SIN 472

Issue 2.8

June 2017

Suppliers' Information Note

For The BT Network

BT Wholesale Broadband Connect (WBC) Products

Service Description

This document does not form a part of any contract with BT customers or suppliers.

Users of this document should not rely solely on the information in this document, but should carry out their own tests to satisfy themselves that terminal equipment will work with the BT network.

BT reserves the right to amend or replace any or all of the information in this document.

BT shall have no liability in contract, tort or otherwise for any loss or damage, howsoever arising from use of, or reliance upon, the information in this document by any person.

Due to technological limitations, a very small percentage of customer interfaces may not comply with some of the individual characteristics, which may be defined in this document.

Publication of this Suppliers' Information Note does not give or imply any licence to any intellectual property rights belonging to British Telecommunications plc or others. It is your sole responsibility to obtain any licences, permissions or consents which may be necessary if you choose to act on the information supplied in the SIN.

This SIN is available in Portable Document Format (PDF) from: http://www.btplc.com/sinet/

Enquiries relating to this document should be directed to: sinet.helpdesk@bt.com

Registered Office 81 Newgate Street LONDON EC1A 7AJ
Registered in England no.1800000

CONTENTS

1. Introduction

1.1 Definitions

WBC Customer: The Legal Entity who purchases the WBC service from BT and sells or provides it to End Users (EUs). Also referred to as Communications Provider (CP) and Internet Service Provider (ISP).

End User: The person wishing to connect to an IP network via the WBC service.

The document will refer to ADSL as the following:

ADSL	ITU-T Recommendation G.992.1[1]
ADSL2	ITU-T Recommendation G.992.3[2]
ADSL2plus	ITU-T Recommendation G.992.5[3]
ADSL2plus Annex M	ITU-T Recommendation G.992.5[3]
FTTP	ITU-T Recommendation G.984 series

1.2 This SIN

This Suppliers' Information Note (SIN) provides service description information about WBC and its related components, End User Accesses (EUAs), Aggregation Points (APs) and Extensions Paths (EPs) for use by BT and WBC Customers.

This SIN should be read in conjunction with SIN 346[20] detailing BT's G.DMT interface. Documentation specific to the Openreach GEA Product that forms the access part of the WBC over FTTP offering can be found in SIN 477. Current editions of BT's SINs are available at http://www.btplc.com/sinet/

For a brief service description of the TV Connect option, please refer to ANNEX C. For further technical information about TV Connect, please refer to the TV Connect SIN 511 available at http://www.btplc.com/sinet/

Further information is available by reading the Customer Request Forms (CRF) developed for the service - these can be found at www.btwholesale.com/pages/static/Products/Broadband/Wholesale_Broadband_Connect//index.htm

For further information on the commercial aspects of this service, please refer to:

- Your BT National Account Manager
- BT Broadband Helpdesk via e-mail to: broadband.enquiries@bt.com
- BT Broadband Helpdesk on 0800 0283663
- ADSL website http://www.btwholesale.com > Broadband community
- BT Wholesale website http://www.btwholesale.com
- BT Terms website http://www.bt.com/terms
- WBC website www.btwholesale.com/pages/static/Products/Broadband/Wholesale_Broadband_Connect//index.htm

If you have inquiries relating to this document, then please contact: sinet.helpdesk@bt.com

2. Product Handbook

For a description of the WBC service please see the Wholesale Broadband Connect Operational Handbook (www.btwholesale.com/pages/static/Products/Broadband/Wholesale_Broadband_Connect//index.htm)

3. End User Access (EUA)

3.1 General

The End User Access is a 'wires only' interface. The service is provided either via:

- DSL;
- FTTP (In trial) or
- FTTC (Limited availability)

See following sections for details of each option in 3.2 and 3.3.

3.2 DSL

Openreach provide a copper line with an NTE box.

The WBC Customer or end user provides an ADSL filter which plugs into the NTE and an ADSL/ADSL2plus compatible CPE (i.e. an ADSL/ADSL2plus modem or router) which should comply with this SIN, in order to interoperate successfully with the WBC service. This SIN is published to help WBC Customers and suppliers of CPE devices to develop their offerings for these products.

The WBC product will only be available with a line provided with WLR.

The WBC service will only be provided if a BT Public Switched Telephone Network (PSTN) line is available,

End Users can only receive a single Broadband Service over a PSTN line.

If PSTN service is ceased then the WBC service is also ceased.

A WBC ADSL2+ Annex M service has been launched which offers the possibility of greater upstream line rates by sacrificing some downstream line rate. This will require ADSL2+ Annex M compatible CPE for the service to work. As the primary purpose of the service is to allow greater upstream bandwidth the performance tester will allow the testing of upstream throughput for Annex M service lines. In all other respects apart from the upstream and downstream line rates the service will work as standard WBC ADSL2+.

3.2.1 Line Negotiation

WBC exchange equipment will auto sense between ADSL1, ADSL2, and ADSL2plus depending on the End User's CPE. Note that if "ADSL" is selected on the CRF then the line will not train up using ADSL2 or ADSL2plus.

3.2.2 Line rates – service speed up and downstream

The following table gives a comparison of data rates of the products discussed within this SIN:

		Downstream Line Speed	**Upstream Line Speed standard product**	**Upstream Line Speed *uncapped* product**
WBC Max (ADSL1)	DSL line rate (kbit/s)	Adaptive 288[(1)]– 8128 kbit/s	Adaptive 64-448 kbit/s	Adaptive 64 kbit/s – 1280Kbit/s
	ATM payload rate (kbit/s)	Adaptive 260 - 7361 kbit/s	Adaptive 57-405 kbit/s	Adaptive 57 kbit/s – 1159Kbit/s
WBC Max (ADSL2/2plus)	DSL line rate (kbit/s)	Adaptive 288[(1)]– 8128 kbit/s	Adaptive 64-448 kbit/s	Adaptive 64 kbit/s – 1280Kbit/s
	ATM payload rate (kbit/s)	Adaptive 260 - 7361 kbit/s	Adaptive 57-405 kbit/s	Adaptive 57 kbit/s - 1159Kbit/s

WBC ADSL 2+ (ADSL2plus Annex A)	DSL line rate (kbit/s)	Adaptive 288[(1)]– 24384 kbit/s	Adaptive 64-448 kbit/s	Adaptive 64 kbit/s – 1280Kbit/s
	ATM payload rate (kbit/s)	Adaptive 260 -22083 kbit/s	Adaptive 57-405 kbit/s	Adaptive 57 kbit/s – 1159Kbit/s
WBC ADSL2+ Annex M	DSL line rate (kbit/s)	Adaptive 288 – 24384 kbit/s		Adaptive 64 kbit/s – uncapped (Note 2)
	ATM payload rate (kbit/s)	Adaptive 260 – 22083 kbit/s		Adaptive 57 kbit/s - uncapped

Table 1 – WBC Data Rates for DSL services

Note 1 *A line can Rate Adapt down to 160 kbit/s but BT will only support lines able to work stably at 288 kbit/s and above.*

Note 2 *In practice Annex M will give line rates from 64 kbits/s up to approximately 2.5Mbits/s (if the end user is very close to the exchange). In addition, if interleaving is enabled then there is a maximum upstream rate limit imposed within the ITU recommendation (G.992.5) of around 1.8Mbit/s (see section 3.2.3.2 for more details).*

A Fixed Rate product variant is available for the WBC ADSL2+ Annex A product. The fixed rates (DSL line rates) are 576kbit/s, 1152kbit/s and 2272kbit/s downstream. The upstream rate for the 576kbit/s downstream product is rate adaptive from 64kbit/s to 288kbit/s. The upstream rate for the 1152kbit/s and 2272kbit/s downstream products is fixed at 288kbit/s. The Fixed Rate product variant is not available for ADSL2+ Annex M.

The ATM payload rate is defined as the rate available for the transport of IP packets in the payload of ATM cells; note however that each IP packet will incur ATM adaptation layer, PPPoA/oEoA and encapsulation layer overheads. The maximum IP throughput rate that can actually be achieved will hence vary considerably depending upon the IP packet size. No allowance is required for the overhead introduced by the ATM layer, as this has already been taken into account in these figures.

3.2.3 Dynamic Line Management (DLM)

3.2.3.1 General

DLM enables the 'tuning' of a line's performance to ultimately provide an improvement in the level of service delivered to the end user. Historical performance data will be gathered periodically from lines and this will be used to identify those which are performing badly. These lines will then be re-configured automatically (if possible) to give an improvement in their overall ADSL performance and stability. This re-configuration will result in a short break (typically 20 seconds) in end user service.

It is possible that re-configuration of a line could occur daily until a stable configuration is found although it is expected that the majority of poorly performing lines will only require 1 (one) to 2 (two) re-configurations before reaching a stable point.

The DLM process may result in a decrease in line rate, but this will only occur where a line is identified as performing badly at a higher rate. DLM will also use interleaving (see sec. 3.2.3.2 below) to fix problem lines (unless the WBC Customer has opted out of Interleaving for these lines), and this will result in an increase in the delay over the connection, which may affect some delay sensitive services (e.g. gaming, Video, VoIP).

Notification to the customer (via Broadband Customer Reports) will be provided whenever a change is made as part of the DLM process. The DLM process may also be applied manually as part of the standard repair process following an end user fault report.

DLM works within the rate bounds of the chosen WBC EUA option

Note that the Fixed Rate product variants will not utilise DLM.

3.2.3.2 Interleaving

Interleaving is a standard feature of BT's CMSANs, and part of the ADSL standards:

ADSL ITU-T Recommendation G.992.1[1]

ADSL2 ITU-T Recommendation G.992.3[2]

ADSL2plus ITU-T Recommendation G.992.5[3]

Interleaving can be enabled on any individual line to improve its overall (error) performance and stability. Interleaving introduces error correction algorithms, which in some cases can make the difference between a line working well or not at all. In addition, interleaving may provide a significant improvement in the quality of service experienced by the end user particularly for error sensitive services like video.

Interleaving may be applied to individual lines automatically by the DLM function and may be used where a line is identified as performing badly. BT's interleaving implementation will increase the delay over the Broadband connection which can affect some delay sensitive services like gaming.

It is possible to opt-out of allowing BT to apply interleaving automatically as part of the DLM process or on an individual line basis at or after the time of order. BT may however temporarily apply interleaving during the repair process to see if it could improve service.

If a line has not been "opted out" and BT has applied interleaving, a Modify Order facility via the BT ordering systems will allow a WBC Customer to "opt out" the line at any time. This will have the effect of removing the interleaving on that line and prevents BT permanently applying interleaving in the future. WBC Customers should consider carefully before selecting this 'opt-out' option, since interleaving offers the potential of significant improvements to end user service. Without using interleaving to stabilise a line, other methods may be used, which could lead to a reduction of the achievable line rate.

For ADSL2 and ADSL2plus services DLM can control the minimum INP (Impulse Noise Protection) value that can be set. This controls the amount of data corruption that the error correction routines are able to repair and as such can trade data rate for protection. Currently we have a limitation on some of our equipment not supporting an optional set of interleaving parameters which means the use of some INP values can cause the achievable rate to drop very quickly with greater effect on the higher rates.

For Annex M lines the framing limitations imposed within the standard (see below) mean that the maximum achievable upstream rate will be limited to around 1.8Mbit/s when interleaving is enabled. This limitation has no effect for non-Annex M lines.

The effects of the min INP value where only the mandatory set of framing parameters are supported are shown in the following tables within the ITU-T ADSL2 and ADSL2plus standards documents.

For the ADSL2 constraints of only mandatory framing parameter support (Within ITU-T G.992.3 standards document):

- Table K.3a/G.992.3- Maximum downstream net data rate (kbit/s) for various values of delay_min = delay_max and INP_min and
- Table K.3b/G.992.3 - Maximum upstream net data rate (kbit/s) for various values of delay_min = delay_max and INP_min

and

For ADSL2plus constraints of only mandatory framing parameter support (Within ITU-T G.992.5 standards document):

- Table K.3a/G.992.5- Maximum downstream net data rate (kbit/s) for various values of delay_min = delay_max and INP_min and

- Table K.3b/G.992.5- Maximum upstream net data rate (kbit/s) for various values of delay_min = delay_max and INP_min

3.2.3.3 Stability options

DLM will support stability options of Standard, Stable, Super Stable and Custom. These allow the WBC Customer to influence DLM by changing the performance region, in terms of errors and retrains, which DLM will manage the line towards. The effect of the higher stability options is to increase the stability of the line, potentially by sacrificing line speed, increasing latency, and breaching the Fault Threshold Rate (FTR). The Custom option allows a CP to specify the thresholds which DLM will manage the line towards.

Note: Stability options are not available to the Fixed Rate product variants.

For ADSL2+ products, banded configurations are used, which are automatically invoked by DLM as required. These configurations will be controlled by using a target margin, interleaving depth, minimum Impulse Noise Protection (INP) and upper and lower rate limits. This improves the ability of DLM to resolve line issues.
Note: This approach is not used for ADSL1 products.

3.2.4 ADSL CPE

SIN 346[20] describes the interface presented at the end of an ADSL line at an end user's premises. It describes three possible implementation options using either a BT supplied ADSL Line box adapter or CPE filters.

The WBC products will utilise the CPE filter options, which will be 'self-installed' by the WBC Customer or End User.

The physical G.DMT interface is described in SIN 346[20].

In order to correctly interoperate with the services the CPE device(s) should:

- conform to the requirements in SIN 346[20]
- support PPP over AAL5 (RFC 2364[14]) using VC based multiplexing.
- use ATM VPI/VCI pair 0/38 for data transmission/reception
- support PPP (RFC 1661[9])
- support User Authentication (PAP/CHAP) with PPP (RFC 1334[6], RFC 1994[13])
- support IPCP (RFC 1332[5])
- support PPP IPCP Extensions for Name Server Addresses (RFC 1877[12])
- carry out upstream traffic shaping to the ATM rate corresponding to the IP data rate
- not use Bi-directional authentication

The CPE device(s) may (optionally):

- support network renumbering (RFC 1631[8]).
- Support Passive LCP mode
- Support optional minimum INP parameters (G.992.3 Amendment 1 (09/2005) Table K-6, page 12)
- Support optional Reed Solomon code word parameters G.992.3 Amendment 1 (09/2005) Table 7-8, page 7)
- Support Low Power mode as defined in Sections 6.8, 7.12, 9.5, 10.3 & Table 8-4 of ITU-T Recommendation G.992.3

Support of the optional ADSL2/2plus INPmin and interleaving parameter will improve line rate performance over the implementation of the mandatory parameters.

Once connected to the line, the CPE for WBC services will be configured by the BT network according to the requirements of the service.

For information, the configuration parameters shown in the table below will be requested of the End User CPE.

Parameter	WBC Max	WBC ADSL 2+	WBC ADSL2+ Annex M
Operating modes	G.992.1	G.992.1 G.992.3 G.992.5	G.992.1 G.992.3 G.992.5 Annex A G.992.5 Annex M NOTE 2
Downstream user rates	Up to 8128 kbit/s	Up to 24384 kbit/s	Up to 24384 kbit/s
Upstream user rates	Adaptive between 64 kbit/s and 448 kbit/s or uncapped (1280Kbit/s)	Adaptive between 64 kbit/s and 448 kbit/s or uncapped (1280Kbit/s)	Adaptive between 64 kbit/s and approx. 2500Kbit/s
Fast mode (single latency)	NOTE 1	NOTE 1	NOTE 1

Table 2 - End User CPE Parameters

NOTE 1. Interleaving may be enabled as part of the Dynamic Line Management process.

NOTE 2. ATU-R transmit PSD template compliant with upstream mask number 7 (EU-56) as specified in G.992.5 Annex M.

3.2.4.1 USB modem configuration

USB devices can have two different transfer modes, bulk mode, and isochronous mode. Bulk mode uses the available USB bandwidth, and isochronous mode reserves the bandwidth it requires. This means that modems operating in bulk mode can suffer from rate limiting, as the modem is unable to match the USB rate to the DSL rate. This is due to a USB1 only having ~11Mbit/s half-duplex bandwidth for all devices on the channel.

The fix may be to switch the modem from bulk mode into isochronous mode, however not all modems support the isochronous mode. Some modems may be able to be upgraded to resolve the problem; others may have to be replaced if the best throughput is to be achieved.

It may be that by changing these settings on a particular modem could also results in a reduction of speed, or other problems. Therefore it is highly recommended that WBC Customers contact their modem supplier to ensure that any USB modems that are to be use are configured correctly to give the best possible throughput.

3.2.4.2 Router & Modem Drivers/Firmware

It is important that WBC Customers check that they have the latest drivers for their CPE to ensure it works at the optimum level. It has been found that this is a major factor in incorrectly reported faults. It is the WBC Customer's responsibility to ensure the correct firmware and/or driver are supplied to the end user before faults are reported to BT. Therefore it is highly recommended that WBC Customers contact their CPE supplier to ensure that any CPE that they are using the correct drivers and or firmware to give the best possible throughput and to support the applicable WBC service.

3.2.5 ADSL2/2plus Low Power Mode

BT Cool Broadband® gives the capability for any line supporting the ADSL2+ technology type, on a fully rate adaptive profile, to enter a reduced power mode also referred to as power "Level 2" (L2). This is in order to support a greener broadband network by supporting ADSL2/2plus low power mode.

ADSL2/2plus low power mode is defined within Sections 6.8, 7.12, 9.5, 10.3 & Table 8-4 of ITU-T Recommendation G.992.3. Low power mode support introduces the capability for the "always on" DSL service to become an "always available" service i.e. the circuit maintains connectivity to the network at all times, but in a reduced power mode when "idle". If supported by both the MSAN and Customer Premise Equipment (CPE) then a circuit can enter a reduced power state that will reduce the downstream transmitted power and hence reduce the downstream available bit rate. This will achieve a reduction in power consumption at the MSAN. If either end does not support L2 mode, or does not accept a request to enter L2 mode, then the line will continue to operate in normal L0 mode. There is currently no low power mode for VDSL2.

If the bandwidth of the data that is being carried on a circuit supporting the low power mode operation drops below 128kbps for more than a defined period, then the circuit will request a transition to L2 mode. If accepted, the power will be reduced and the line will enter L2 mode. This power reduction will result in a reduced bit rate capacity on the line which will be negotiated by both ends of the circuit during the power trim procedure. This negotiated rate will be somewhere between the 128kbps and the original "level 0" (L0) rate. The circuit will carry data while in the L2 mode up to this negotiated rate and remain in low power mode. As soon as the requested rate is higher than the negotiated L2 rate, then the line will transition back to full power mode and will request its previous L0 mode parameters. These transitions between full to low power (L0-L2) and low to full power (L2-L0) should not lose traffic. While in low power mode, the MSAN and CPE can negotiate a lower than required rate than is required to support the power drop, and this translates to an increase in reported downstream Signal to Noise Ratio compared to that reported while in full power mode.

The settings utilised for controlling the operation of low power are discussed within the Broadband Forum TR-202 and the values utilised within the WBC product are within those suggested in the technical report.

[Note: Enablement of BT Cool Broadband will be dependent on verification of L2 functionality for a range of modems through technical interoperability testing]

3.3 Fibre Access

The Wholesale Broadband Connect product consumes Openreach's fibre access products. For further information please refer to the Fibre Access SIN`s 477, 495 and 498 at http://www.btplc.com/sinet/ These should be read in conjunction with all WBC over FTTP and FTTC documentation at www.btwholesale.com/pages/static/Products/Broadband/Wholesale_Broadband_Connect//index.htm

Upstream	Downstream	Notes
135kbps	**135kbps**	**Symmetric Generic Ethernet Access - Voice Enablement Product (GEA - VEP). BTW propose to incorporate the 135k Openreach VEP in the WBC portfolio as a strictly data product variant.**
500kbps	**500kbps / 2.5Mbps PIR**	**Symmetric GEA Data Product (GEA - DP)**
2Mbps	**10Mbps**	**GEA Data Product (GEA - DP). 2Mbps upstream assured**
2Mbps	**10Mbps CIR / 30Mbps PIR**	**A downstream bursting feature, which can be applied to the Data Product. 2Mbps upstream assured**
2Mbps	**10Mbps CIR / 100Mbps PIR**	**A downstream bursting feature, which can be applied to the Data Product. 2Mbps upstream assured**

Table 3 – WBC Data Rates for FTTP services

Please note the speed associated with FTTC is determined dynamically by the VDSL line rate.

Please note that there is no DLM on FTTP services.

3.3.1 FTTP CPE

The definition of the CPE requirements is as stated in the Openreach documentation in SIN 477, available at http://www.btplc.com/sinet/

3.4 PPP Layer

The WBC Customer's terminating equipment must support PPP conforming to RFC 1661 [9], RFC 1994 [13] and RFC 1877 [12]. Authentication using CHAP (Challenge Handshake Authentication Protocol) will be requested during the Link Establishment phase.

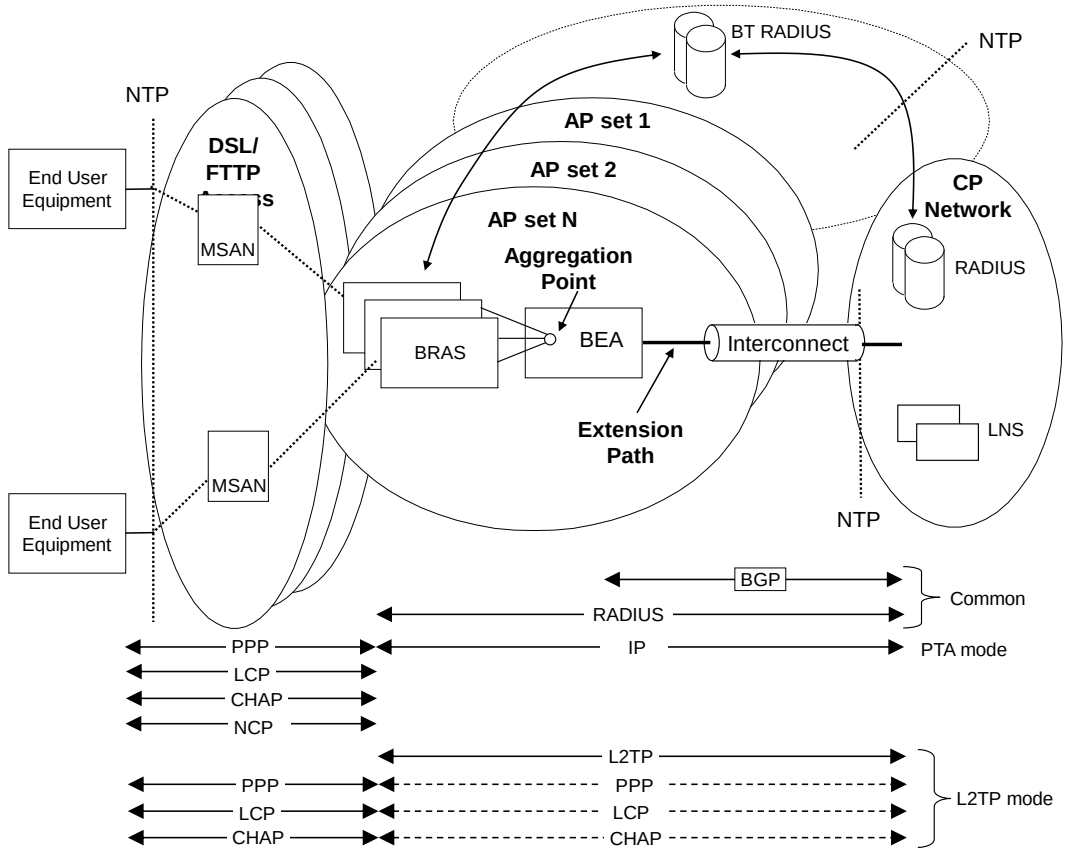

Figure 1. PPP Stages for WBC

Note that only one PPP session per End User is allowed.

3.5 PPP layer Aspects

WBC presents a Protocol and encapsulation auto-sensing PPP layer.

The following combinations of protocol and encapsulations are supported:

- PPPoA VC-mux (DSL delivery only),
- PPPoA LLC/SNAP (DSL delivery only)
- PPPoE LLC/SNAP - single session per line.

The WBC BRAS will automatically configure itself to the protocol presented by the End User's CPE.

Some CPE is itself auto-sensing – in order to ensure that the MSAN and CPE do not constantly chase each other, changing protocols and encapsulations, the MSAN implements a minimum time period before which it will change encapsulation. Following an encapsulation change the MSAN always starts by requesting the creation of a PPPoA session. (DSL only)

The following RFC describe the Protocols and encapsulations used:

- Multi-Protocol encapsulation - RFC 1483[7]
- PPP - RFC 1661[9]

- PPP over AAL5 (ATM) - RFC 2364[14]
- PPP over Ethernet - RFC 2516[15]

It should be noted that the use of LLC/SNAP and PPPoE incur a higher protocol overhead and therefore will provide a lower throughput than the use of PPPoA/VC-mux.

3.5.1 PPPoE Aspects

If the WBC Customer intends their End User's to use PPPoE, they must ensure that the End Users are provided with a PPPoE client compliant to RFC2516[15] and RFC1661[9].

Note. It is BT's experience that most of the popular PPPoE clients are not compliant to this standard.

PPPoE offered as part of the WBC service has no shared broadcast domain with other End Users on the WBC product.

With PPPoE the Maximum Transmission Unit (MTU) is 1492 bytes (i.e. the 1500 byte maximum data size of Ethernet less the PPP header overhead).

WBC Customers allowing the use of PPPoE should be aware that there is potential for some PPPoE traffic to be discarded (so-called "black-holing" - see RFC2923[19] Section 2.1 for further information). In brief: there are servers on the Internet that send out 1500 byte UDP packets with the "don't fragment" bit set. MTU discovery does not work well across the Internet as these servers are often behind firewalls that block ICMP and therefore prevent the MTU discovery process from working. In this case, packets larger than 1492 bytes will not meet the PPPoE MTU requirements and so will be discarded. This is not an issue for UDP packets of 1492 bytes or below. This may not be an issue for TCP as the MSS negotiation should ensure packets are small enough.

3.6 Use of ids

3.6.1 Service Identifier (SID)

The format of the SID is the prefix 'BBEU' followed by an 8 digit string i.e. BBEUnnnnnnnn.

3.6.2 User names

To establish a session to the WBC Customer the End User CPE can support the CHAP protocol which requires a name field to be populated. The name will be passed transparently through to the WBC Customer RADIUS and are not used by the WBC Platform RADIUS server.

3.7 OAM

WBC uses ATM over DSL for the link between the end user and the MSAN but Ethernet between the MSAN and the BRAS. Ethernet does not support OAM. It is therefore not possible to use OAM to ensure the end to end circuit is complete.

3.8 Test User Account

The test user account is a log on ID which allows the test head to check IP connectivity from the EvoTAM to the BRAS. This helps prove that part of the network, which helps us isolate faults.

3.9 EUA interactions via Aggregation Point e.g. Tags Terminations

All EUA interactions with BRAS will be done via hosted web pages e.g. when a WBC Customer has stopped answering End User calls.

These situations are likely to be:

- WBC Customer not responding for MAC
- WBC Customer service has been ceased due to non payment of bill
- WBC Customer has no agreed Aggregation Point (AP) connectivity

4. WBC End User Handover

4.1 Introduction

WBC can be run in two customer handover modes from the AP:

- PPP Termination and Aggregation (PTA), whereby the End User PPP Session is terminated on the WBC BRAS and Routed IP is used to handover the End User traffic to the CP;
- Layer 2 Tunnelling Protocol (L2TP) handover, whereby the End User PPP session is tunnelled to a CP-owned L2TP Network Server (LNS) where the PPP session is terminated.

4.1.1 Flex POSI

BT intends to introduce a new feature to WBC called Flex POSI. This feature is only possible on certain types of BRAS within the WBC network. The Flex POSI WBC option will co-exist for some time with the standard, fixed mapping version of WBC. Where the two options differ this will be highlighted within the SIN. Customers wishing to use the Flex POSI option will also need to continue to support the original WBC interconnect option.

Flex POSI allows the traffic from certain BRASs to route through a WBC interconnect of the customers choosing. Flex POSI and the differences it brings to the product specification is described in further details in Appendix B

4.2 L2TP handover

The L2TP handover solution uses IPv4 source & destination addressing.

It should be noted that if a CP opts to use L2TP handover for its end users' sessions the L2TP control traffic ("keep alives") must be prioritised by the CP's LNS. If this traffic is not prioritised and the link is congested to the point where some of the control traffic is dropped then the CP runs the risk that the L2TP tunnel will be dropped. This will result in the users' sessions on that tunnel being terminated and the users having to start a new PPP session.

4.3 PPPoE and WBC L2TP Pass through (for information)

If a WBC Customer's End Users are using PPPoE then they should be aware of the following behaviour with L2TP Pass through.

PPPoE clients should behave as per RFC2516[15] and RFC1661[9] – if so:-

- PPPoE End User client should send an MRU of 1492 Bytes or less to the BT LAC in the BRAS.
- The BT LAC (in the BRAS) will send an MRU of 1492 Bytes to the PPPoE client.
- The PPPoE client and BT LAC will agree on the lower value MRU.
- The BT LAC will then proxy on this value towards the WBC Customer's LNS.

If the PPPoE client does not obey RFC2516[15] and RFC1661[9] (which occurs regularly), one of the following circumstances will occur:

- If the PPPoE client and BT LAC do not agree a valid MRU then the BT LAC will proxy a value of 1500 Bytes towards the WBC Customer's LNS.
- If the PPPoE client does not send an MRU, but agrees the BT LAC MRU of 1492 Bytes then the BT LAC will proxy that value towards the WBC Customer's LNS for that End User to the WBC Customer.
- If the PPPoE client and BT LAC do not agree a valid MRU then the BT LAC will default to proxying a value of 1500 Bytes towards the WBC Customer's LNS which will advertise this as the IP MTU for that End User to the WBC Customer.
- If the End Users PC is behind the device raising the PPPoE session (for example a routed LAN), then the IP MTU of the PC must be set to 1492 Bytes or lower.

The WBC Customer's LNS should be set up with the correct MTU depending upon the service they wish to offer.

- If they wish to offer just PPPoE then they could set up their LNS with a statically configured IP MTU of 1492 Bytes or lower.
- If they wish to offer just PPPoA then they could set up their LNS with a statically configured IP MTU of 1500 Bytes. However, if an End User chose to use a PPPoE client then packets over 1492 Bytes will be carried over the WBC network but then dropped by the client.

If the WBC Customer wishes to offer a mixed PPPoE and PPPoA service then their LNS should be able to accept the MRU proxied on by the BT LAC and assign that on a per End User basis. However, this will not always be accurate if the PPPoE client does not obey RFC2516**[15]** and RFC1661**[9]** (as detailed in the information above).

4.4 PPP Termination and Aggregation (PTA)

4.4.1 Introduction

All WBC PTA End Users must be assigned an IPv4 address. There are additional options for IPv4 which are described below.

An additional option is being rolled out to assign a globally routable IPv6 prefix to a PTA End User in addition to the IPv4 address. The IPv6 for PTA Mode option is explained in more detail below.

4.4.2 End User IP Layer

In the case of PTA, this layer must conform to RFC 791 [4].

If IPv6 for PTA is used, this layer should also conform to the following RFCs (this list is correct at the time of publication):

- RFC 5072[25]
- RFC 4443[26]
- RFC 4861[27]
- RFC 2460[28]
- RFC 5095[29]

4.4.3 IPv4 address support for PTA

In PTA Mode, IPv4 addresses are dynamically assigned to the End User CPE during the set-up phase of a PPP session.

Three types of IPv4 Address are supported for PTA by WBC:

- Public IPv4 Address;
- Private IPv4 Address;
- Shared Address Space IPv4 address.
  (IANA have introduced a Shared Address Space IPv4 address range for use in "Carrier Grade NAT" functionality, which is a way to mitigate against IPv4 address exhaustion. Shared Address Space is defined in RFC 6598 [24]).

Notes:

1. In order to ensure uniqueness of PTA End User IPv4 addresses across the BT Wholesale platform, private IPv4 address and Shared Address Space ranges will be allocated by BT Wholesale in response to CP requests. This applies to both BT allocated dynamic pools of private or Shared Address Space IPv4 addresses, and CP allocated private or Shared Address Space IPv4 addresses.
2. Addresses in the range 172.16.0.0 – 172.31.255.255 are used by BT Infrastructure and will not be allocated to End Users to avoid clashes.

The WBC Customer (i.e. the owner of the WBC Customer RADIUS) selects who provides the IPv4 address during RADIUS Authentication. The options are:

- BT Wholesale allocates the address from a dynamic pool supplied to BT Wholesale by the WBC Customer;

- The CP returns the address in the RADIUS authentication response. This may be used to allocate Static/Consistently Served IP Addresses or No NAT subnets to an end user. Please see Section 5.3 for route limits.

If the CP wishes BT Wholesale to allocate the address, the IPv4 address would be dynamically selected from a pool of addresses supplied to BT Wholesale by the WBC Customer.
Notes:

1. The CP may provide a single default pool, a number of named pools or both. Pools must be applied across all WBC nodes where the CP has an Aggregation point of presence. It is possible to have both a default pool and named pools. If named pools are used the name must be unique across the WBC service. The pool name is limited to 24 characters and BT recommends all pools used by a particular CP should start with the same prefix e.g.: CPNamePool1, CPNamePool2. At present it is not possible to use the same Pool Name on WBC and IPstream Connect. This will be addressed at a later date.
2. BT is currently exploring the options for reporting IP address utilisation. Reports will be available for utilisation per pool per AP.
3. There is an IETF standard RADIUS attribute (88 Framed-Pool) that should be used to select a named pool in an Access-Accept message. For compatibility reasons a Redback Vendor Specific Attribute (36 Framed-Named-Pool) which achieves the same purpose must also be supplied. BT plan to remove the need for the VSA at a later date but named pools will initially be made available by using both the Redback VSA and IETF attribute. By requiring both attributes initially BT are providing a low impact migration path to the eventual VSA removal.
4. As groups of addresses are added to particular AP's from the pool, WBC will inject routes into the BGP routing to the WBC Customer to inform them which AP the group has been added to.

If the WBC Customer allocates the address:

- The WBC Customer returns the IPv4 address to be used in the WBC Customer RADIUS authentication response.
- The WBC Customer can use any valid public IPv4 address that has been advised to BT Wholesale.
- The WBC Customer can use private IPv4 addresses or Shared Address Space IPv4 addresses from ranges that have been agreed with BT Wholesale.

4.4.4 IPv6 address support for PTA End Users

WBC is introducing support for IPv6 addressing for PTA Mode.

IPv4 is the standard offering and is used for all WBC services. IPv6 for PTA customers will be overlaid on the existing IPv4 infrastructure as an optional feature, creating a dual stacked infrastructure. If requested, the CP's service including their existing AP, and all EP VLANs connect to the AP, will be enabled for IPv6. It is recommended that a CP upgrading to IPv6 does so on all the APs they own (there will be a migration period during which a subset of the APs will be enabled).

An End User Access service on IPv6 will have 2 address types:

- A dynamic WAN address using Link Local. This will be automatically generated by the CPE and does not need to be provided by the CP. This address is not globally routable and not advertised out to the CP;
- A globally routable IPv6 delegated prefix for use on the CPE LAN. The IPv6 addressing for the IPv6 delegated prefixes will be provided by the CP. The CPE will obtain an IPv6 Delegated Prefix from the BRAS by running DHCPv6. The delegated prefixes will be advertised out to the CP.

Like IPv4, there are two ways of allocating an IPv6 delegated prefix to an end user:

- Allocation from a dynamic pool of addresses provided by the CP to BT and installed on the WBC network;
- Allocation of specific addresses from an IPv6 address range by the CP. This may be used to allocate Static/Consistently Served IP Addresses to an end user. Please see Section 5.3 for route limits.

The WBC Customer selects who provides the IPv6 address during their End User RADIUS authentication response. If they wish BT Wholesale to allocate the address dynamically from the IPv6 Delegated Prefix Pool, this will be advised in the RADIUS authentication response. If the WBC Customer wishes to allocate the address, the Customer RADIUS returns the IPv6 Delegated address prefix to be used in the WBC Customer RADIUS response (the WBC Customer can use any valid globally routable IPv6 address prefix from a range that BT has been advised of).

It is expected that the CP will provide up-front all the IPv6 addressing that is ever likely to be needed for the End Users on their WBC service – this includes addressing for an IPv6 Delegated Prefix Pool and any addressing required for CP RADIUS allocated addresses. For the IPv6 Delegated Prefix Pool, BT will require a specific address block size based on the CP's current and forecast volumes. BT will allocate all the addressing needed from the CP pool to cover the current network. As the network grows, additional ranges will be allocated from the pool as required.

A CP will need to choose a fixed size of Delegated Prefix when they request IPv6 for PTA Mode on their service. The chosen delegated prefix size will be used for all the CP's IPv6 End Users, including those with addresses allocated via a pool, and those with addresses allocated via the CP RADIUS. Valid delegated prefix sizes are from /64 to /48.

Notes:

- End Users with IPv6 addressing must also be provided with an IPv4 address;
- A single default IPv6 Delegated Prefix pool is supported at this time;
- IPv6 End Users will be allocated one IPv6 delegated prefix during Authentication;
- Initially IPv6 for PTA Mode will not be supported on every BRAS and will therefore not be available for every End User.
- For Dynamic allocation of addresses from a pool, if the CP has signed up for IPv6 for PTA Mode, and their EU's CPE asks for an IPv6 address, an address will be provided – the presence or not of an IPv6 Address Pool in the RADIUS reply will not influence whether the EU is given an IPv6 address.

4.4.5 Session Terminations

The End User session to the WBC Customer can be terminated by the End User. It will terminate when the End User establishes a new session. The session is also terminated if the End User CPE is switched off. In exceptional circumstances, such as abusive usage by an End User, the WBC Customer may request BT to remotely disconnect one of their End Users' sessions forcibly.

4.4.6 Customer Responsibilities

It is the WBC Customer's responsibility for PTA:

- To provide Public IPv4 addresses for their PTA End Users;
- To request Private IPv4 address ranges for PTA from BT Wholesale;
- To request Shared Address Space ranges for PTA from BT Wholesale;
- To provide globally routable IPv6 Addresses for their PTA End Users if required;
- To install, configure and administer their RADIUS servers.

5. Aggregation Point (AP) & Extension Paths (EPs)

There are 20 WBC PoSI Nodes. When a CP orders the WBC service at a PoSI Node, they will order two elements (in addition to the individual End User accesses):

- An Aggregation Point (AP). The AP is a logical CP-specific virtual router, in which the CP's traffic is managed.
- One or more Extension Paths (EPs). EPs are handover links from the WBC AP to the CP's infrastructure. Each EP is made up of a physical link, a VLAN, and eBGP peering between the WBC AP and the CP.

Alongside the introduction of the Flex POSI feature, BT will be making available a new type of Extension Path called Direct Extension Path. This will provide connectivity directly between the WBC AP and the CP's infrastructure and may consist of a single physical link or a number of physical links with the option of using LAG across the group of links. Where LAG is used, a single VLAN will be associated with the LAG bundle. Availability of the direct extension path is subject to certain network infrastructure activities having taken place, please check with your account manager before ordering.

5.1 Infrastructure IP Addressing

5.1.1 Introduction

This section explains how to define addressing for WBC.

For IPv4, WBC supports two methods of addressing for WBC APs & EPs:

1) Default private addressing allocated by WBC;
2) Customer-allocated private or public IP addressing.

Note: If the customer allocates their own private addressing:

- Prefixes from the range 172.16.0.0 – 172.31.255.255 must not be allocated and must not be advertised into WBC, to avoid clashes with BT infrastructure addressing.
- If the customer is also using private IP addressing for PTA End Users, the PTA IP addresses must not clash with the infrastructure IP addresses used by the CP for the WBC APs & EPs. Note: Shared Address Space IP addressing should not be used for WBC AP & EP infrastructure addressing.

For customers who wish to use IPv6 for PTA mode, IPv6 infrastructure addressing will be required and will be allocated by BT.

Please see below for more details.

5.1.2 BRAS to Broadband Edge Aggregator (BEA) IP Addressing

The customer can specify alternative addressing. Per aggregation point (AP), the customer needs to specify five IPv4 subnets. Two /28 subnets are required for each BRAS logical POP. One /28 is required to address the 4 x MSE AP VLANs per CP between the BEA and the MSEs associated with the AP.

Subnets must be unique across the WBC Customer's network. Ideally the 5 /28 subnets should be contiguous. Note that the WBC Customer cannot change the IPv4 address mask.

For each AP, addressing to populate the IPv4 address tables shown below is required from the WBC Customer to override the default addressing.

AP	Logical POP	Network Address/Mask	BRAS Starting Address	IPv4 address for BEA1 VLANs
Birmingham	1st	10.16.2.0/28	10.16.2.1	10.16.2.14
Birmingham	2nd	10.16.3.0/28	10.16.3.1	10.16.3.14

Table 4 - Sample BRAS to BEA1 VLAN Interface Addressing Override

AP	Logical POP	Network Address/Mask	BRAS Starting Address	IPv4 address for BEA2 VLANs
Birmingham	1st	10.17.2.0/28	10.17.2.1	10.17.2.14
Birmingham	2nd	10.17.3.0/28	10.17.3.1	10.17.3.14

Table 5 - Sample BRAS to BEA2 VLAN Interface Addressing Override

AP	1st Downstream VLAN subnet	1st Upstream VLAN subnet	2nd Downstream VLAN subnet	2nd Upstream VLAN subnet
Birmingham	10.18.7.0	10.18.7.4	10.18.7.8	10.18.7.12

Table 6 - Sample MSE VLAN Interface Addressing Override

The IPv4 BRAS to BEA VLAN addressing will be advertised out to the CP, the MSE VLAN addressing will not be advertised out.

5.1.3 BEA to WBC Customer IP Addressing

WBC-allocated default private IPv4 addressing is supported for up to 8 EPs associated with an AP. If the customer wishes to have more than 8 EPs associated with an AP, they will be required to allocate their own IPv4 addressing.

For the WBC Customer to override the default IPv4 addressing of the EPs the WBC Customer would be required to fill in a table with appropriate values. The tables shown below illustrate the sort of information which would be required from the WBC Customer for each AP and EP.

AP	EP1		EP2	
	AP End	WBC Customer End	AP End	WBC Customer End
1	10.16.10.1/30	10.16.10.2/30	10.16.10.5/30	10.16.10.6/30
2	10.17.1.1/30	10.17.1.2/30	10.17.1.5/30	10.17.1.6/30

Table 7 - Sample BEA to WBC Customer VLAN Interface Addressing for EP1 and EP2

AP	EP3		EP4	
	AP End	WBC Customer End	AP End	WBC Customer End
1	10.16.10.9/30	10.16.10.10/30	10.16.10.13/30	10.16.10.14/30
2	10.17.10.9/30	10.17.10.10/30	10.17.10.13/30	10.17.10.14/30

Table 8 - Sample BEA to WBC Customer VLAN Interface Addressing for EP3 and EP4

If a customer requires PTA mode with IPv6 enabled, an IPv6 /127 is required for each Extension Path and this will be allocated by BT.

5.1.4 B-RAS Loopback addressing

The customer can specify alternative IPv4 addressing. Per aggregation point (AP) the customer needs to specify a contiguous block of unique (unique with respect to the WBC Customer) IPv4 addresses. The number of loopbacks required depends on the number of BRAS devices that will be covered by the aggregation point, an address is required for every BRAS. Note that the WBC Customer cannot change the IPv4 address mask.

For each AP, an IPv4 address table as shown below is required from the WBC Customer to override the default addressing for the BRASes at the AP. The WBC Customer needs to provide BT with a contiguous range of IPv4 addresses which can be summarised with a /29 mask. For example.

Logical POP	IP range /29
1	10.1.1.0
2	10.1.1.8

Table 9 - Sample BRAS Loopback Addressing Override

With the above settings the LAC loopback for each BRAS in each Logical POP is as follows:

Logical POP	BRAS1 Loopback	BRAS2 Loopback	BRAS3 Loopback	BRAS4 Loopback	BRAS5 Loopback	BRAS6 Loopback	BRAS7 Loopback	BRAS8 Loopback
1	10.1.1.0/32	10.1.1.1/32	10.1.1.2/32	10.1.1.3/32	10.1.1.4/32	10.1.1.5/32	10.1.1.6/32	10.1.1.7/32
2	10.1.1.8/32	10.1.1.9/32	10.1.1.10/32	10.1.1.11/32	10.1.1.12/32	10.1.1.13/32	10.1.1.14/32	10.1.1.15/32

Table 10 - BRAS Loopback Addressing

In addition, a /25 subnet per AP is required to provide loopback addresses for each MSE associated with the AP. The addresses will be allocated as /32s to the MSE. The table below provides an illustration using an example subnet of 10.1.2.0 /25.

MSE	Loopback Address
MSE1 Loopback	10.1.2.0 /32
. . .	. . .
MSE128 Loopback	10.1.2.127 /32

Table 11 - MSE Loopback Addressing

Ideally the two /29s and the /25 subnets should be contiguous.

Each BRAS and MSE will advertise an IPv4 Tunnel End Point so the number of tunnel end points advertised to the CP will increase as BRASes and MSEs are added to the AP. BT will advertise a /32 per BRAS/MSE as a tunnel end point, so the number of prefixes advertised will also increase.

5.2 Address Summarisation

For WBC Customers with only PTA end users using dynamic addresses, the BRAS's VR should only pass infrastructure and PTA pool routes and these routes should be summarized at the BRAS's VR where possible.

For the PTA service if a WBC Customer wishes to use consistently served IPv4 addressing for their end users, or no NAT addressing, then summarization for these end users will not be possible and individual /32 host routes and no NAT subnets will need to be learnt from the BRASes. PTA dynamic end user routes should still be received from the BRAS devices as summary routes.

Similarly, if the CP chooses the IPv6 for PTA Mode option, if a WBC Customer wishes to use consistently served IPv6 addressing for their end users, then summarization for these IPv6 delegated prefixes will not be possible and individual delegated prefixes will need to be learnt from the BRASes. Dynamic end user routes should be received from the BRAS devices as summary routes as for IPv4.

For WBC Customers with only L2TP end users the BRAS's VR should only pass infrastructure routes and these routes should be summarized at the BRAS's VR where possible.

Note: It is possible for the WBC Customer to have both PTA and L2TP end users on the BRAS's VR at the same time.

5.3 AP & EP Route Limits

There are two places that the routes can either be limited or monitored to protect the BEA from reaching the maximum amount of routes that can be supported.

i. The eBGP peerings to the CP (one v4 peering for each EP, optionally one v6 peering for each EP) has an inbound prefix filter that limits the number of routes received. This is currently set to a 1000 routes and alarm notification is set to 80%. If 1000 routes is exceeded the BGP peering will be brought down. This action is required to protect other users of the WBC service from configuration errors within the CP's network flooding the WBC service with routes and thereby disabling it.

ii. The number of routes within a VRF will also be monitored. Once the VRF route limit is reached, an alarm notification will be raised. Initially the VRF route limit will be set to 50,000 in total which includes both IPv4 and any IPv6 routes. Once this threshold is breached, we will review the total number of routes configured across all CPs and may request that the number of routes for that CP to be reduced.

Routes counting toward this 50,000 route limit include:

1. Address Pools - one route for each PTA address block;
2. Address Supernets - one route for each Summary Route;
3. Routes pushed by the CP's RADIUS - one for each static end user and one for each LAN route;
4. Inbound Routes from the CP in i) above.

5.4 Asymmetric Routing

If asymmetric routing is perceived to be a problem, it is the responsibility of the WBC Customer to address this.

5.5 Policy Based Routing for PTA Customers

WBC Customers who need EUs to be configured in PTA mode can supply range(s) of addresses against which the EU traffic can be 'Policy Based Routed'. When supplied this feature ensures that the WBC customer is sent all traffic from their EUs and that traffic may not route between EUs in the AP or BRAS. Only traffic going from a PTA customer of the CP to another PTA customer of the same CP is policy based routed.

Addresses can either be supplied to BT for allocation from a pool assigned to each BRAS or provided at connection time by returning the IP address in a radius attribute. The address range(s) must include addresses allocated from dynamic address pools plus any static addresses or LAN routes. If an EU is configured with an address that is outside of the Policy Based Routing (PBR) address range, then that EU will not be able to send or receive data.

The solution is as follows:

- The CP may provide an IP address in the CP's network for use as a PBR next hop. The CP may supply a unique address or the same address for each POSI. The CP may use any address with the exception of the 172.16.0.0 – 172.31.255.255 range;
- The PBR address will be learnt at the BEA via the BGP peering with the CP, and will be subject to the same load balancing as all other non-PBR traffic destined to the CP. It will be possible for the CP to use the MED attribute when advertising the PBR prefix to steer traffic between the CP's PTA End Users over particular EPs;
- If a PBR next hop is not provided, a default BT address of 172.31.255.227 will be used and PBR traffic will follow the Default Route advertised by the CP;
- If all the EPs are down, or a CP withdraws the PBR next hop and the default route, then traffic between the CP's PTA End Users will be dropped at the BEA.

The change of functionality is being rolled out node by node so there will be a period where both methods are in use.

Note: The above functionality is applicable for IPv4. For IPv6, there will be a partial solution for day one, which is that Policy Based Routing will be used on the BRASes to prevent local traffic turnaround. IPv6 traffic between a CP's End Users that are on different BRASes at the same WBC node will not be Policy Base Routed.

5.6 AP Policing

Traffic is policed and metered at the Aggregation Point for billing purposes:-

- Total Traffic is measured and passed through (no policing is applied).

- The rate at which real time traffic is transmitted downstream (between the Communications Provider and the End Users) is policed - please see below for details of the Real Time policing rate. Traffic in excess of the policing rate (described below) will be discarded.

Measurements and policing apply to the aggregate of traffic arriving via the Extension Paths. The exception to this is policing of Real Time traffic which will be done on a per EP basis. The level at which traffic will be restricted is set to 115% of what is ordered to allow a degree of bursting. Above this limit Real Time packets will be dropped. For customers with multiple EPs, they will be able to send more than this across multiple links however excess charging will apply.

Reporting of AP bandwidth utilisation will be through the BBCR reports.

5.7 Best Efforts traffic classes

It is possible for CPs to mark Best Efforts traffic as:

- Best Effort (Priority)
- Best Efforts (Normal)
- Best Effort (Sub)

Under congestion BT may need to discard some of the Best Efforts traffic going to a particular user. The Sub/Priority mechanism allows the CP to indicate the relative priority of packets. This feature does not affect the amount of bandwidth allocated to a user under congestion. However it does affect the weighting of the traffic streams. This means for example that under congestion the Sub Best Effort traffic will be given the lowest priority and the Priority Best Effort traffic will be given the highest priority of traffic in the Best Efforts queue.

5.8 Real Time QoS service

Real Time QoS is supported on WBC. There are two elements to this service:

i) The CP orders a contracted amount of Real Time bandwidth to be applied at the AP;
ii) The CP orders an amount of Real Time bandwidth for each End User line that requires the service.

A CP's Real Time traffic, providing it is within the EP policing limits described in 5.6, will be passed through to the BRAS. Any excess above the policing bandwidth will be dropped at the BEA. For End User Access with Real Time QoS enabled the traffic will be prioritised with the highest priority on egress from the BRAS. Any excess traffic over the End User Access Real Time QoS configured bandwidth will be dropped. The Real Time bandwidth configured for the End User will be the smaller of:

- Contracted Real Time bandwidth;
- CVLAN rate – 100 Kbps.

This is to ensure a certain amount of BE traffic and more importantly, control traffic, can get through.

If the End User Access has not been enabled for Real Time QoS then the traffic will be treated as Best Efforts.

There are a range of Real Time QoS End User Access downstream bandwidth rate options, from 220kbit/s to 4.9Mbit/s.

Real Time QoS is available for End Users on the Fixed Rate product variants. Only the 220kbit/s rate is supported for Fixed Rate product variants.

Note: In the case of L2TP traffic, decisions on handling Real Time QoS will be based on the outer (L2TP carrying) DSCP marking.

5.9 Packet markings

Marking the traffic for each class is performed using the DSCP (Differentiated Services Code Point) field in the IP packet header. The BT network can then recognise the packets and treat them appropriately (this is known as DSCP classification).

Class	DSCP codepoint	Binary value	Decimal value
Real-Time	EF	101 110	46
Best Effort (Priority)	AF21	010 010	18
Premium Forwarding	AF31	011 010	26
Best Effort (Normal)	DE	000 000	0
Best Effort (Sub)	EXP/LU	000 011	3

Table 12 – DSCP code points

Where the PPP session is handled as L2TP, the WBC Customer will have to mark the packet header of the IP packets carrying the L2TP towards BTW's network.

By only classifying the traffic at the DSCP level, BT is able to disassociate itself from the application layer (IP address, Port number, etc.) and leaves that to the WBC Customer layer. When the WBC Customer changes its application layer protocols, functionality etc. it does not then require changes from BT. It is then also up to the WBC Customer layer to control the interaction between applications on a single EUA.

Packets with DSCP markings other than described in the table above will be treated as Normal Best Efforts.

The DSCP marking AF31 was introduced for Content Connect. Any WBC traffic marked with AF31 may be re-marked by WBC as Best Efforts. In the case of L2TP traffic, the re-marking will be performed on the outer (L2TP carrying) DSCP marking.

If IPv6 for PTA Mode is chosen, it will be handled in the same way as IPv4. There will be no change in the classes, DSCP markings or the way that classes are treated.

5.10 Extension Path (EP)

5.10.1 Network Terminating Equipment (NTE)

WBC 1G & 10G standard EP Customers must have an MSIL and equipment that complies with the MSIL requirements. Please see MSIL product documentation for more details at http://www.btwholesale.com/pages/static/Products/Data/MSIL.html

40G EP Customers and 10G & 100G Direct EP Customers do not require a MSIL. The interface for these EPs will be via an Openreach CableLink. The hand-off will be In-Building in the CP's area of the WBC POSI Node, and may be either unterminated or terminated in a rack using SC/APC connectors and 40GBase-LR4 Single Mode fibre presentation for the 40G variants, 10GBase-LR for the 10G variant and 100GBase-LR4 for the 100G variant.

5.10.2 CP equipment

The WBC Customer will present traffic over VLANs from a router (or L3 switch). It is recommended that WBC Customers should configure their network device that connects to WBC to operate at Layer 3 and NOT as a Layer 2 switch. It is possible to place intermediate Layer 2 switches between WBC and the Layer 3 switch but this will result in slower detection of link failures as it will no longer be possible to rely on transmission link failure detection and it will instead be necessary to rely on Layer 3 mechanisms. If intermediate Layer 2 switches are used, the first Layer 3 device will terminate the VLAN.

5.10.3 BGP

The WBC product will support eBGP routing towards the WBC Customer. A BGP peering will be established over the Extension Path (EP) VLAN between the AP and the WBC Customer. Where there are multiple EP VLANs between the AP and the WBC Customer then there will always be a BGP peering over each EP VLAN. BGP4 is required for all services. If a customer takes the IPv6 PTA option, a dual stack solution will be used over the same VLAN - an eBGP v6 peering will also be run between the WBC network and the CP.

The WBC Customer will provide an AS number. AS65099 & AS65100 are the chosen BEA WBC AS numbers taken from the private AS numbering range, so the WBC Customer will not be allowed to use AS65099 or AS65100 as their AS number. Prefixes from WBC will have an AS path of either "65100" or "65100, 65099". The same AS number will be used for both IPv4 and IPv6.

Note that these AS numbers are applied per virtual router (VR) and each WBC Customer is given their own VR per AP so there is no number clash between WBC Customers. A BGP standard community will be placed on the IPv6 delegated prefixes that are advertised to the CP. This will allow the CP to identify from which WBC POSI node the delegated prefix originates. The list of communities will be available on request.

For 1G services there will be a maximum of 8 EPs per AP, so in essence a maximum of 8 eBGP peers (16 peers if IPv6 for PTA is also used). For 10G, 40G and 100G services, the maximum will be 16 EPs per AP.

Note: For 1G, 10G and 100G Eps (standard and direct), the CP can choose the VLAN tag to use for each EP. For 40G EPs, the VLAN tag will be allocated by BT from the range 1800 to 1999.

The WBC Customer must always send an IPv4 default route prefix. If the IPv6 for PTA option has been chosen, the Customer must also send an IPv6 default route prefix. This is in addition to any other routes the customer sends.

Load sharing across EPs will be set by default and the BEA will load balance traffic over up to 16 equal cost paths. If a CP increases the number of EPs, they may need to make changes in their equipment should they require traffic to be load balanced over multiple paths.

In the case of multiple link topologies, if a CP does not require load sharing across their EPs there are 2 options for influencing traffic flows:

1) The CP requests BT to assign a MED/Local Preference to an Extension Path when they place their order. The MED and local preference will be set at the BT end on the BEA. MED/Local Preference cannot be set at the BEA on a per prefix basis, so this will apply to all prefixes over the EP. Different MEDs & Local Preference may be set for IPv4 and IPv6 if required. In this scenario it will be the CP's responsibility to ensure that their end of the network is set-up correctly.
2) The CP sets MED/Local Preference within their own network to influence their traffic flows. MEDs will be honoured by BT. This option will allow individual prefixes to be handled differently over the EPs.

BT will set the hold time to 180 seconds and keep alive to 60 seconds for both v4 and v6 peerings. If the WBC Customer wishes to reduce the convergence time then they can set lower values for the hold time and “keep alive” time. If lower values are selected then BT will automatically honour them. It is a requirement that WBC Customers do not use BGP timers below a hold time of 30 seconds and a “keep alive” time of 10 seconds.

MD5 authentication should be configured to the WBC Customer router.

16 bit vs. 32 bit AS numbers

The WBC service currently supports 16 bit AS numbers only. RIPE no longer issues 16 bit AS numbers. They operate AS Number assignments from an undifferentiated 32-bit AS Number allocation pool. If WBC supports 32 bit ASNs in the future, then this will be reflected in an updated SIN.

5.10.4 Session Based Load Balancing

Session based load balancing only applies to L2TP implementation and relies on the WBC Customer sending tunnel endpoints with different MEDs over the multiple peers. The WBC Customer will then use RADIUS to respond with different tunnel end points and the routing will ensure that the sessions are established over different EPs/peers. The WBC Customer will be responsible for load-balancing the downstream traffic.

With multiple eBGP peers load-balancing is the default behaviour. This will permit the WBC Customer to load balance either PTA or L2TP traffic. Limitations with load balancing lead to inefficient use of available bandwidth if the number of source/destination pairs this typically will not be a problem in PTA implementation, however L2TP by its nature results in very few source/destination pairs. If the WBC Customer is running L2TP and they detect poor efficiency they can arbitrarily increase the number of tunnel endpoints they send to BT by configuring their RADIUS and other network elements appropriately.

5.10.5 Frame sizes

The maximum MTU size on the BT Broadband network for End User data packets is 1500 bytes. The WBC infrastructure ensures that End User data packets up to and including 1500 bytes are carried without fragmentation.

This feature applies at the layer being carried by WBC PPP. If tunnelling protocols (such as PPTP or IPSEC) are being used over this connection WBC Customers should advise end users to configure the MTU on these higher layers correspondingly smaller to ensure that the packets being carried by the WBC PPP layer do not exceed 1500 bytes.

WBC Customers using PPPoE should read section 4.3 to assess its impact on MTU size.

For each Extension Path (EP) the CP should set their MTU large enough to allow PTA and L2TP packets in the downstream direction without fragmentation. WBC recommends that the first networking device in the CPs network should terminate the EP layer 3 VLAN and the eBGP peering. The BEA will have its MTU set to 1900 bytes on the EP, and eBGP updates may be sent as 1900 byte packets, therefore the CP should ensure that there are no devices in their network that would prevent packets of that size.

MSIL may add further restrictions on the MTU size, please refer to their documentation to clarify.

Customers should be aware that when purchasing interconnect products, such as an MSIL, for use with WBC, the maximum usable throughput on the WBC product is normally less than the headline rate on the interconnect product due to network protocol overheads. The exact maximum usable throughput will depend upon packet size and whether L2TP is used or not. As a guideline BT Wholesale suggests that customers do not exceed 90% of the headline bandwidth of the interconnect e.g. 900M for a 1G MSIL and 9G for a 10G MSIL.

5.10.6 EP Policing

- For 1G & 10G EPs, each EP is policed downstream at the MSIL end and upstream at the BEA end to conform to EP bandwidth; EP policing and traffic discard are independent of QoS marking and will be done on Ethernet frames;
- There is no EP policer for 40G Eps, and for the direct EP variants (both 10G and 100G)

5.11 Resilience under Catastrophic Node Failure

A CP can only achieve full Catastrophic Nodal Disaster (CND) protection at a Point of Service Interconnect (PoSI) node if they purchase WBC EPs/MSILs parented to all 20 PoSI nodes. Any less than this and a (sliding scale) subset of End User customers parented to a node that fails will not be recovered to other nodes around the country.

Customers who use aggregation products to deliver WBC from a subset of POSI nodes will not be able to connect to their EU's who are re-parented to those nodes in the event of a CND. In this event, it may be possible to re-configure EU's to connect via the same aggregation product, but over an undetermined period of time.

The MSIL product has launched a Secure MSIL product option where a pair of MSILs will be connected to separate EEAs. It is recommended that the Secure MSIL offering is purchased to provide additional resilience taking the WBC service for 1G & 10G EP hand-offs. For 40G EPs, it is a requirement that 40G EPs are bought in pairs for resilience. For the Direct EP product the recommendation is that the product is bought in pairs at each node for resilience.

6. RADIUS protocol & DNS

6.1 General

RADIUS Authentication is necessary for the normal operation of WBC. RADIUS Accounting is provided as an option.

RADIUS traffic is delivered from BT RADIUS servers connected to BT's Internet Core network. WBC Customers will therefore need an Internet connection to handle RADIUS traffic.

A single shared secret is required, common to all BT and WBC Customer RADIUS servers used with the service.

A CP may have up to two RADIUS servers (working as Primary/Secondary) associated with their WBC service.

The WBC product supports the following RADIUS packet types:

ID	Packet Type
1	Access-Request
2	Access-Accept
3	Access-Reject
4	Accounting-Request
5	Accounting-Response

Table 13 - Packet type id

All RADIUS requests from the BT RADIUS must be acknowledged. If a BT RADIUS server does not receive a response within a defined timeout period to a request it will re-try a number of times to the same server. If no response is received and the WBC Customer has nominated a back-up RADIUS server, this will be tried. The values for the timeout and number of retries may be defined, within the constraints of the product, by the WBC Customer when placing their order.

If no response is received from either WBC Customer RADIUS server, the End User's PPP session will be terminated locally on the BRAS and assigned a local IPv4 address specifically for this purpose. End User HTTP requests may be re-directed to a local server informing them of a network problem. All other IP traffic will be filtered out.

A RADIUS server should always respond to a request from a valid RADIUS client. A silent discard is not appropriate, as the platform will attempt a retry. The only occasion where a silent discard is warranted is where the authenticator fails to match. Otherwise a response should always be made to prevent the BT RADIUS server from re-transmitting

The WBC RADIUS will support IPv6 for PTA. Please refer to RFC 3162[30] and RFC 4818[31].

Note: The WBC RADIUS infrastructure will use IPv4 addressing.

6.2 RADIUS Authentication

An incoming PPP session to a BRAS will trigger a RADIUS Access-Request following the response to the CHAP challenge from the End User, using attributes shown in RFC 2865[17].

Other attributes that may be included in the Access-Request should be ignored.

The WBC Customer's RADIUS must respond with attributes from

Table 15 - PTA Access-Accept Attributes

or

Table 16 – L2TP Access-Accept Attributes

Note 1: The option specified in RFC 2865[17], where the CHAP challenge is a 16 bit value, is used. I.e. the Access-Request authenticator contains the CHAP challenge and the CHAP-Password attribute contains the CHAP identity and response string.

Note 2: For information, the RADIUS Access-Request will include Vendor Specific Attributes.

Note 3: If IPv6 for PTA is chosen, no additional RADIUS Access-Request attributes will be present.

The WBC Customer RADIUS can return either an Access-Accept, or an Access-Reject in response to an Access-Request.

The set of attributes supported in an Access-Reject are shown below. A RADIUS Accounting start or stop record will not be sent in response to an Access-Reject.

No	Attribute	Value	Comment
33	Proxy-State	unique string	Exact copy of Proxy-State from the corresponding Access-Request. Please refer to section 6.4 RADIUS Client

Table 14 – Access-Reject Attributes

The range of Access-Accept attributes supported is shown in

Table 15 - PTA Access-Accept Attributes

or

Table 16 – L2TP Access-Accept Attributes

Two distinct BRAS operating modes are supported, selected by the set of RADIUS attributes returned. End User sessions can be either:

- **PTA mode**: Terminated locally on the BRAS. End User IP data is routed over the WBC Customer interface; or
- **L2TP mode**: Extended to an LNS using L2TP Version 2, where the PPP session is terminated.

No	Attribute	Value	Comment
6	Service-Type	Framed	Required
7	Framed-Protocol	PPP	Required
8	Framed-IP-Address	<IP Address>	Required (Note 3)
22	Framed-Route	String	Optional (Note 4)
25	Class	unique string	Optional (Note 5)
27	Session-Timeout	>= 7200	Optional (Note 6)
28	Idle-Timeout	>= 7200	Optional. (Note 6)
33	Proxy-State	unique string	Exact copy of Proxy-State from the corresponding Access-Request. Please refer to section 6.4 RADIUS Client
26	Vendor Specific	Redback VSA 36 Framed-Named-Pool	This is a vendor specific attribute required by the Redback BRAS to select a named IP address pool in an access accept message It should be supplied with attribute 88 and the string should be identical (Note 7)
88	Framed-Pool	unique string	This is the IETF standard Radius attribute used to select a named IP address pool in an access accept message (Note 7)
123	Delegated-IPv6-Prefix	OctetString	Optional (Note 8)

Table 15 - PTA Access-Accept Attributes

Note 3: *The value255.255.255.255 (0xFFFFFFFF) which would allow the End-User's CPE to assign a locally configured IP address is not supported. If the value of Framed-IP-Address is 255.255.255.254 then an address from a local IP pool will be assigned.*

WBC Customers may assign individual (static) IP addresses to End Users via the Framed-IP-Address attribute. These addresses will not be summarised into BGP which could lead to an un-acceptable number of host addresses in the BGP routing tables.

Note 4: *A framed route may be injected into the BRAS routing table, via the IP address assigned using Framed-IP-Address. If the value of Framed-IP-Address is 255.255.255.254 to assign an address from a local IP pool, the next hop address in the Framed-Route value string should be 0.0.0.0*

Note 5: *Only a single Class attribute per packet is supported by WBC.*

Note 6: *Absolute session and idle timeouts may be returned but the value of these should not normally be less than 7200 seconds to avoid possible congestion within BT's infrastructure.*

Note 7: Attribute 88 and Redback VSA 36 are only required if named address pools are used. If they are not sent the default address pool will be used.

Note 8: *Optional for CPs who choose to take IPv6 for PTA Mode. If Delegated-IPv6-Prefix is returned, the value will be used as the IPv6 delegated prefix for the End User. If not provided, the End User will be dynamically provided with an IPv6 delegated prefix from a Default IPv6 pool.*

No	Attribute	Value	Comment
25	Class	unique string	Optional (Note 8)
33	Proxy-State	unique string	Exact copy of Proxy-State from the corresponding Access-Request. Please refer to section 6.4 RADIUS Client
64	Tunnel-Type	L2TP	
65	Tunnel-Medium-Type	IPv4	
67	Tunnel-Server-Endpoint	LNS IP address	(Note 9)
69	Tunnel-Password	password	Optional (Note 10)
82	Tunnel-Assignment-ID	String	Optional (Note 11)
83	Tunnel-Preference	Integer	Optional (Note 12)
90	Tunnel-Client-Auth-ID	String	Optional (Note 13)

Table 16 – L2TP Access-Accept Attributes

Note 8: Only a single Class attribute per packet is supported by WBC.

Note 9: Only the dotted decimal notation format required in RFC 2868 is supported.

Note 10: This attribute is used to populate the L2TP Challenge AVP as described in RFC 2661[16]. Reference should be made to the security considerations for the Tunnel-Password attribute described in RFC 2868[18].

Note 11: This attribute allows sessions to be grouped in separate tunnels between the same endpoints. Creating a large number of tunnels between the same end points can be detrimental to both LNS and BRAS performance, so should be used with caution.

Note 12: This attribute is used to group tagged attributes as described in RFC 2868[18]. Tagging is only required if more than one Tunnel-Server-Endpoint is used.

Note 13: The Tunnel-Client-Auth-ID is used to populate the L2TP Host Name AVP as described in RFC 2661[16]. If the Tunnel-Client-Auth-ID attribute is not used, the default host name from the BRAS will be used. The format of this host name is unspecified but will be unique for any given BRAS.

Session and idle timeouts values could conceivably be returned with the L2TP attributes but this is inappropriate. Session and idle timeouts if required should be applied to an End User session on the LNS.

6.3 RADIUS Accounting

RADIUS Accounting is optional. If configured, Accounting records will be passed to the CP.

If IPv6 for PTA is being used, IPv4 & IPv6 will be recorded in the same accounting records i.e. there will be one set of statistics covering both IPv4 and IPv6. The one difference is the presence of the IPv6 delegated prefix in the accounting record.

The following accounting records will be generated:

1) An Accounting start record is generated after an End User has been authenticated.

- If PTA mode is used, the Accounting start record is delayed until a Network Control Protocol (NCP) stack has completed and an IP address has been assigned to the End User. In the case of dual-stack (IPv4 & IPv6) working, the accounting start may be sent when the IPv4 address has been assigned and before the IPv6 address has been assigned or vice-versa.
- In L2TP mode, an Accounting start record is sent following L2TP session establishment.

2) When a second NCP completes and an additional IP address is assigned to an existing PPP session, an Event Triggered interim accounting record will be generated. It will include the IP address information (IPv4 address or IPv6 delegated prefix in the case of IPv6 for PTA). This will indicate that IPv4 or IPv6 is now enabled for the PPP session. The counters in these interim accounting records will not be incremented.
3) An interim accounting record will be generated every 2 hours after the start record.
4) An Accounting stop record is sent when the End User disconnects.

No	Attribute	Value	Comment
1	User-Name	CHAP username	As entered by End User
4	NAS-IP-Address	BRAS source IPv4 address	
5	NAS-Port	BRAS local port	
6	Service-Type	Framed	
7	Framed-Protocol	PPP	
8	Framed-IP-Address		IPv4 address. PTA mode only
25	Class	unique string	If set in Access-Accept
27	Session-Timeout		If set in Access-Accept
28	Idle-Timeout		If set in Access-Accept
31	Calling-Station-Id	String	Service Id
32	NAS-Identifier	String	Unique BRAS id
33	Proxy-State	unique string	Set by BT RADIUS
40	Acct-Status-Type	Start (1) Stop (2) Interim (3) Acct-On (4) Acct-Off (5)	
41	Acct-Delay-Time		
42	Acct-Input-Octets		
43	Acct-Output-Octets		
44	Acct-Session-Id		
45	Acct-Authentic		
46	Acct-Session-Time		
47	Acct-Input-Packets		
48	Acct-Output-Packets		
49	Acct-Terminate-Cause		
52	Acct-Input-Gigawords		
53	Acct-Output-Gigawords		
55	Event-Timestamp		
61	NAS-Port-Type	Virtual (5)	
123	Delegated-IPv6-Prefix	OctetString	The Delegated IPv6 Prefix associated with the PPP session.

Table 17 - Accounting Attributes

Note: For information, the RADIUS Accounting attributes will include Vendor Specific Attributes.

6.4 RADIUS Client

The RADIUS packets are proxied via BT RADIUS Servers. WBC Customers will need to configure these addresses as their RADIUS clients.

6.5 DNS

CPs who take the PTA option will need to provide DNS servers. Up to two DNS servers will be provided, with an option to work as load-balanced or primary/secondary.

Dual stacked PTA end users (IPv4 and IPv6) will use the DNS servers on IPv4 infrastructure using AAAA records.

7. Product Enhancements

Please see **http://www.btwholesale.com** for the latest information on product enhancements / Developments.

8. References

[1]	G.992.1	Asymmetric digital subscriber line (ADSL) transceivers
[2]	G.992.3	Asymmetric digital subscriber line transceivers 2 (ADSL2)
[3]	G.992.5	Asymmetric digital subscriber line (ADSL) transceivers – Extended bandwidth ADSL2 (ADSL2plus)
[4]	RFC0791	Internet Protocol: DARPA Internet Program Protocol
[5]	RFC1332	The PPP Internet Protocol Control Protocol (IPCP)
[6]	RFC1334	PPP Authentication Protocols
[7]	RFC1483	Multiprotocol Encapsulation over ATM Adaptation Layer 5
[8]	RFC1631	The IP Network Address Translator (NAT)
[9]	RFC1661	The Point-to-Point Protocol (PPP)
[10]	RFC1771	A Border Gateway Protocol 4 (BGP-4)
[11]	RFC1772	Application of the Border Gateway Protocol in the Internet
[12]	RFC1877	PPP Internet Protocol Control Protocol Extensions for Name Server Addresses
[13]	RFC1994	PPP Challenge Handshake Authentication Protocol (CHAP)
[14]	RFC2364	PPP Over AAL5
[15]	RFC2516	A Method for Transmitting PPP Over Ethernet (PPPoE)
[16]	RFC2661	Layer Two Tunnelling Protocol "L2TP"
[17]	RFC2865	Remote Authentication Dial In User Service (RADIUS)
[18]	RFC2868	RADIUS Attributes for Tunnel Protocol Support
[19]	RFC2923	TCP Problems with Path MTU Discovery
[20]	SIN 346	BT's G.DMT ADSL interface
[21]	WBC CRFs	See www.btwholesale.com/pages/static/Products/Broadband/Wholesale_Broadband_Connect//index.htm for copies of the latest Customer Requirement Forms
[22]		See http://www.openreach.co.uk/orpg/products/newproducts/fttp/fttp.do for Openreach GEA documentation
[23]	TR-202	ADSL2/ADSL2plus Low-Power Mode Guidelines
[24]	RFC 6598	IANA Reserved IPv4 Prefix for Shared Address Space
[25]	RFC 5072	IP Version 6 over PPP
[26]	RFC 4443	Internet Control Message Protocol (ICMPv6) for the Internet Protocol Version 6 (IPv6) Specification
[27]	RFC 4861	Neighbor Discovery for IP version 6 (IPv6)
[28]	RFC 2460	Internet Protocol, Version 6 (IPv6) Specification
[29]	RFC 5095	Deprecation of Type 0 Routing Headers in IPv6
[30]	RFC 3162	RADIUS and IPv6
[31]	RFC 4818	RADIUS Delegated-IPv6-Prefix Attribute

SINs are available from http://www.btplc.com/sinet/

For copies of referenced documents please see the contacts on the Document Sources page at http://www.btplc.com/sinet/

9. Abbreviations

Acronym	Expansions
AAAA	Authentication, Authorisation, Accounting & Auditing
ADSL	Asymmetric Digital Subscriber Line
AP	Aggregation Point
AS	Autonomous System
ASN	Autonomous System Number
ATM	Asynchronous Transfer Mode
ATU-r	ADSL Terminating Unit – remote
AVP	Attribute-Value Pair
BBCR	BroadBand Customer Reporting
BEA	Broadband Edge Aggregator
BGP	Border Gateway Protocol
BRAS	Broadband Remote Access Server
BT	British Telecommunications plc
BTW	BT Wholesale
CHAP	Challenge Handshake Authentication Protocol
CMSAN	Copper Multi Service Access Node
CP	Communications Provider
CPE	Customers' Premises Equipment
CRF	Customer Requirements Form
DARPA	Defense Advanced Research Project Agency [USA]
DHCP	Dynamic Host Configuration Protocol
DLM	Dynamic Line Management [BT system]
DNS	Domain Name System/Server
DSLAM	Digital Subscriber Line Access Multiplexer
DSL	Digital Subscriber Line
DSCP	Differentiated Services Code Point [IETF]
eBGP	Exterior Border Gateway Protocol
EF	Expedited Forwarding
EP	Extension Path
EU	End User
EUA	End User Access
FTTC	Fibre To The Cabinet
FTR	Fault Threshold Rate
FTTP	Fibre To The Premises
G.DMT	G-series Discrete Multi-Tone [ITU-T]
HEC	Header Error Check
HTTP	HyperText Transfer Protocol
IANA	Internet Assigned Numbers Authority
ICMP	Internet Control Message Protocol [IETF]
IEEE	Institute of Electronic and Electrical Engineers
IETF	Internet Engineering Task Force
IGP	Interior Gateway Protocol
IGMP	Internet Group Management Protocol
IP	Internet Protocol
IPCP	Internet Protocol Control Protocol [IETF]
IPR	Intellectual Property Rights
IPstream	BT IPstream
IPsC	BT IPstream Connect
ISDN	Integrated Services Digital Network
ITU	International Telecommunications Union
ITU-T	International Telecommunication Union - Telecommunications Standardization Sector

Acronym	Expansions
kbit/s	Kilobits per second
L0	ADSL2/2plus power mode indicator showing the ADSL link is fully functional (full power mode)
L2	ADSL2/2plus power mode indicator showing the ADSL link is active but a low power signal conveying background data is sent from the ATU-C to the ATU-R. A normal data carrying signal is transmitted from the ATU-R to the ATU-C (low power mode)
L2TP	Layer 2 Tunnelling Protocol
L3	ADSL2/2plus power mode indicator showing that there is no signal transmitted at the U-C and U-R reference points. The ATU may be powered or unpowered in L3. (out of sync)
LAC	L2TP Access Concentrator (function in BAS)
LAN	Local Area Network
LCP	Link Control Protocol [IETF]
LLC	Logical Link Control
LNS	L2TP Network Server (e.g. Home Gateway)
LTS	Layer 2 Tunnelling Protocol (L2TP) Tunnel Switch
MAC	Migration Authorisation Code
Mbit/s	Megabits per second
MD5	Message Digest version 5
MED	Multi-Exit Discriminator
MRU	Maximum Receive Unit
MSAN	Multi-Service Access Node
MSE	Multi-Service Edge router
MSIL	Multi Service Interconnect Link
MSS	Maximum Segment Size
MTU	Maximum Transmission Unit
NAS	Network Access Server
NAT	Network Address Translation
no NAT	Optional removal of NAT for Static IP Addressing
NCP	Network Control Protocol
NTE	Network Termination Equipment
NTE5	Network Terminating Equipment No. 5
OAM	Operations and Maintenance
OLO	Other Licensed Operator
ONT	Optical Network Termination
OSPF	Open Shortest Path First protocol
OSS	Operational Support System
PAP	Password Authentication Protocol [IETF]
PAT	Port Address Translation
PBR	Policy Based Routing
PF	Premium Forwarding
POP	Point of Presence
PPP	Point to Point Protocol
PPPoA	PPP over ATM
PPPoE	PPP over Ethernet
PSD	Power Spectral Density
PSTN	Public Switched Telephone Network
PTA	PPP Termination and Aggregation
PVC	Permanent Virtual Connection
QoS	Quality of Service
RADIUS	Remote Authentication Dial In User Service
RAM	Random Access Memory
RAS	Remote Access Server

Acronym	Expansions
RFC	Request for Comment
RIPE	Réseaux IP Européens
RJ	Registered Jack
RJ45	Registered Jack 45
RT	Real Time
SID	Service Identifier
SIN	Suppliers' Information Note
SNAP	Sub Network Access Protocol
STIN	Suppliers' Trial Information Note
TCP	Transmission Control Protocol [IETF]
TCP/IP	Transmission Control Protocol/ Internet Protocol
UDP	User Datagram Protocol [IETF]
USB	Universal Serial Bus
UTP	Unscreened Twisted Pair
VC	Virtual Connection [ATM]
VCI	Virtual Connection Identifier [ATM]
VLAN	Virtual Local Area Network
VP	Virtual Path
VPI	Virtual Path Identifier [ATM]
WAN	Wide Area Network
WBC	Wholesale Broadband Connect
WLR	Wholesale Line Rental

10. History

STIN Issue 1.0	19/07/07	Draft release to give industry early visibility of the WBC interfaces
STIN Issue 1.1	11/09/07	Redrafted version to cover the ST3 service trial, including addendum at start
STIN Issue 1.2	14/12/07	Amended in order to improve clarity, eg removal of duplicated information
STIN Issue 1.3	25/01/07	Added that the WBC Customer is required to advise BT which address pools it will be using,so that traffic from that WBC Customer's users will be routed to that WBC Customer. Also minor amendments to improve clarity, including expansion of list of abbreviations.
STIN Issue 1.4	29/02/08	Updated to add in references to FTTP trial.
SIN Issue 1.0	02/04/08	Updated to move to SIN issue
SIN Issue 1.1	08/05/08	Three Accounting Attributes added to Table 17.
SIN Issue 1.2	19/12/08	FTTP trial section deleted Clarification: - Limiting the amount of Routes in a VRF clarified. - Policy Based Routing clarification - QoS marking support and policing behaviour on L2TP clarified. - Real Time Service - The introduction of more QoS services will result in different behaviour of the BT network Removed WBC Banded rates table Removed Access request Attributes Table Corrected table 3 AP Routes have had more clarification on their limits
SIN Issue 1.3	29/05/09	Section 5.1 Added section on resilience routing for Catastrophic Node Disasters. All section 5 subsections have incremented by one. Section 5.7 Added discussion of 16bit vs 32 bit ASN numbers. Section 5.9 PBR now available and clarified. Section 5.11 Release dates for Priority and Sub best efforts QoS updated
SIN Issue 1.4	04/11/09	Section 3.2.3.2 moved interleaving opt –in statement. Replaced tables V.1/G992.3, V.2/G992.3, V.1/G992.5 and V.3/G992.5 Inserted new section 4.2 to add OAM statement.

		Section 5.11 Release dates for Priority and Sub best efforts QoS updated
SIN Issue 1.5	13/01/10	Added Description of new ADSL2+ Annex M service in section 3.2, 3.22 and 3.24
SIN Issue 1.6	05/05/10	Section 3.2.4 added note Section 3.4.1 Moved section 4.2 to new section 4.1.1 New 4.2 section Section 4.4.2 clarification of eBGP default route requirement Section 5.2 revised points 2 and 3 Section 7.1 headline bandwidth on WBC with Interconnect Section 3.1 Status change FTTC & FTTP Section 3.2 Minor change, status change Annex M and note 2 added Section 3.3 Reworded for clarity and status change Section 4.2 Routing behaviour of L2TP Section 4.3 Guidance of IP address range used by BT Infrastructure Section 4.4.2 Default routes Section 5.1 Secure MSIL added Section 5.2 Physical connection points 2 and 3 reworded Section 5.11 Complete reworded and status updated Section 5.12 Complete reworded and status updated Section 5.14 Test account, more info added Section 6.4 Status change Section 7.1 Overhead statement added Section 8 New URL
SIN Issue 1.7	15/09/10	Changes to support multiple named address pools for PTA usage Changes to section 4.3 and 6.2
SIN Issue 1.8	22/10/10	Section 1.1, 1.2, 5.13 & ANNEX B - Changes to introduce Content Connect Section 5.12 – Clarified the handling of Real Time traffic for L2TP
SIN Issue 1.9		Section 3.2.2 Clarification of upload speed for ADSL1 uncapped product Section 4.3 Added a current restriction on Multiple Named Address Pools for PTA
SIN Issue 1.10	June 2011	Section 3.2.2 – Minor change to clarify Annex M upstream line speed Section 4.3 & 5.3 – Clarified that 172.16.0.0 – 172.31.255.255 is a restricted range Section 4.3 – Clarified support for PTA private IP addressing Section 4.4.2 - Updated to reflect the increase to 16 EPs per AP for 10G EPs Section 5.3 – Clarified allocation of addressing for WBC APs & EPs Section 5.10 – Clarified Real Time “policer”

		Section 5.11 – Modified the QoS code rollout & Assured Rate traffic handling statements Section 5.12 – Clarifications to the Real Time QoS solution
SIN Issue 1.11	Oct 2011	Minor corrections Addition of BT Cool Broadband® low power mode information
SIN Issue 1.12	Jan 2012	Section 1.2 & Annex B – Removed references to Content Connect Basic; removed references to SIN 504 and introduced reference to CDN Integration and User Guide Section 1.2 & Annex C – Introduction of TV Connect Section 4.3, 5.3.1 – Introduced new rules for allocation of private addressing for PTA end users Section 4.4.2 – Added in AS65099 as a WBC AS number Section 5.3.2 & 5.3.4 – Updated addressing to reflect introduction of new WBC BRAS (MSE) & additional wording updates Section 5.3.5 – Removed section as not relevant to the WBC CP interface Section 5.11, 5.12, 5.13 – Changes to reflect rollout of BRAS QoS code and introduction of Real Time QoS Section 6.2 – Added a note that Authentication Requests include Vendor Specific Attributes; added a note that a single class attribute per packet is supported; re-arranged the numbering of the Notes. Section 6.3 – Added a note that Accounting includes Vendor Specific Attributes Section 6.4 – Revised wording for RADIUS client configuration Annex B – Removed the private address limitation for Content Connect; minor change to wording for handling of IP Address Ranges & Content Connect charging
SIN Issue 1.13	March 2012	Section 3.2.2, 3.2.3.2 – Added a clarification regarding the maximum upstream rate for Annex M when used with Interleaving.
SIN Issue 1.14	July 2012	Section 4.3 & 5.3.1 – Updated to reflect the introduction of support for Shared Address Space for PTA IP Addressing Section 4.4.2, 4.4.5, 5.1, 7.1 – Updated to reflect the introduction of 40G EPs Section 4.4.2 – Clarification of the MTU size used by BT Section 5.9 – Changes to the Policy Based Routing mechanism
Issue 1.15	September 2012	Annex C updated.
Issue 2.0	October 2012	General – changed references to the WBC website General – minor editorial changes Sections 3, 4, 5. & 7. – Restructured to aid readability. Section 4 changed to cover End User Handover. Merged Section 7 (Extension Path) into Section 5.

		Section 3.2.2 – Merged the Annex M table with Table 1 & changed orientation of the table Section 3.2.3.3 – Added Custom stability option; updated the information related to banded configurations. Sections 3.2.2, 3.2.3, 3.2.3.3, 5.8, 5.9 – Updates to reflect introduction of Fixed Rate product Sections 4.4.3, 6.2 & 6.4 – Removed references to RADIUS Double Dip Sections 4., 5., 6. & 8. – Updates to reflect the introduction on IPv6 for PTA Mode. Section 5.4 – reworded the statement about asymmetric routing Section 5.11.1 – Added information about physical hand-off for 40G EPs Section 6. – Added a DNS sub-section Section 6.1 – Removed statement that interim accounting records are optional – interim accounting will be included if RADIUS accounting is selected Section 8. – Added references for IPv6 RFCs Annex B & C – Clarified that IPv6 is not applicable to Content Connect and TV Connect
Issue 2.1	December 2012	Annex C updated to include TV Connect over WBC Copper
Issue 2.2	January 2013	Annex C amendment to clarify the availability of the service
Issue 2.3	March 2013	Annex C updated to include TV Connect for L2TP handover variant of WBC
Issue 2.4	April 2013	Section 1,2, 5.6, 5.8, 5.9, Annex B – Removed references to Assured Rate traffic. Section 6.2 – Clarified that the Access-Accept attributes must be returned.
Issue 2.5	June 2013	Section 5.10.5 – Clarified the MTU size for the EPs.
Issue 2.6	July 2013	Section 3.2.2 – Updated to reflect that the Fixed Rate product variant is now launched. Section 5.6 – Updated to reflect that Real Time QoS is now launched. Section 5.8 – Revised to add in the range of end user bandwidth options for Real Time QoS. Updated to reflect that Real-time QoS is now supported on the Fixed Rate product variants. Annex C – Revised to remove the caveats about TV Connect support for WBC Copper and L2TP; removed reference to future priority markings; minor editorial changes.
Issue 2.7	December 2015	Section 1., 5.9 & 9 – Changes to reflect the withdrawal of Content Connect Annex B – Content Connect Annex removed Annex C – Wording changes in line with withdrawal of Content Connect Change SINet site references from http://www.sinet.bt.com to http://www.btplc.com/sinet/
Issue 2.8	June 2017	Details of Flex POSI added

ANNEX A – Warning regarding throughput limitation imposed due to RWIN settings and increased latency.

In TCP networking the RWIN (TCP Windows size) is the amount of data that can be accepted by a computer without being acknowledged. If the sender of the data hasn't received an acknowledgement for the first packet it has sent, then it will stop and wait and once this wait exceeds a limit, it may retransmit. This is how TCP achieves reliable data transfer. Even with no data loss in a network, this RWIN limitation can impose a throughput limit according to the following formula depending on the overall Round Trip Time (RTT):

Throughput <= RWIN (bytes) / RTT (seconds)

For example, on a DSL system running in fast mode, with a round trip time of 8ms and a default RWIN size of 65535bytes, the throughput is limited to a maximum of:

Throughput max =65535/0.008 = 8.19Mbps

If this line is placed in interleaved mode with an interleaving delay of 16ms, then potentially the round trip time could be increased to between 32-40ms dependant on the delay imposed by the DSLAM. If we recalculate with the round trip time at 40ms, then

Throughput max = 65535/0.04 = 1.64Mbps

In order to address this restriction, it is possible on most systems to increase the default RWIN size to improve this. For example, increasing the RWIN to 327600 would return the maximum throughput back to the original figure:

Throughput max = 327600/0.04 = 8.19Mbps

It should also be noted that this restriction is imposed only to a single TCP stream. It is possible to utilise the entire available DSL bandwidth through the use of multiple TCP streams and or UDP traffic.

ANNEX B Flex POSI

Background

The original WBC network design was based around the BRAS (Broadband Remote Access Server) function being located at the 20 WBC interconnect nodes. Whilst some of these BRASs are still in service, BT Wholesale has undertaken a program of work to introduce new hardware supporting the BRAS function closer to the end user, these new BRASs are often referred to as MSE (Multi-Service edge) devices.

Currently a VLAN exists between the MSAN (ADSL) or Cabinet interconnect point (FTTx) and the BRAS. For the original BRAS this provides a data path all the way to the interconnect node but for the MSE, as the MSE is not located at the interconnect node and is closer to the end user, the VLAN only goes part way and a VPN is used to provide the connectivity from the MSE back to the interconnect node. The original product specification dictates that all traffic from a given exchange needs to be routed through a single fixed interconnect node – this aligns with those instances where the BRAS is at the interconnect node with a layer 2 connection between the BRAS and MSAN or cabinet but for the case where a MSE BRAS is used the VPN part of the connectivity could be configured to be more flexible in how the traffic is routed.

Flex POSI options and constraints

BT Wholesale intends to make available in 2017 a more flexible WBC interconnect product feature, called Flex POSI. The feature will allow the routing of MSE based broadband traffic to take advantage of the routing capability of the VPN and will effectively allow the CP to determine which interconnect node to be used. In order to achieve this, configuration changes need to be applied to the VPN carrying the CPs traffic and to the router devices at the interconnect nodes.

As a result, when the product feature is first made available it will not be possible to offer the feature across the entire WBC platform due to constraints with the underlying network architecture. The feature can be supported for existing WBC CPs where the following criteria are met:

- The CPs connectivity will need to have undergone a network infrastructure level migration (planned for 2017) to reconfigure how the broadband traffic is routed at the interconnect node. Each current direct customer of WBC will be contacted to discuss the timing of this migration, whilst any new WBC customers will be configured on the new infrastructure (though will be subject to the standard lead times for the provision of the product).
- The traffic will need to route through a MSE BRAS – non MSE BRAS traffic will continue to be routed to the current WBC interconnect node as described in WBC dataset 37.
- The CP will need to have opted in to the Flex POSI product offering so that BT can reconfigure the VPN(s) associated with the customer to route the traffic appropriately

The network infrastructure migration activity is required to re-route traffic through a new Broadband Edge Aggregator (BEA) device (this new device may also be referred to as MSC or Multi-Service Core device). Within the constraints and criteria listed above, Flex POSI will allow the CP to dictate which WBC interconnect node the WBC traffic will route through.

For flex POSI the following CP topologies can be supported:

1. Groupings of MSE selected by the CP handing off at a set of POSI's of the CP's choosing.
2. Primary/secondary POSI.
    - A group of MSE (selected by the CP) will always have a primary site that they are mapped to for handoff to a CP.
    - Within a group of MSEs, a subset or all of them can be selected to handoff at a secondary POSI site, by applying a backup POSI ID. Again the secondary POSI grouping will be chosen by the CP. Only a single Secondary POSI can be selected per MSE.

- The secondary POSI will only be used if the primary POSI fails. This will be controlled by community IDs; local preference on the MSE; and MED on the eBGP peer to the CP.

It is expected that CPs that use L2TP will prefer example 1 where all MSE are mapped to a group of POSI's. This gives the CP to control the egress point by advertising specific tunnel endpoints with the BGP MED set. Example configurations are shown below:

MSE to POSI mapping using Option 1 – No resilience

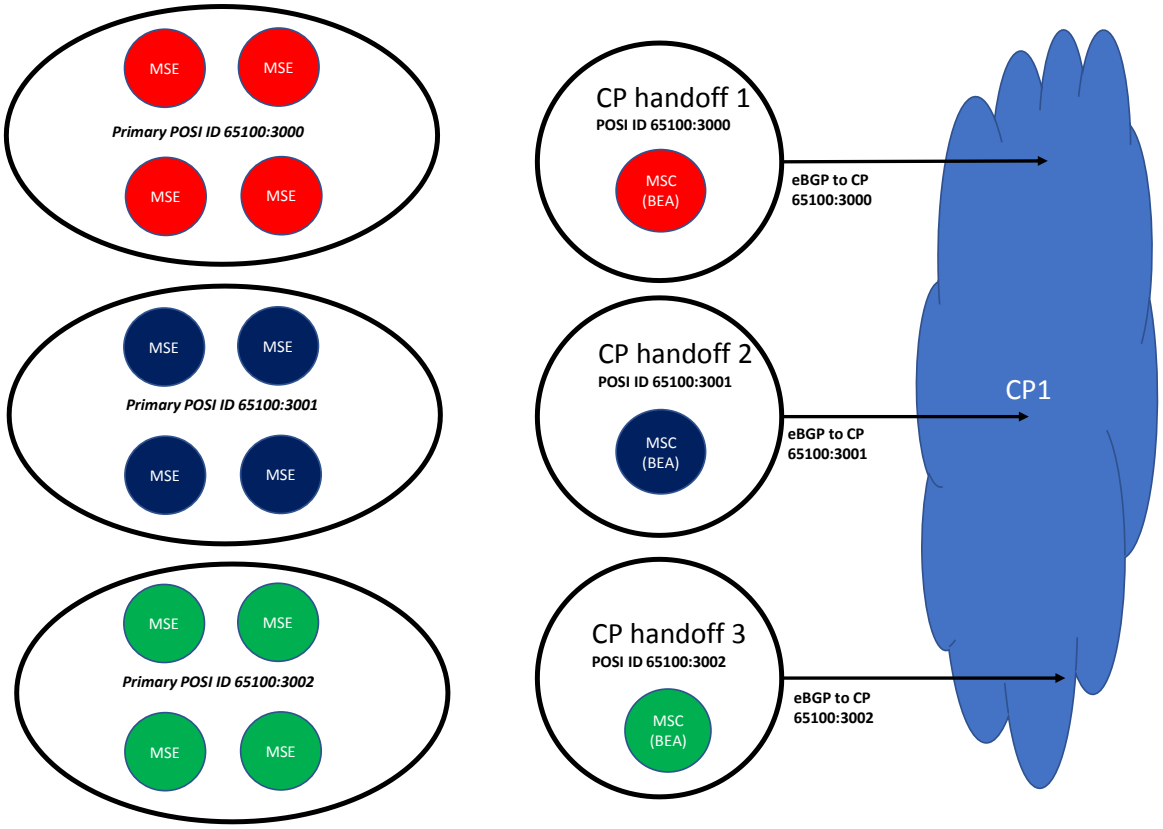

This configuration maps each MSE to a single interconnect node, much like the original WBC setup though with Flex POSI it is possible to pick different interconnect nodes for each MSE than those defined in the WBC dataset 37. Whilst possible, this configuration is not recommended due to the lack of resilience offered – loss of an interconnect to the CP would result in a loss of connectivity to the MSE BRASs associated with that interconnect.

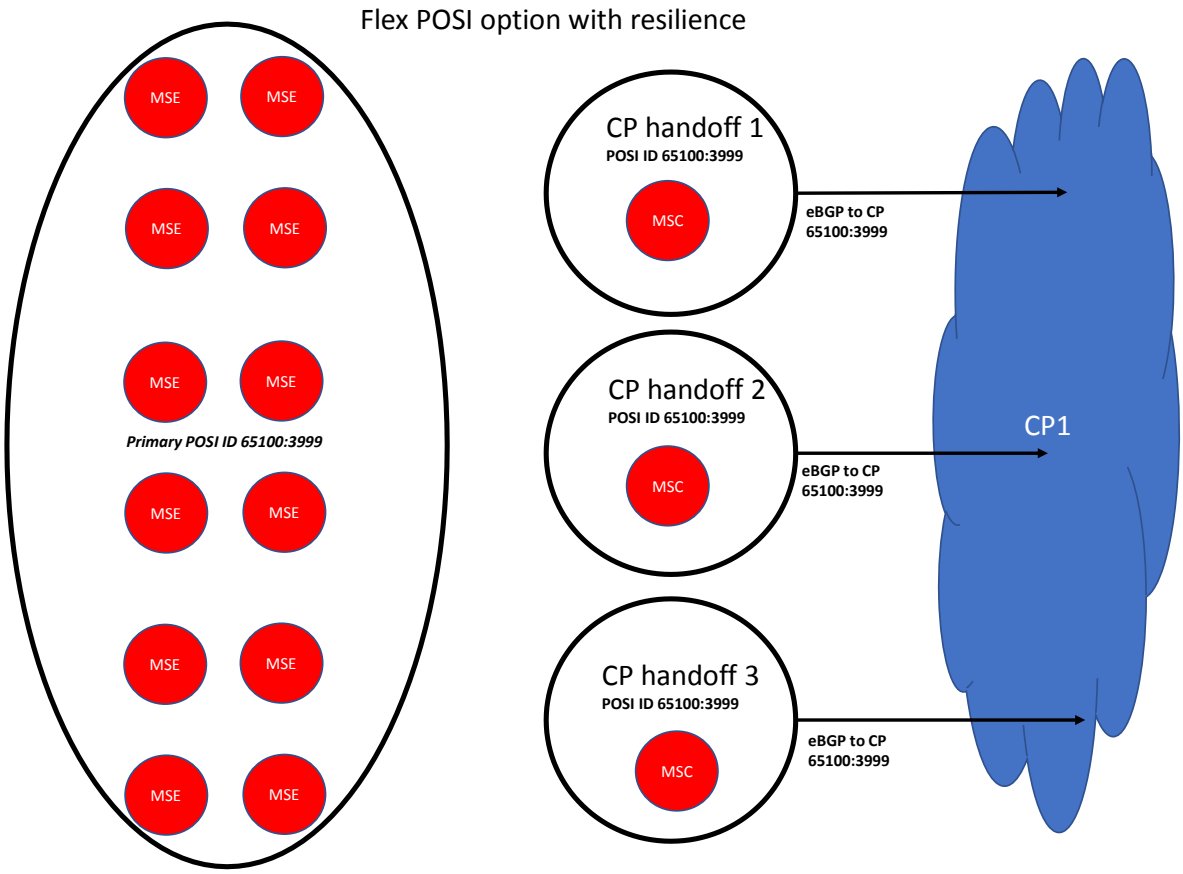

In this configuration, all MSE are will receive a default route and more specific prefixes that are advertised by the CP. All of the MSCs will see all MSE routes and advertise these to the CP. Traffic will choose the best route across the VPN. This is particularly well suited to CPs that use L2TP

Flex POSI with primary/secondary MSC option

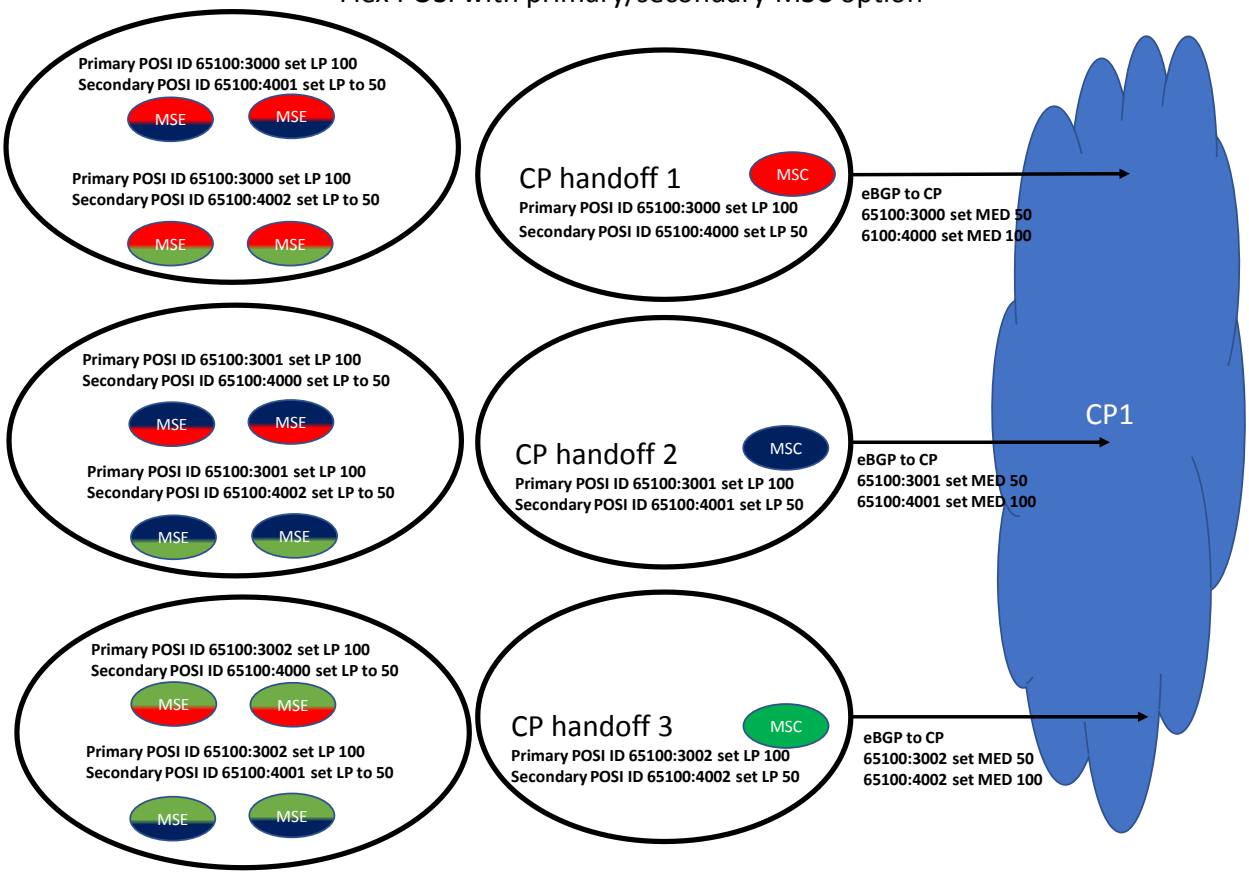

This configuration represents an example of the primary/secondary POSI resilience approach:

- Each group of MSE have a primary POSI. That group of MSE have distributed secondary POSI for use under failure of the primary POSI.
- Upstream end user traffic prefers higher local preference.
- Downstream end user traffic from CP prefers lower MED.

Flex POSI interconnect details

Flex POSI will be supported on the existing 1G and 10G MSIL based interconnect options and additionally on 2 new “Direct EP” options where connectivity is provided directly onto the MSC device performing the BEA function offered at 10G and 100G speeds.

For the direct EP options, it is expected that the CP will connect to both BEAs at a given POSI. The physical interface will be configured in a LAG bundle to make growth easier All interfaces within a bundle must be of the same physical access type i.e. all 10GE or all 100GE.

The interface will be an Ethernet interface that is set to carry 802.1Q tagged frames. A single VLAN known as an extension path (EP) is carried over the physical interface, and this EP VLAN will support either IPv4 only (single stacked), or it will support both IPv4 and IPv6 (dual stacked).

A single stacked EP VLAN will be allocated an IPv4 /30 or /31 subnet. A dual stacked EP VLAN will be allocated an IPv4 /30 or /31 subnet and an IPv6 /127 subnet.

Routing information must be exchanged between the BEA and the CP. The BEA needs to advertise end user addresses and L2TP IPv4 endpoints, and the CP must send a default route. .The CP will be allowed by default to send up to 1,000 prefixes per eBGP peering. The routing protocol used between the BEA and the CP is eBGP for both IPv4 and IPv6.

Each VLAN will require a BGP peering. The peering will be established between the VLAN IP addresses. There will be an IPv4 eBGP peer for single stacked VLANs, and for dual stacked VLANs there will be an IPv4 eBGP peer and an IPv6 eBGP peer per VLAN. The AS number used by WBC is 65100.

eBGP peering details:

- CP to send default route.
- CP allowed to send up to 1,000 prefixes.
- MSC to advertise the BRAS IPv4 L2TP loopback.
- MSC to advertise the BRAS dynamic pool addresses.
- MSC to advertise the BRAS static/non NAT IP addresses.
- MSC to advertise the IPv6 delegated prefix (pool or static address)
-

It will be optional for a CP to request BFD for the eBGP peer for direct handoff connections. EPs delivered via the MSIL will not support BFD.

All PTA traffic will be sent to the CP. The CP can then allow the traffic to route back to the BEA and to any end user if required.

Only PTA traffic will be policy base routed, and L2TP traffic should be normally routed.

ANNEX C – TV Connect

Introduction

This annex provides a brief service description for Wholesale Broadband Connect customers who wish to utilise the TV Connect option.

TV Connect is an optional feature of BT's Wholesale Broadband Connect (WBC) service. It provides live streaming of TV channels to End Users on WBC for an ISP. Fully managed TV channels are delivered from a Head End to End Users by using multicast technology within the 21C core.

TV Connect is not available for the IPstream Connect service.

Note: TV Connect will use IPv4 addressing; it will not use IPv6 addressing.

TV Connect is illustrated in Figure 2 below. For further technical information about TV Connect, please refer to the TV Connect SIN 511 available at http://www.btplc.com/sinet/

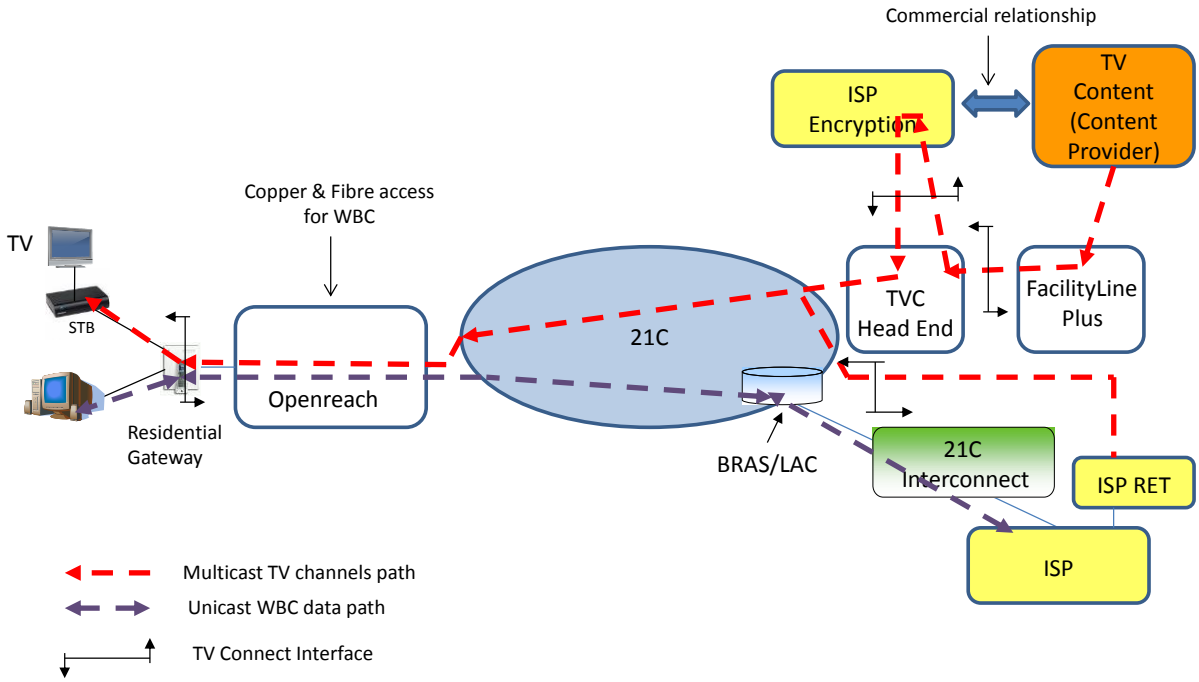

Figure 2. TV Connect service architecture

Onboarding for TV Connect

An ISP is introduced to TV Connect using the onboarding process. During the onboarding process, BT will be advised of the TV channel content providers with which the ISP has a commercial relationship.

ISP programme sources for TV Connect will be presented to the Head End via FacilityLine Plus.

The standard BT Wholesale FacilityLine Plus product will be used to deliver programme sources to the Head End. The FacilityLine product falls outside the scope of TV Connect. Technical information on FacilityLine Plus is available from BT Wholesale at http://www.btwholesale.com.

The ISP may use connectivity methods other than a Facility Line or Facility Line Plus.

The ISP must advise BT of its intended connectivity method prior to order placement. This will be assessed by BT on a case by case basis and BT will advise the ISP if the method requested can be supported or not.

The TV Connect Head End is designed to accept dual feeds of each ISP programme source from the ISP in order to ensure high availability of TV content.

The ISP will supply their own method of encrypting channels to support conditional access (CA).

TV Connect Channel Types

Standard Definition and High Definition channel types are supported with a range of options. Please consult the TV Connect SIN 511 available at http://www.btplc.com/sinet/

Impact on existing WBC components

WBC End User Access Component

The TV Connect channels are delivered over a separate path from the broadband data path. As a consequence there are the following impacts on the End User Residential Gateway (home router) device.

For an End User connected via Openreach GEA, the Residential Gateway will support:

– IPoE for multicast traffic and PPPoE for broadband traffic

– VLAN tag ID of 0 or no VLAN tag ID for multicast and broadband traffic

– Fork IGMP requests up multicast and broadband paths

For an End User conncted via an MSAN (WBC Copper), the Residential Gateway will support

– A dual VC architecture with TV Connect traffic on ATM VP/VC 0/35 with IPoE and broadband traffic on ATM VP/VC 0/38 with PPPoE

– Fork IGMP requests up multicast and broadband paths

The End User Residential Gateway will act as an IGMP proxy routing agent, support IGMP snooping and will fork IGMP messages up the TV Connect and broadband paths.

For further information, please refer to the TV Connect SIN 511 available at.
http://www.btplc.com/sinet/

QOS in relation to WBC

For an End User connected via Openreach GEA, an IPoE path is used for multicast traffic and a PPPoE path for broadband traffic.

For an End User connected via WBC copper an IPoE path is used for multicast traffic and a PPPoE path for broadband traffic.

In both cases IGMP requests are forked up the multicast and broadband paths. This feature enables the BRAS (or LNS for an ISP taking the L2TP handover option for WBC) to shape the broadband traffic to a rate equal to the line rate minus the sum of the multicast traffic.

Broadband traffic will be prioritised at the BRAS in accordance with the traffic markings as outlined earlier in this SIN [sections 5.6 to 5.9].

All TV Connect traffic will be given higher priority than all data traffic from the GEA CableLink to the ONT / DSLAM and similarly for the MSAN. TV Connect traffic is given a (.1p) priority marking of 3.

Test & Diagnostics

TV Connect has its own suite of test & diagnostics tools specifically developed to analyse, diagnose and ensure the delivery of TV channels across the network.

<END>