service
    g) New Interfaces with the IPstream service to support ADSL SID
2) Complexity of Branch Router choosing between multiple network types. This will cause a high number of test cases.
3) Highly variable quality of service in GPRS networks. May need to use lab facilities from Wireless WAN provider to provide sufficient test cases.
4) Integration testing with Application Layer
5) Resilience of Branch application under ACE blade failure
6) Security Testing – Firewalls, IDS with Application and Systems Management traffic

A facility will be provided to prove DR protocol; To invoke the Bootle to Wigan campus fail-over we manually close the network connections within Bootle down (i.e., isolating Bootle) apart from the SRDF link and essential Support links (e.g. IRE11) [This is covered in CS/PLA/025]. To validate HNG-x Data-centres and contracted fail-over times we need to be able to isolate the primary HNG-x Data-centre almost instantaneously - will such a feature be available?

11 Risk and assumptions

This section captures the major risks and assumptions in this architectural space.

- HTTPS support access is possible to network appliances through SAS. For example in the case of Firewall support need to use https is for management access to firewalls using Advanced Security Device Manager (ASDM). Need to ensure this will function correctly via SAS as it uses a Java applet downloaded from the firewall to the client (in this case SAS server).
- BT have make ADSL SID available (see section 2.7.1). If this is not the case then the provisioning approach will need to be redesigned.
- New (technology works as expected (refer to section 10 where this is enumerated).

(DN: Need to identify other significant risks in the network design in particular around migration complexity).

12 Requirements traceability

This section is not required as requirements traceability is handled external to the document.

A Universal Service Obligation

This Appendix provides the rationale why PSTN cannot be assumed as a service available at every Post Office Branch.

PSTN
It is the case that PSTN is covered under the "universal service obligation", meaning BT (and Kingston) cannot refuse a request for connection. However the threshold at which BT can apply excess charges is £3,400.

Ref [Ofcom] "Review of universal service obligation"
Ref [BT] "BT'S RESPONSE TO OFCOM'S CONSULTATION "REVIEW OF THE UNIVERSAL SERVICE OBLIGATION" 2005

From [Ofcom]
"BT and Kingston are each required to provide access to basic telephone services upon reasonable request and at uniform prices, irrespective of location. Where installation of a new line costs £3,400 or less, BT's makes a standard charge (£74.99 for residential, £116.33 for business). Where installation costs over £3,400, BT requires the user to pay the excess costs (plus its standard charge)."

Also note that the above only applies to the first narrowband line only – HNG-X would be the 2nd;

Statistics;

From [Ofcom]
"BT has advised that, over a recent twelve month period, it provided connections under its USO involving excess charges to 28 customers. In the same period, BT completed a total of two million orders for PSTN lines. This figure covers USO and non-USO situations, including orders for first and additional lines, switched lines, and connections provided at non-served premises, eg site offices." Against the total of two million new lines, the issue of excess charges appears to be relatively minor. However, there is a need for caution about coming to this conclusion. BT has not retained data on requests for connection where an excess charge is involved, only those where the customer agreed to pay the excess. Ofcom does not know how many cases have arisen where costs would exceed £3400 but a customer has decided not to proceed. Nor does Ofcom know the full range of excess charges that have been quoted. BT has now agreed to keep full records over a three month period. This data will be useful but will not reveal the extent of 'suppressed demand' i.e. people who have previously tried to get connection but who did not pay the excess charge.

From [BT]
"During the period August to November last year out of 8 cases where the £3,400 rule was triggered, 50% of customers decided not to proceed. Thus, during that period, 50% of customers required to pay excess construction charges cancelled their orders and could be added to the already existing level of suppressed demand. These cancelled charges ranged from £19,326 to £112,930 and totalled £166,829. This suggests that the level of suppressed demand is relatively high."

Predictions based on the above
A) Of 2 million orders for PSTN, 28 were fulfilled under USO
B) On average 50% of orders were USO is triggered are not proceeded with
C) Therefore assume of 2 Million orders, 56 trigger USO so about 1 in 35,000.

To allow for the fact that PO openings may have a greater geographic dispersal than the above sample, the figure of 1 in 17,500 is used.

With about 200 OBC changes per annum that require a communications service then there is over an 8 year period;
A 8.2% chance that exactly one PSTN order will cause USO to be triggered
A 0.4% chance that exactly two PSTN orders will cause USO is triggered
Higher values are negligible.

B Network Component characteristics

B.1 Cisco Catalyst 6513 multilayer switch

The model specified has a 720 Gbps backplane and can switch 400 Mpps at Layer 2 / Layer 3.

In terms of configuration the 6513 is specified which is a 13 slot chassis.

Slot	Usage
7,8	Slot 7,8 Fabric card – 720 Gbit /sec backplane
9 – 13	Line cards in slots 9, 10, 11, 12, 13, each Line card 40 Gbits /sec onto backplane; Slots 9 through 13 - 48 port 10/100/1000 Total ports = 288
1	ACE
2	FWSM
3	
4,5,6	96 port 10/100 cards, shared 32 G bps bus Total FE only ports = 288

Max available GE Ports 288 (Design target is 100, leaving 188)

FE only Ports 288 + GE left over ports (188) = 476 (Design Target is 380)

Key Points

- Important to note that these are already in use on Horizon hence there is design and operational experience.
- If extra FE ports are required during Dual running then GE ports will be used to link to Access switch.
- Need to evaluate whether ACE module should be used to free up slot. Also has an advantage in terms of virtual context facility as this simplifies administration.

CSM Blade

Connections

1 million concurrent TCP connections (Design target is 34,954)

165,000 connection setups per second—Layer 4 (Design target is 828)

Throughput

Total combined throughput of 4 Gbps (client to server and server to client) (Design target is < 32 Mbps)

1.25 million packets per second (Design target is <17,000)

16,384 Real servers (Design Target < 40)

Key points

- Risk of not meeting design target is very low.
- Important to note that these are already in use on Horizon hence there is design and operational experience.

SSL Blade

Connections

2,500 new connections per second (Design Target is 583)

50,000 simultaneous connections (Design Target is 34,954)

Throughput

The SSL Module can support up to 300 Mbytes /sec of bulk encryption (design target is 4 Mbytes /sec)

Key points

- Risk of not meeting design target is very low.
- No experience within POA of these. This will be mitigated by prototyping in lab environment.
- Multiple blades can be load balanced to increase performance

FWSM

The Cisco FWSM provides;

* 100,000 connections per second (design target is 828)

* 5 G bps throughput (design target is 35 M bps)

* 1million concurrent connections (design target is 34,954)

Key points

- Risk of not meeting design target is very low.
- Important to note that these are already in use on Horizon hence there is design and operational experience.

B.2 Cisco ASA 5500 Adaptive Security Appliances

The model specified is the 5540 and this is used for all Firewall clusters.

Feature	Description
Firewall throughput	Up to 650 Mbps
Concurrent sessions	400,000
Maximum connections/second	20,000
IPSec VPN peers	5000

Feature	Description
Security contexts	Up to 50
Interfaces	4 Gigabit Ethernet ports and 1 Fast Ethernet port
Virtual interfaces (VLANs)	200
Scalability	VPN clustering and load balancing
High availability	Active/Active, Active/Standby