VectaStar 2

TECHNICAL SPECIFICATION

Cambridge Broadband Networks Limited
D001058 Technical Specification Version 3

June 2011

Details are correct at time of publication and are subject to change without notice

CONFIDENTIAL: PDF TO BE RELEASED UNDER NDA ONLY

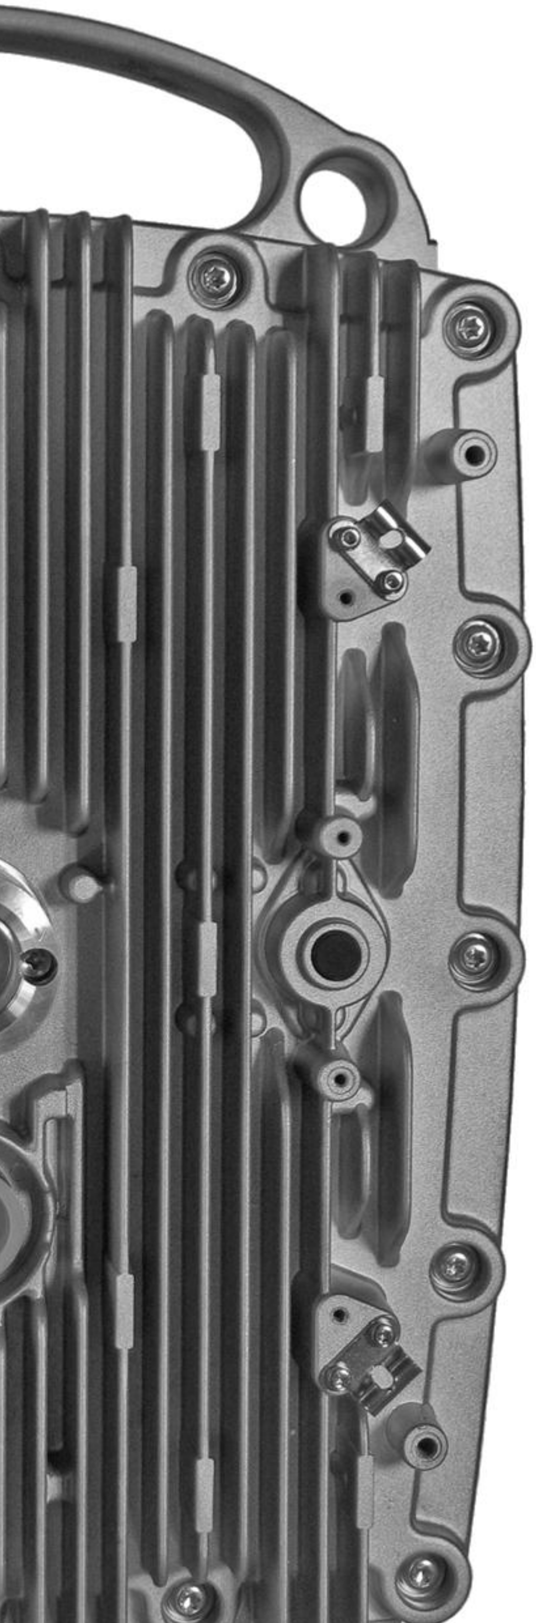

CONTENTS

VECTASTAR OVERVIEW

Cambridge Broadband Networks' VectaStar is a highly flexible carrier-class Point-to-MultiPoint transmission platform, offering a compelling business case advantage over Point-to-Point radio and/or leased lines. With its unique mix of features, VectaStar offers a competitive solution for high capacity business access and cellular backhaul networks.

Operating in the standard ETSI 10.5GHz, 26GHz and 28GHz frequency bands VectaStar can achieve throughputs of over 175Mbps (Gross) in a 28MHz channel. A single platform seamlessly supports Ethernet, TDM and ATM backhaul and has been proven to support the migration from TDM / ATM to an all IP backhaul. VectaStar helps operators minimise their investment in backhaul technology, reduce operational and network management expenditure, and provides the extra benefit of supporting additional revenue generating services.

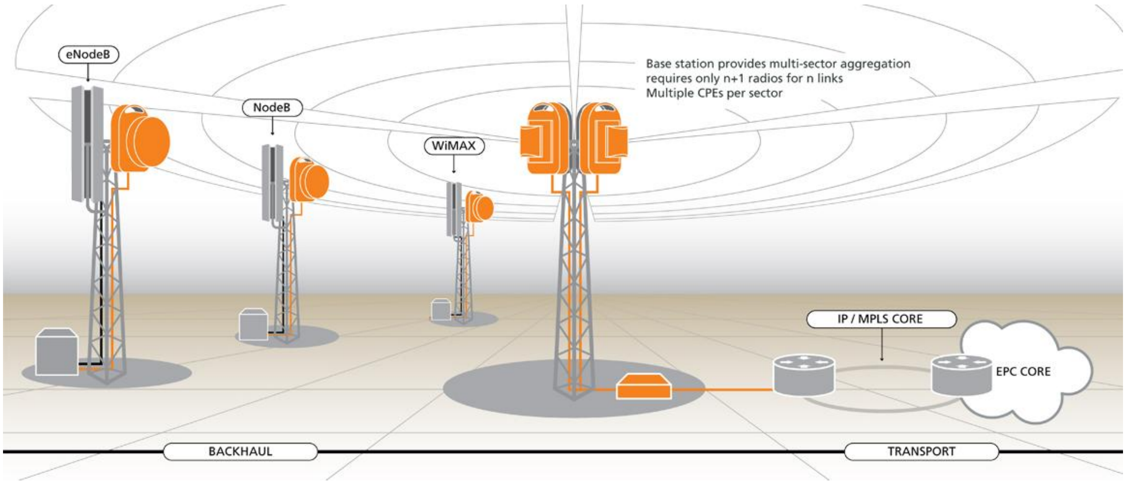

FIGURE 1: VECTASTAR NETWORK OVERVIEW

KEY FEATURES

- Over 175Mbps Gross throughput (350Mbps full duplex) per sector (28MHz channel)
- Class-leading spectral efficiency of 6.2 bits/second/Hz (Gross)
- 7-state Hitless Adaptive Modulation (ACM) from QPSK up to 256QAM with Trellis Code Modulation (TCM)
- Full Quality of Service (QoS) support for IP/Ethernet, E1 and TDM services
- Sub-millisecond latency for delay-sensitive applications
- Built-in Statistical Multiplexing and Traffic Aggregation
- Integrated optimisation for TDM (E1), 2G (Abis), 3G (Iub) and IMA services
- Flexible synchronisation schemes
- Fully redundant Base Station (1+1)
- Fully symmetric radio links (power, modulation & bandwidth)

VECTASTAR DESCRIPTION

SYSTEM OVERVIEW

VectaStar is the world's most advanced Point-to-MultiPoint (PMP) microwave radio system designed and optimised for backhaul of high capacity services. VectaStar is used around the world to backhaul 2G, 3G, HSPA and WiMAX networks as well as to provide carrier-grade corporate access.

At the hub of a VectaStar network is the base station comprising multiple sector radios (Access Points) connected to the indoor equipment which provides management, service termination, protocol conversion, monitoring, alarms and network interfaces. Typically, Access Points (AP) are 90-degree sectors so four APs are used for full 360 degree coverage. The diagram below shows the VectaStar architecture (for simplicity only one AP is shown):

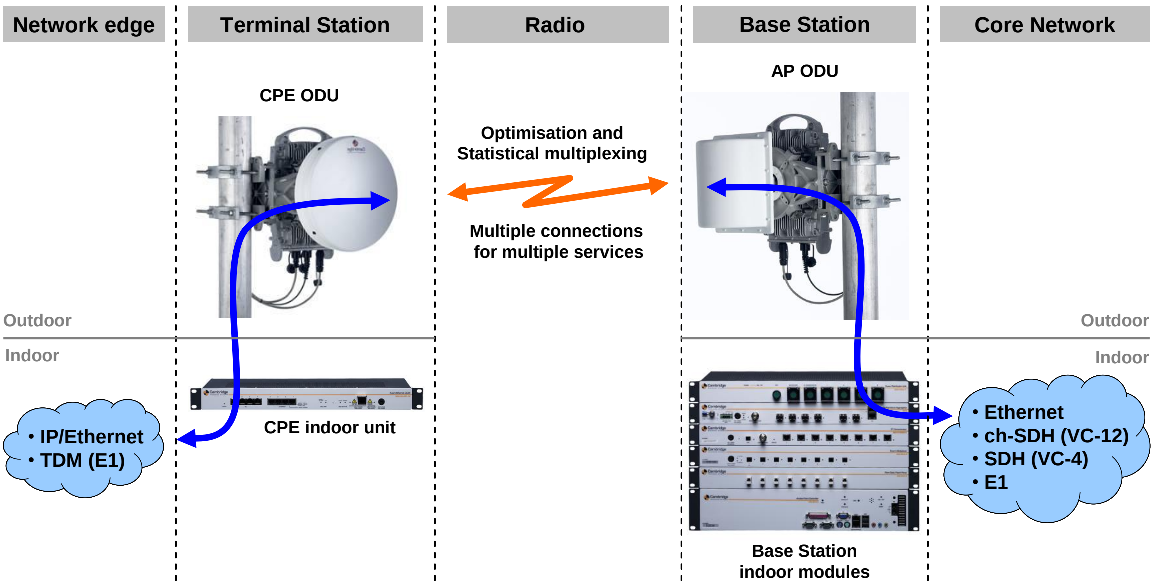

FIGURE 2: VECTASTAR SYSTEM

On the left, the CPE connects to the terminal station equipment, e.g. BTS, NodeB or WiMAX Base Station by means of either Ethernet or E1 interfaces. On the right, the VectaStar base station aggregates multiple sectors' traffic before terminating it in the core network.

VectaStar networks are managed using the VectaStar Network Management System (VNMS), which provides a graphical based environment for operators and service providers to configure, manage and monitor their VectaStar network. The VNMS provides a suite of applications which cover Fault, Configuration, Accounting, Performance and Security (FCAPS).

BASE STATION EQUIPMENT

The VectaStar base station terminates IP/Ethernet, E1 and ATM traffic, providing simultaneous support for both business access and cellular backhaul applications.

A typical VectaStar base station comprises indoor Base Station modules and multiple sector radios. The sector radio, or Access Point (AP), is an externally mounted antenna, radio, modem and network interface. Indoor equipment has been designed to be radio band-agnostic with the radio band of each sector determined by the AP. Therefore, a single Base Station can support APs on different frequency bands, e.g. one sector providing 26 GHz radio coverage and another providing 28 GHz or 10.5 GHz.

FIGURE 3: VECTASTAR BASE STATION INDOOR EQUIPMENT

VectaStar Base Station indoor units are industry standard 19" rack mountable modules designed for installation in weather-protected telecommunication equipment cabinets. The indoor equipment comprises a minimum of:

- Access Point Controller (APC)
- Multiplexer (MUX)
- fibre optic patch panel
- Multi-Protocol Aggregator (MPA), which provides comprehensive connectivity to different backhaul networks such as Gigabit Ethernet, channelised SDH, etc.

Optional indoor units are:

- Power Distribution Units (PDU) for distributing -48V to the APs and Base Station modules
- E1 Concentrators (not shown above) for terminating E1 services at the Base Station

The VectaStar Base Station is truly scalable due to its modular design, which allows additional modules (e.g. APs, E1 Concentrators, MPAs, etc.) to be added when increased capacity or coverage is required. Each AP aggregates traffic from the CPEs in its sector and the VectaStar indoor equipment then aggregates traffic from all sectors without the need for additional aggregation or optimisation equipment.

VECTASTAR2 ACCESS POINTS – 26 AND 28 GHZ

The outdoor radio units located at a VectaStar Base Station are called Access Points (AP) and comprise antenna, radio, modem, Medium Access Control (MAC), baseband board and interfaces for data and power. Each 26 or 28 GHz AP provides radio coverage for one sector and can deliver over 130Mb/s Ethernet throughput to a maximum of 30 CPEs per sector. A Base Station can support up to a total of five APs (redundant or non-redundant), with each AP operating in the same or different frequency bands e.g. 10.5 GHz, 26 GHz or 28 GHz.

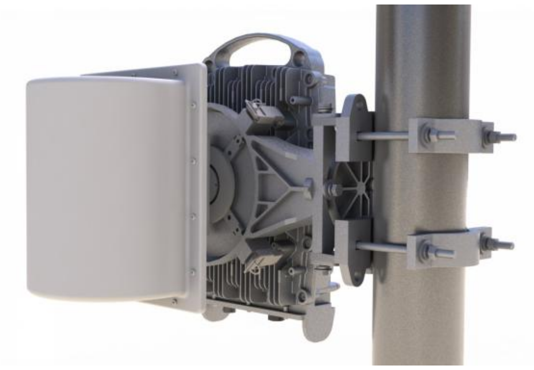

FIGURE 4: FRONT ¾ VIEW OF 26/28 GHZ AP

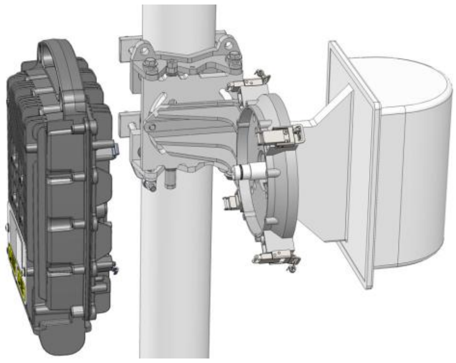

FIGURE 5: SLIP-FIT DESIGN OF 26/28 GHZ AP

Operators can benefit from the 1.5 km extended-reach capability of the multi-mode fibre interconnect, which enables the use of existing fibre optic cable infrastructure to provide the maximum reuse of existing sites. This capability removes the need for environmentally controlled rooms near the antenna mast or tower, which adds additional flexibility for the acquisition or leasing of antenna locations. It also allows the sectors to be positioned in a distributed manner, eliminating potential shadows and facilitating the rapid and low-cost deployment of new antenna masts, using small footprints on building roofs.

26 and 28 GHz VectaStar2 AP Technical summary	
Data Interfaces	1 x STM-1 LC Multi Mode (to controller)
Redundancy	Equipment-based with sub-5 second failover to paired AP
MTBF	>50 years
Power input	–48 VDC
Power consumption	35 W
Dimensions	26 GHz: 431 x 266 x 134 mm (Height x Width x Depth) 28 GHz: 431 x 266 x 139 mm (Height x Width x Depth)
Weight	8.5 kg

VECTASTAR2 ACCESS POINTS – 10.5 GHZ

The outdoor radio units located at a VectaStar Base Station are called Access Points (AP) and comprise antenna, radio, modem, Medium Access Control, baseband board and interfaces for data and power. Each 10.5 GHz AP provides radio coverage for one sector and can deliver over 65Mb/s Ethernet throughput to a maximum of 30 CPEs per sector. A Base Station can have multiple APs, with each AP operating in the same or different bands e.g. 10.5 GHz, 26 GHz or 28 GHz.

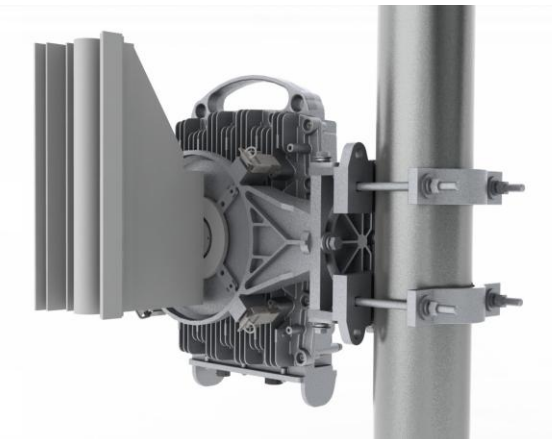

FIGURE 6: 10.5 GHZ AP WITH SLIP-FIT VERTICAL ANTENNA

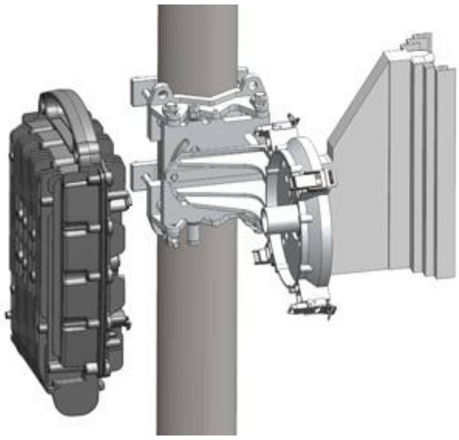

FIGURE 7: 10.5GHZ AP WITH VERTICAL ANTENNA EXPLODED VIEW

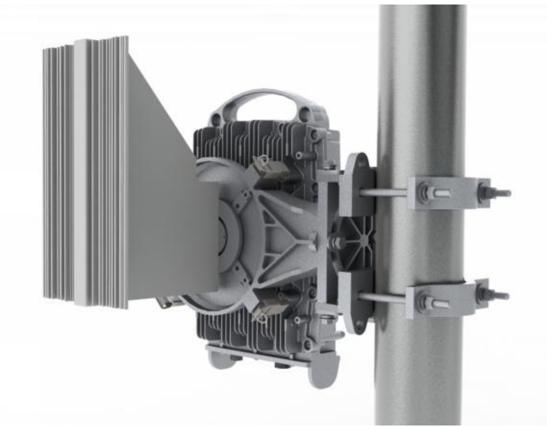

FIGURE 8: 10.5 GHZ AP WITH SLIP-FIT HORIZONTAL ANTENNA

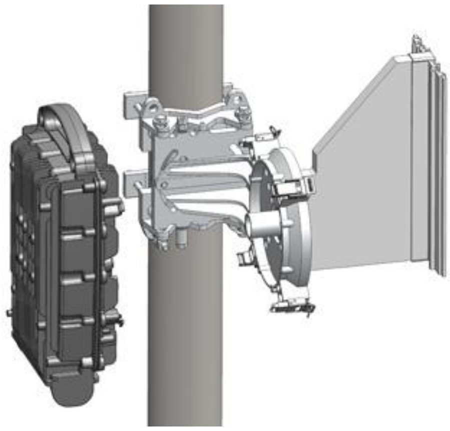

FIGURE 9: 10.5GHZ AP WITH HORIZONTAL ANTENNA EXPLODED VIEW

Operators can benefit from the 1.5 km extended-reach capability of the multi-mode fibre interconnect, which enables the use of existing fibre optic cable infrastructure to provide the maximum reuse of existing sites. This capability removes the need for environmentally controlled rooms near the antenna mast or tower, which adds additional flexibility for the acquisition or leasing of antenna locations. It also allows the sectors to be positioned in a distributed manner, eliminating potential shadows and facilitating the rapid and low-cost deployment of new antenna masts, using small footprints on building roofs.

10.5GHz VectaStar2 AP Technical summary	
Data Interfaces	1 x STM-1 LC Multi Mode (to controller)
Redundancy	Equipment-based with sub-5 second failover to paired AP
MTBF	>50 years
Power input	–48 VDC
Power consumption	35 W
Dimensions	431 x 266 x 127mm (Height x Width x Depth)
Weight	8.5 kg

Cambridge BROADBAND NETWORKS

ACCESS POINT CONTROLLER

The Access Point Controller (APC) is a dedicated embedded computer running a High Availability Linux operating system on Flash storage for optimum reliability. A hard disk drive is used for logging purposes. The APC connects to the Multiplexer (BSC-MUX-8) to manage and configure all connected Access Points and E1 concentrators. The unit is suitable for installation in weather-protected telecommunication equipment cabinets.

APC Technical summary	
Data Interfaces	1 x STM-1 (OC-3) SC Multi Mode
	2 x Fast Ethernet (one for Management, one for Redundancy connection)
	1 x USB for factory reset
	2 x Serial connection
Redundancy	Equipment-based with hitless failover to paired APC
	Fast Ethernet connection to paired APC
MTBF	>50 years
Power input	–48 VDC
Power consumption	35 W
Dimensions	90 x 481 x 260 mm (2U 19")
Weight	5.5 kg

MULTI-PROTOCOL AGGREGATOR

The Multi-Protocol Aggregator (MPA) provides a range of high capacity interfaces and protocol conversions, enabling a VectaStar network to interface seamlessly with both legacy SDH networks and Gigabit Ethernet Networks. Both channelised (VC-12) and un-channelised (VC-4) SDH networks are supported.

The MPA supports VectaStar optimisation technology, enabling operators to maximise the efficiency of their backhaul networks by carrying optimised Abis, IuB and S1 interfaces across their SDH transport networks. Multiple MPAs can be installed at a VectaStar Hub site providing a scalable capacity solution. Typically, MPAs are installed in pairs as part of a Protected Solution, utilising MSP 1:1 for interconnectivity protection.

Small Form Pluggable (SFP) connectors allow for support of the full range of optical and electrical interfaces and remove the need for Single Mode to Multi Mode fibre converters.

Management is via the VectaStar NMS suite and the MPA seamlessly integrates into a VectaStar network allowing full FCAPS network management functionality.

MPA Technical summary	
Data Interfaces	4 x STM-1 SFP 2 x Gig Ethernet SFP 2 x 10/100BaseT Ethernet
Synchronisation options	External station clock (G703) or STM-1 frame (channelised or unchannelised)
Fail-over	Clock fail over between redundant sources with user-configurable clock reversion
Redundancy	Both Port based (single MPA) or Equipment based (Redundant MPA) STM-1 interfaces: MSP 1:1 Ethernet: Loss of link Protection
MTBF	>50 years
Power	–48 VDC
Power consumption	20 W
Dimensions	45 x 481 x 244 mm (1U 19")
Weight	3.0 kg

MPA Service terminations	
Ethernet	ATM Interworking Function (IWF) : • RFC1483 (MPOA)
VC-12 (E1)	ATM IWF: • Circuit Emulation Service (structured and unstructured) • Abis Optimisation • IuB (G.804) Optimisation
ATM	VC-4

MPA Performance	
Ethernet	VLAN filtering 802.1pq QinQ support 1000 VLAN terminations 450 Mbps at 64 byte packet throughput non-redundant or 135 Mbps redundant
VC-12 (E1)	Cross connect up to 63 CES ATM VCs to VC-12 container
ATM	Line rate (STM-1) cross connect
Optimisation	Abis: Up to 63 E1s (VC-12) g804: Up to 63 E1s (VC-12)

MULTIPLEXER

The Multiplexer (MUX) is a low cost 8 port STM-1 switch. It combines management and network traffic to and from the network interfaces and the VectaStar equipment attached to its ports. Additional interface modules can be connected to the Multiplexer to increase the functionality e.g. additional 8 E1 Concentrators can be added. Multiplexers can be concatenated (daisy-chained) up to 3 tiers.

MUX Technical summary	
Data Interfaces	8 x STM-1 (unchannelised only) Multi Mode MTRJ
Synchronisation options	STM-1 frame
Fail-over	Clock fail over with automatic clock reversion
Redundancy	Both port-based redundancy (on a single MUX) or equipment-based redundancy as part of a fully 1+1 Base Station STM-1 interfaces: MSP 1:1
MTBF	>50 years
Power input	–48 VDC
Power consumption	20 W
Dimensions	45 x 481 x 244 mm (1U 19")
Weight	3.0 kg

POWER DISTRIBUTION UNIT

The PDU is a dedicated Power Distribution Unit for base station control equipment. The PDU takes -48VDC from a suitable communications equipment power source and provides independently switched, fused and surge protected -48VDC outputs suitable for base station control equipment.

Power Distribution Unit Technical Summary	
Interfaces	Front panel: 1 x APC indicator, 6 x press switches Rear panel: 1 x input, 7 x output Neutrik sockets
MTBF	>50 years
Power	-48VDC
Power consumption	10 W
Dimensions	45 x 481 x 244 mm (1U 19")
Weight	3.0 kg

E1 CONCENTRATOR

The E1 Concentrator is an optional module, which connects to the Multiplexer and provides eight RJ48 E1 interfaces, supporting Circuit Emulation Service (CES), Inverse Multiplexing for ATM (IMA). The E1 Concentrator also supports advanced optimisation algorithms for 2G (Abis), 3G (Iub), TDM (G704), IMA and CDMA services, allowing substantial bandwidth savings and advanced dynamic bandwidth allocation / statistical multiplexing of E1 services.

E1 Concentrator Technical summary	
Data Interfaces	1 x STM-1 Multi Mode MTRJ for connection to a Multiplexer 1 x BNC for 2.048MHz reference clock input/output 8 x RJ48 balanced 120 Ohm differential pairs for E1/IMA
Synchronisation options	STM-1 frame from Multiplexer 2.048MHz reference clock from External source Any one of the eight RJ48 ports
Fail-over	Clock fail over with automatic clock reversion
MTBF	>50 years
Power input	–48 VDC
Power consumption	5 W
Dimensions	45 x 481 x 244mm (1U 19")
Weight	3.0 kg

FIBRE OPTIC PATCH PANEL

The Fibre Optic Patch Panel is a passive unit that provides a termination for the outdoor grade industrial fibres, from the Access Points, and provides an indoor fibre patch panel to connect network equipment together.

Fibre Optic Patch Panel Technical Summary	
Interfaces	Front panel: 8 x MTRJ Multi Mode Rear panel: 8 x SC Multi Mode
MTBF	>50 years
Power	Not applicable
Dimensions	45 x 481 x 244 mm (1U 19")
Weight	3.0 kg

CUSTOMER PREMISES EQUIPMENT (CPE)

CPE ODU – 26 AND 28 GHZ

The VectaStar CPE comprises an Outdoor Unit (ODU) and associated indoor equipment that connects to the customer building/cell site. The CPE ODU provides a high gain radio link to the Access Point at the Base Station and can deliver up to line rate Fast Ethernet for any packet size. The CPE ODU comprises radio, modem, MAC, digital baseband board and interfaces for data and power.

FIGURE 10: 26/28 GHZ CPE ODU WITH SLIP-FIT 30CM (1-FOOT) PARABOLIC ANTENNA

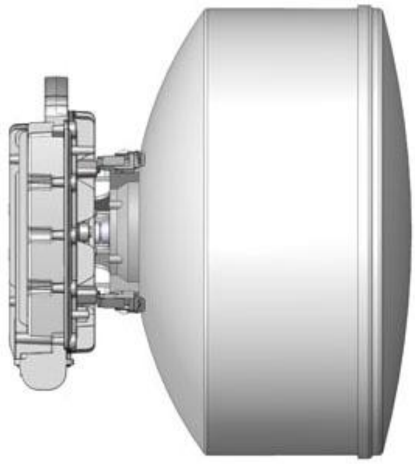

FIGURE 11: 26/28 GHZ CPE ODU WITH SLIP-FIT 60CM (2-FOOT) PARABOLIC ANTENNA

Antennas are linearly polarised and there are different sizes available for each frequency variant. A larger antenna will increase the link budget for both the downlink and uplink. The CPE antenna can be rotated in the field to move between Vertical and Horizontal polarisation.

VectaStar is the only PMP solution that supplies an “outdoor only” variant for the CPE. This is due to the VectaStar CPE design that puts all the intelligence in the ODU along with an Ethernet interface for localised Ethernet connection. Therefore, if only Ethernet is required at the CPE site e.g. for an Ethernet NodeB, then VectaStar provides the simplest network solution since all that is required is to power the CPE ODU with -48V DC and connect the Ethernet output from the ODU to the end equipment.

The connection between the Outdoor and optional Indoor units is a single CAT-5E cable carrying power and data (Ethernet).

26/28 GHz CPE Technical summary	
Data Interfaces	1 x RJ-45 (carrying 10/100Base-T Ethernet and -48 VDC)
MTBF	>50 years
Power input	–48 VDC (coupled with Ethernet on CAT-5E from LP-CODU, WB or IDU)
Power consumption	35 W
Dimensions, weight	431 x 266 x 134 mm (Height x Width x Depth), 8.5 kg

CPE ODU – 10.5 GHZ

The VectaStar CPE comprises an Outdoor Unit (ODU) and associated indoor equipment that connects to the customer building/cell site. The CPE ODU provides a high gain radio link to the Access Point at the Base Station and can deliver up to line rate Fast Ethernet for any packet size. The CPE ODU comprises radio, modem, MAC, digital baseband board and interfaces for data and power.

FIGURE 12: 10.5 GHZ CPE ODU WITH SLIP-FIT 30CM (1-FOOT) PARABOLIC ANTENNA

FIGURE 13: 10.5 GHZ CPE ODU WITH SLIP-FIT 60CM (2-FOOT) PARABOLIC ANTENNA

Antennas are linearly polarised and there are different sizes available for each frequency variant. A larger antenna will increase the link budget for both the downlink and uplink. The CPE antenna can be rotated in the field to move between Vertical and Horizontal polarisation.

VectaStar is the only PMP solution that supplies an “outdoor only” variant for the CPE. This is due to the VectaStar CPE design that puts all the intelligence in the ODU along with an Ethernet interface for localised Ethernet connection. Therefore, if only Ethernet is required at the CPE site e.g. for an Ethernet NodeB, then VectaStar provides the simplest network solution since all that is required is to power the CPE ODU with -48V DC and connect the Ethernet output from the ODU to the end equipment.

The connection between the Outdoor and optional Indoor units is a single CAT-5E cable carrying power and data (Ethernet).

10.5 GHz CPE Technical summary	
Data Interfaces	1 x RJ-45 (carrying 10/100Base-T Ethernet and -48 VDC)
MTBF	>50 years
Power input	–48 VDC (coupled with Ethernet on CAT-5E from LP-CODU, WB or IDU)
Power consumption	35 W
Dimensions, weight	431 x 266 x 127mm (Height x Width x Depth), 8.5 kg

CPE INDOOR UNIT (IDU)

All CPEs can be enhanced to provide more user interfaces by the addition of an indoor unit. The ODUs connect to the indoor units by means of CAT-5E cable, which carries -48V DC power and Fast Ethernet (100BaseT). All indoor units serve to couple the data and power onto the CAT-5E cable and include at least one Fast Ethernet interface.

There are two form factors for the indoor unit:

- The indoor Wallbox is a wall-mountable module with ODU status LED indicators and DC power (-48v) and Ethernet (RJ45) interfaces.
- The 1U 19" Indoor Unit (IDU) is a rack mountable version of the Wallbox, but also includes multiple Ethernet and E1 interfaces.

The range of IDUs is listed below:

FIGURE 14: CPE WALLBOX (IDU-WB-PSU)

FIGURE 15: CPE 4E1, 4VLAN 1U INDOOR UNIT (IDU-4VL-4E1)

IDU Summary	IDU-4VL-4E1	IDU-WB-PSU
Data Interfaces	4 x 100BaseT Ethernet 4 x RJ48 E1/IMA 120 Ω balanced differential pairs	1 x 100BaseT Ethernet
ODU interconnect	1 x RJ-45 (carrying 10/100Base-T Ethernet and -48 VDC)	
MTBF	>50 years	
Power input	–48 VDC or 100/240 VAC via optional PSU	
Power consumption	5 W	3 W
Dimensions	45 x 481 x 244 mm (1U 19")	80 x 90 x 30 mm
Weight	3.0 kg	0.5 kg

CABLE RUN LENGTHS

The maximum cable run length for a VectaStar2 CPE is 90m. Note: for installations where a Wallbox is used, this run length is from the CPE ODU to the end-user equipment since the Wallbox is not an Ethernet repeater. The Ethernet run length for the 1U IDU is from the ODU to the IDU as the IDU includes Ethernet switch functionality. Standard Ethernet over CAT-5E cable run limits apply from the IDU to the end-user equipment.

Cable runs in excess of 90 metres up to a maximum of 300 metres are supported through the use of an LP-CODU. See the Lightning Surge Protection section for more details.

LIGHTNING SURGE PROTECTION

VectaStar ODUs and Lightning Surge Protection units are designed for outdoor installation. Typically, outdoor radio units are mounted as high as possible on the mast/building to ensure Line of Sight to the other side of the radio link.

Typical microwave radio installations such as masts and tall buildings are susceptible to lightning strikes. Therefore, measures must be taken to minimise damage to any related equipment in the event of a strike. A proper installation will ensure that the lightning energy dissipates to ground through paths that do not include the VectaStar equipment.

CBNL recommends that the following levels of protection be used for VectaStar outdoor equipment:

- Deploy the outdoor equipment so as to avoid a direct strike. This typically means mounting the ODU inside the “cone of protection” as depicted in the diagram below
- Ground the mounting structure
- Ground the ODU
- Ground the cables on the mounting structure (using a lightning surge protection unit)
- Ground the cables on entry to the indoor enclosure (using a lightning surge protection unit)
- Ground the indoor equipment.

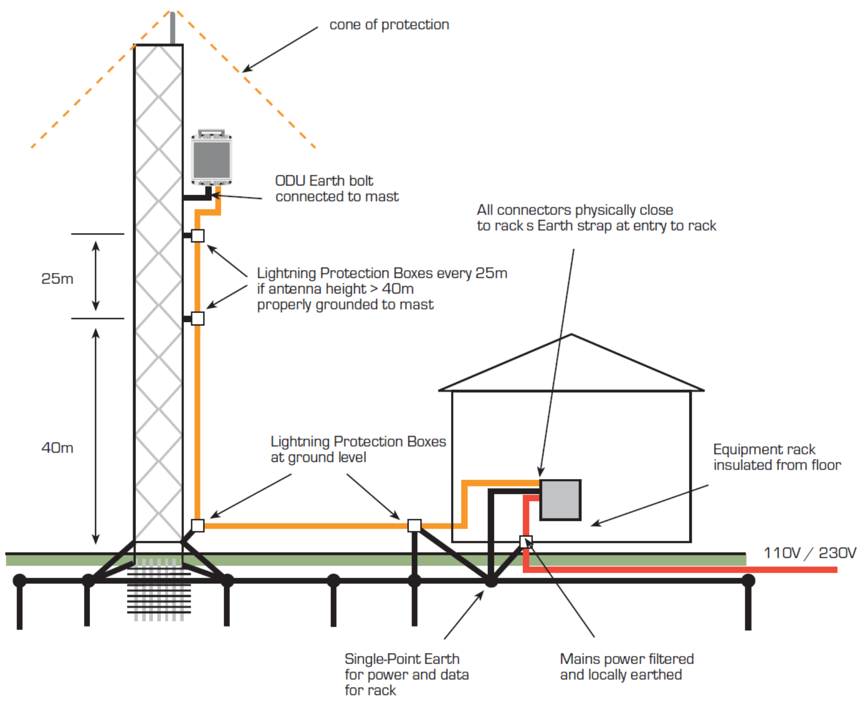

FIGURE 16: EXAMPLE OF TYPICAL LIGHTNING SURGE PROTECTION INSTALLATION

Where only one lightning surge protection unit is being installed in the cable run, this should be placed as near to the building ingress as possible.

Due to the different methods of connecting power and data to the Base Station APs versus the CPE ODUs, there are different types of Lightning Surge Protection boxes[1] for each type of ODU. This is explained below.

LIGHTNING SURGE PROTECTION FOR 2-CORE POWER (E.G. AP)

Base Station APs use 2-core cable for connecting power and fibre for making data connections. Fibre does not conduct lightning energy, therefore only the 2-core power cable run requires surge protection. This is provided in the form of the LP-AP pictured below:

FIGURE 17: LP-AP

LIGHTNING SURGE PROTECTION FOR CAT-5E CPE

CPEs use CAT-5E cable to connect the ODU from the indoor equipment. There are two types of lightning protection modules as detailed below:

- LP-CPE for CAT-5E only cable runs up to 80m.

FIGURE 18: LP-CPE

---

[1] CBNL's lightning surge protections modules provide protection to ITU-T Recommendation K.20, K.21

and BS EN 61000-4-5

- LP-CODU for CAT-5E runs up to 300m. This unit includes Ethernet repeater technology. This unit can also be used where CAT-5E and 2-core power are run in parallel. There is a maximum limit of three LP-CODUs per CPE.

FIGURE 19: LP-CODU

Configuration	Maximum cable length*
Wallbox (WB) or Rack mount IDU only	90m
WB or IDU + 1 LP-CPE	80m
WB or IDU + 2 LP-CPE	60m
WB or IDU + 1 LP-CODU	100m
WB or IDU + 2 LP-CODU	200m (needs 2 core power cable)
WB or IDU + 3 LP-CODU	300m (needs 2 core power cable)

* Cable lengths shown are the total lengths between indoor equipment (WB or IDU) and outdoor equipment (ODU). Note the distances for Wallbox installations have been calculated using a 10m CAT-5E cable to connect to 3rd party equipment

REDUNDANCY

VectaStar supports a number of different redundancy options at the Base Station.

- Full Base Station redundancy, including Access Points
- Indoor redundancy (APC and MPA)
- Outdoor redundancy (APs only)
- Gigabit and Fast Ethernet port redundancy (proprietary link-down protection)
- STM-1 backhaul port (MUX or MPA port) redundancy (sub-5 second 1:1 MSP)

These are detailed in the following sections.

FULL BASE STATION REDUNDANCY

For full Base Station redundancy, two full Base Stations are installed with all hardware paired. One set of Indoor Base Station equipment is connected to the other using redundancy cables as appropriate. This ensures that each hardware item can protect the failure of its partner independently.

The redundancy mechanism uses two nodes (separate, physical APCs) in a cluster. Each node has a unique hostname and management IP address. The cluster has a single IP address that is an alias to whichever node is currently running the VectaStar resource. The fully redundant base station is therefore managed using a single "Cluster" address which has control of the two full Base Stations.

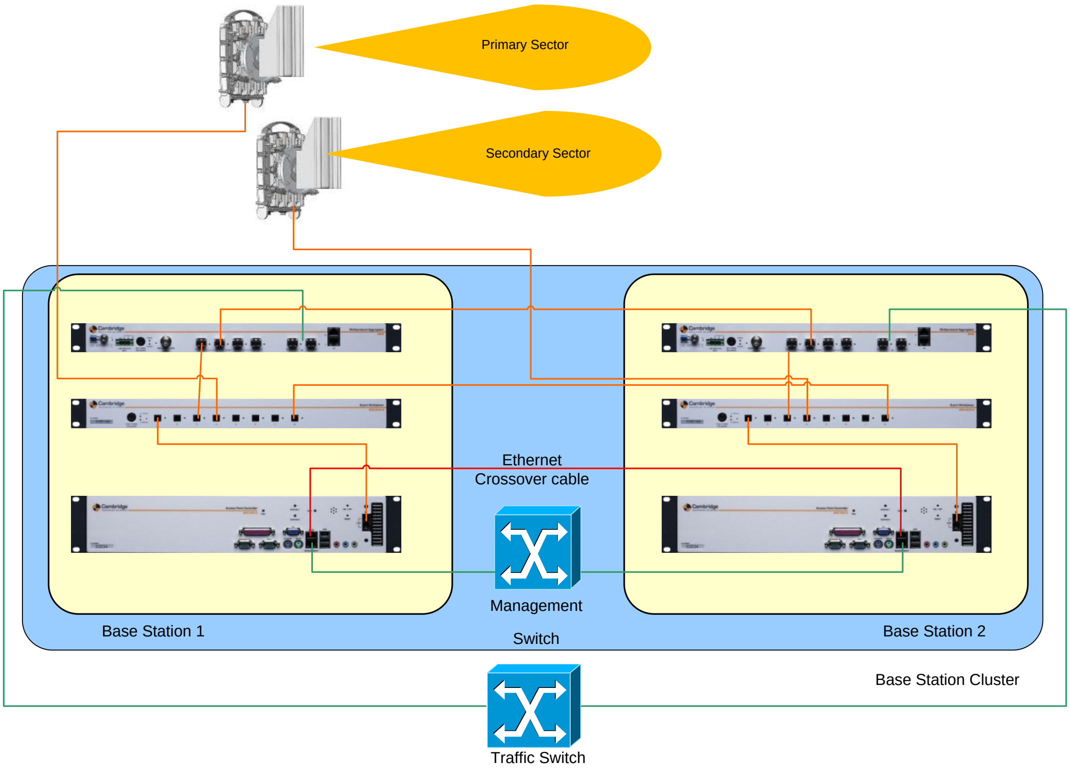

FIGURE 20: FULL BASE STATION REDUNDANCY – PDU AND FOPP NOT SHOWN FOR CLARITY

In the event of a Primary MUX failing, the entire Base Station will fail over to Secondary APs, the Secondary MUX and Secondary APC. All the CPEs will fail over to their allocated Secondary AP. This inherently provides backhaul STM-1 interface redundancy, as VCs on backhaul ports on the Primary MUX are switched to the same interface on the Secondary MUX. Note that backhaul port redundancy can also be provided via a single MUX. Fail over time is under five seconds.

OUTDOOR REDUNDANCY (ACCESS POINTS)

Outdoor redundancy uses AP pairs connected to a single set of indoor equipment. AP pairs (primary and secondary) are allocated in the VNMS. Fail over time is under five seconds. If the total number of APs exceeds the number of ports on a single MUX, a Virtual MUX can be used to increase the number of available ports.

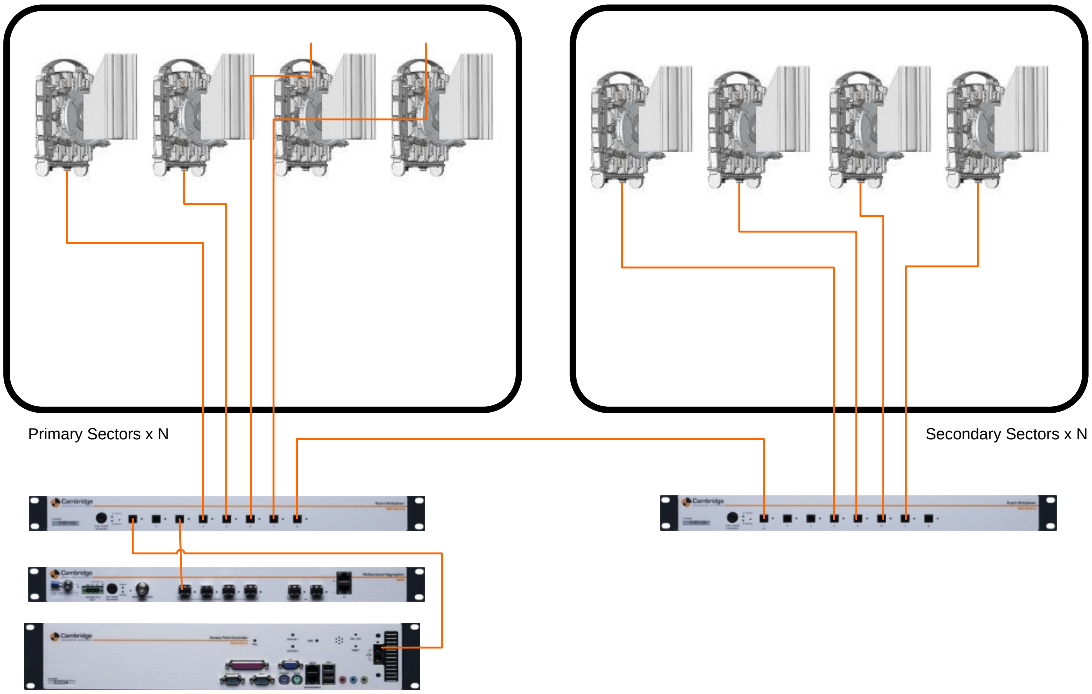

FIGURE 21: AP REDUNDANCY – PDU AND FOPP NOT SHOWN FOR CLARITY

GIGABIT AND FAST ETHERNET PORT REDUNDANCY

This version of redundancy protection uses a proprietary link-down protection. Fail over is under five seconds for VectaStar equipment.

STM-1 PORT BACKHAUL REDUNDANCY

On a Base Station with a single APC, it is possible to configure STM-1 interface redundancy. In the event that the Primary MUX port fails, any service VCs terminating on an interface of that MUX will be disconnected and switched to the nominated Secondary redundant port, with the VPI/VCI unchanged.

SERVICES

OVERVIEW

The VectaStar system is capable of carrying multiple protocol traffic from multiple sites back to a uses an optimised transport layer with different traffic streams carried over unique radio services from a CPE to the Base Station or another CPE. CPEs can support multiple services, each carried on separate virtual services. These services can be similar or different e.g. multiple vbridge services or a combination of E1 and vbridge services.

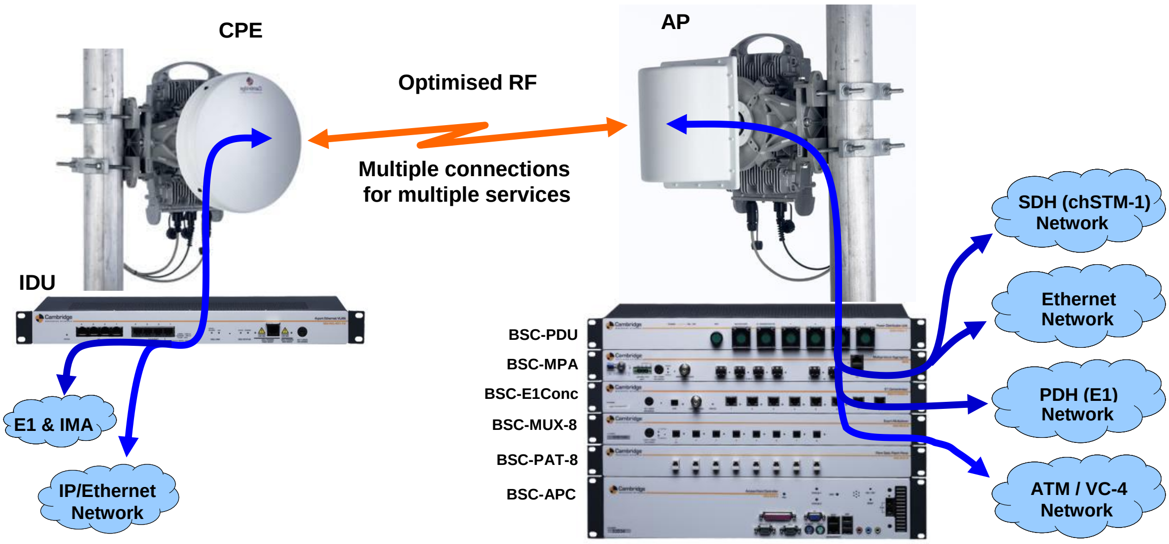

FIGURE 22: VECTASTAR SERVICE TERMINATIONS

At the Base Station, each service may be:

- Terminated at the Base Station in several ways:
  - switched into core network/routers as channelised SDH (VC-12) via the MPA
  - switched into the core network as unchannelised SDH (VC-4) via the MPA or MUX
  - terminated as IP/Ethernet on the MPA
  - terminated as E1 on an E1 Concentrator
- terminated on another CPE registered with the same Base Station (self-backhauled) in two ways:
  - terminated as IP/Ethernet on an IDU or Wallbox
  - terminated as E1 on an IDU

MODULATION PER SERVICE

Modulation is assigned per service. The modulation can be set independently on the Downlink and Uplink. So, different modulations can be in use on a single CPE as well as different CPEs within a sector.

VectaStar supports both static modulation assignment and adaptive modulation. See the **Radio Frequency Planning** chapter for more details on choosing modulation scheme.

SERVICE TYPES

The following service types are available:

Service	Description
vbridge	An RFC1483/2684 service carrying 802.1/802.1Q Ethertype frames which is part of an 802.1d bridge connected to the CPE's 802.1Q Ethernet interface
bridge	An RFC1483/2684 service carrying 802.1 Ethertype frames which is part of an 802.1d bridge connected to the CPE's 802.1 Ethernet interface
E1CES	A Circuit Emulation Service
E1CES (optimised)	ABIS optimised service for more efficient transport of 2G traffic, compared with CES
G.804	A service carrying all ATM cells extracted from a G.804 or IMA E1 interface
IMA	A service optimised for carrying Inverse Multiplexed ATM
Loopback	A VC that loops back in the CPE, used for test purposes

QOS

Quality of Service is defined on a per service basis (similar to modulation) and can be applied at the Radio layer and at either end e.g. with layer 2/layer 3 filtering available for IP/Ethernet services.

From a QoS point of view, there are two distinct types of services:

- Constant Bit Rate (CBR) services, which are E1-based
- non-CBR services, e.g. Ethernet and IP

QOS SUPPORT IN THE MAC

Within the AP's Medium Access Controller (MAC) there are four priority classes (numbered 0 – highest, to 3 – lowest) and each service is allocated to one class. Bandwidth is first offered to the highest priority class and the remaining unallocated bandwidth is then offered to the next priority class and so on. The highest priority class (0) is exclusively reserved for CBR traffic.

Each "virtual service" is queued independently with Peak Information Rate (PIR) limiting the maximum throughput per service. If no Committed Information Rate (CIR) values are set then the MAC implements a round robin scheduler in each priority class.

If CIR values are set, then the round robin scheduler is weighted such that services which are sending data below their CIR settings are scheduled ahead of those who have sent data, on average, above their CIR settings. Once all services have received the configured CIR the different CIR settings are used to weight the remaining bandwidth allocation towards each service up to their PIR limit.

The scheduling is applied using the formula below

$$T_{ST} = CIR_T + BW_{NAC} \times (CIR_S / \sum \text{of all CIRs})$$

Where:

T_{ST}	= Total Service Throughput
CIR_T	= Total CIR targets
BW_{NAC}	= Sector Bandwidth not Allocated to CIR (remaining available)
CIR_S	= CIR of this Service

A diagram showing an example of how the new scheduler works is shown below:

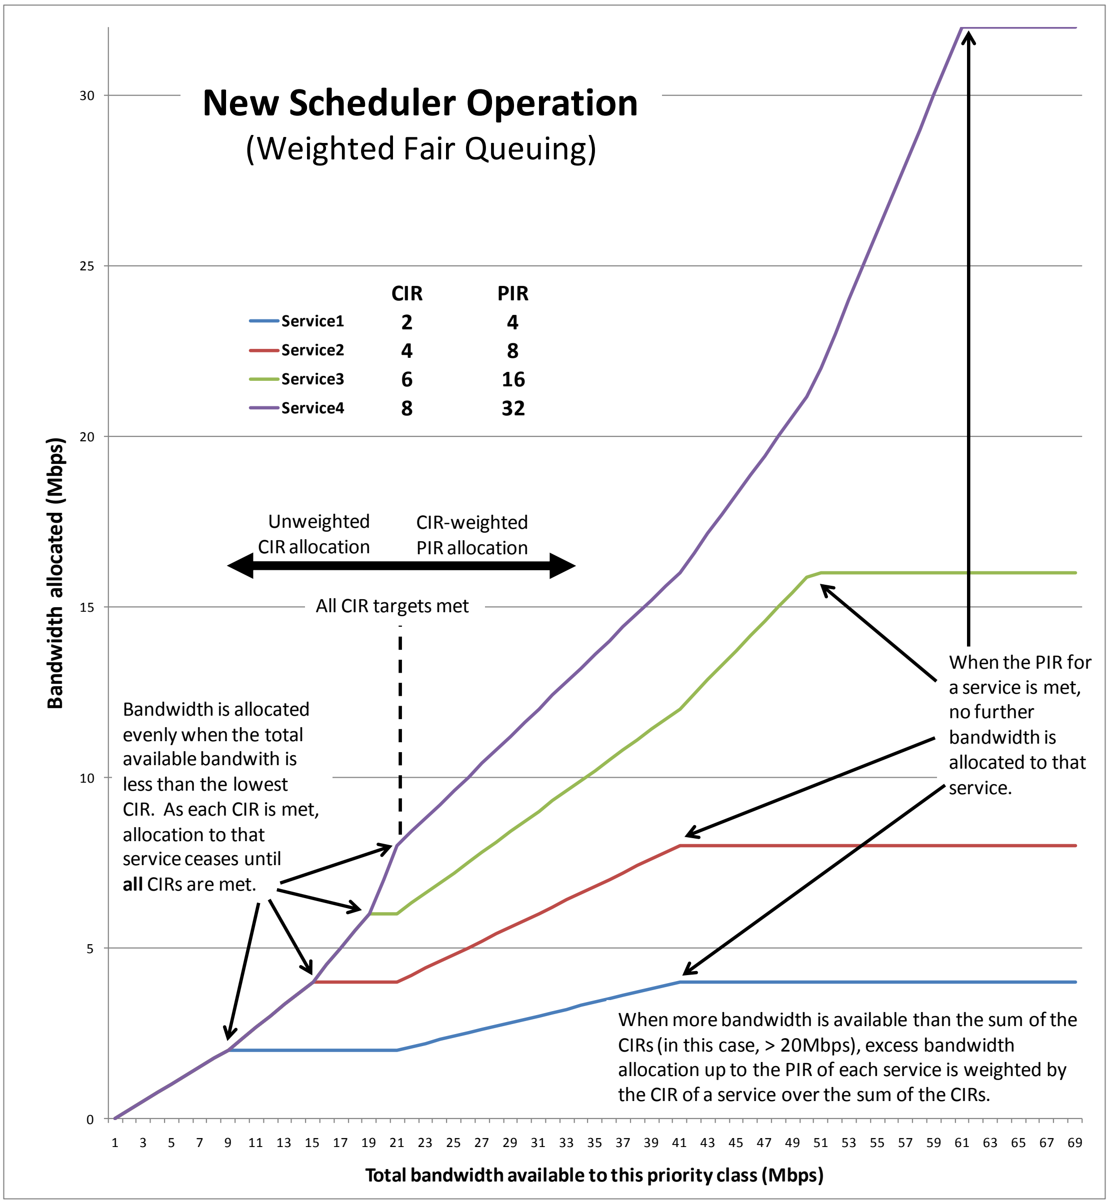

The downstream frame structure is Time Division Multiple Access (TDMA) with the AP's MAC scheduler creating each frame based on service demand and the relative QoS settings. There are two uplink access methods: reservation-based with initial contention is used for non-CBR services and a grant-based system is used for CBR services (to minimise Cell Delay Variation (CDV)).

QOS FOR CBR SERVICES

VectaStar encapsulates CBR services in AAL1 ATM VCs. The VectaStar MAC (Medium Access Controller) assigns these services to the highest priority class with no PIR or CIR rate limiting in order to reduce the CDV (Cell Delay Variation) to a minimum. Polling is used for the uplink MAC access, as this keeps uplink CDV to a minimum. The end point (CPE, E1 Concentrator or third-party equipment) retimes the AAL1 cell stream using an elastic buffer to offer a true CBR service between E1 interfaces.

QOS FOR NON-CBR SERVICES

With non-CBR services, QoS is enforced in two distinct areas:

- The AP Medium Access Controller (MAC) prioritises services and schedules bandwidth according to relative and absolute QoS settings within each sector
- The CPE can perform filtering of IP / Ethernet data onto individual services based on QoS criteria.

CIRCUIT VS PACKET SERVICES

The Medium Access Controller (MAC) treats packet and circuit services differently. Services using AAL5 encapsulation (e.g. Ethernet) initially contend and thereafter have their bandwidth reserved for the uplink access method. For CBR services, the MAC reserves uplink bandwidth based on downlink bandwidth to minimise uplink cell delay variation (CDV). This is what is known as a grant-based system. CDV will occur in the air interface and other parts of the transport system. As E1CES is a synchronous service, the CDV is removed in the CPE by the DSP which performs the Circuit Emulation Service (CES) termination. Cells arrive, with CDV, from the air interface and are placed in an elastic buffer in the DSP, which then paces the data out of the E1 interface as a true CBR service.

EPD SETTING

The VNMS uses the service's Early Packet Discard (EPD) setting to determine whether the service is carrying a CBR service, and therefore the appropriate uplink access method (contention or polling). EPD is automatically set to 'off' for CBR services, e.g. E1CES. However, if an ATM or loopback service is being used for symmetric CBR traffic (e.g. E1) then EPD must be set to 'off'. If the EPD setting is incorrect:

- The MAC may struggle to schedule the data efficiently, which can result in excessive CDV and a reduction in the sector's capacity
- The EPD mechanism assumes AAL5 and will break an AAL1 service
- IMA services should only have the EPD set to off if the service is carrying a symmetric service, e.g. an AAL2 service.

CIRCUIT SERVICES (AAL1, OR AAL1-LIKE)

The following services are carried over AAL1 services:

- E1CES
- optimised E1CES
- g804
- ATM with EPD set to 'off'
- IMA services where the priority class is set to zero (0) and EPD set to 'off'

PACKET SERVICES (AAL5)

The following services are carried over AAL5 services.

- vbridge
- bridge
- clip
- e1hdlc
- ATM with EPD set to 'on'
- IMA services where the priority class is set to 1, 2 or 3 and EPD set to 'on'

OPTIMISED SERVICES

In most cases, E1s are used in non-CBR applications i.e. where the real usage bandwidth is varying continuously, but the transport layer, the E1, does not vary in bandwidth. Such inefficiencies are particularly wasteful in a PMP system where any available bandwidth can be used by multiple CPEs. VectaStar removes this limitation by offering a suite of optimisation options for different types of CBR services. For all these options, the optimisation occurs within the VectaStar system and the original data stream is presented at the edges of the VectaStar part of the network.

One such option is Abis optimisation which reduces the constant rate E1 service to an optimised, variable bit rate service over the sector. This process is invoked in two stages. First, unused timeslots are removed (fractional E1) – this is often referred to as grooming. Second, the remaining timeslots are constantly analysed for activity. Only remaining timeslots carrying new information are carried across the service.

The example below shows the effect of optimising four E1 from a cell site on a live network. 3 of the E1s were used for GSM backhaul (Abis interface) and 1 E1 was used for 3G backhaul (Iub interface).

In this example, you can see that the optimisation reduces the bandwidth carried from 8 Mb/s (4 E1s) down to 1.5 Mb/s (less than 1 E1).

Cambridge BROADBAND NETWORKS

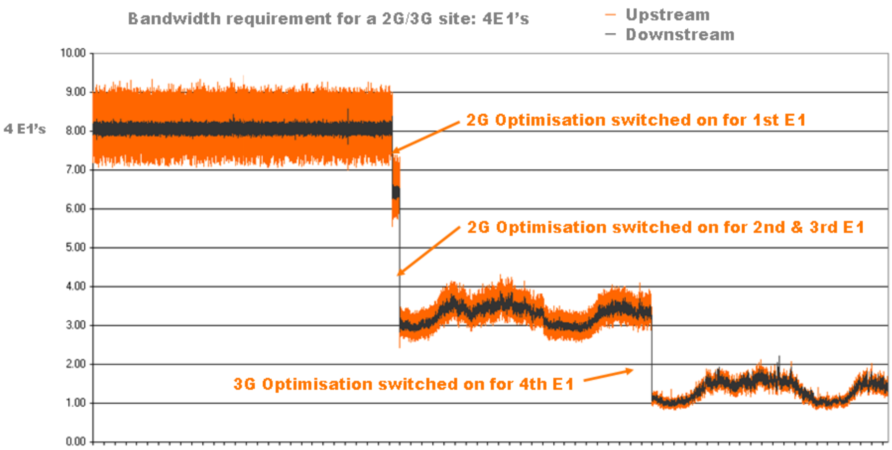

FIGURE 23: SEQUENTIAL IMPLEMENTATION OF OPTIMISATION ON FOUR E1 SERVICES

WHAT IS OPTIMISATION?

Optimisation turns a constant bit rate E1 (2.048 Mb/s) into a variable data rate service, an optimised E1, where unused bandwidth is saved and re-used elsewhere in the sector. Optimisation is transparent and loss-less i.e. the user/end equipment cannot detect it. Optimisation is specific to the type of service on the E1 i.e. PBX voice differs to Frame Relay over E1, as does 2G to 3G; therefore VectaStar has several different optimisation services, each specific to a different traffic type.

WHY USE OPTIMISATION?

Optimisation turns a constant bandwidth service into a variable bandwidth service, freeing up bandwidth. The freed bandwidth will vary with time and when the 2G / 3G Base Station is at peak load the freed bandwidth will be at a minimum. With a point-to-point (PTP) radio link or leased line, there is little benefit gained by optimisation as it is hard to re-use the freed bandwidth effectively. However, with point-to-multipoint systems, any bandwidth saved on one link is available for any other link in that sector.

Given that each link will vary independently of the others, a statistical multiplexing gain can be achieved (similar to standard TDM voice multiplexing using Erlang / blocking analysis).

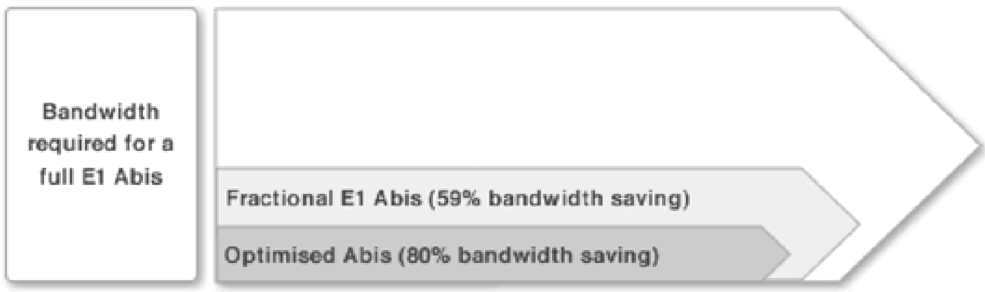

WHAT IS SPECIAL ABOUT VECTASTAR?

VectaStar combines unique optimisation at every E1 interface with exceptionally efficient statistical multiplexing in a Point-to-Multi-Point platform. This combination of features allows Abis and Iub interfaces to be statistically multiplexed and overbooked in a controlled and manageable fashion, realizing a significant effective bandwidth gain in the network.

HOW MUCH BANDWIDTH DOES OPTIMISATION SAVE?

Any permanently unused timeslots on the E1 generate a permanent saving, which can be confidently re-used by other services. On top of this, additional bandwidth savings are achieved by virtue of the fact that the load on the 2G / 3G base station varies over time and therefore the backhaul bandwidth requirement also varies over time. When statistically multiplexing several Abis / Iub interfaces, the overall bandwidth gain will be dependant on the usage statistics of those interfaces. The bandwidth required by a BTS, on the Abis interface, is a function of the instantaneous number of voice and data calls.

HOW IS OVERBOOKING MANAGED?

Optimisation allows a number of services to be statistically multiplexed within a sector. It is therefore possible that the sector could become overbooked occasionally (although correct planning should make this a very rare event). CBNL's optimisation algorithms contain a back pressure mechanism, which is used to manage over booking of Abis interfaces by gracefully reducing the on-air data rate in response to sector overload, by changing various speech Transcoder and Rate Adaptation Unit (TRAU) frames into idle TRAU frames to reduce demand.

ETHERNET SERVICES

There are two services which support Ethernet bridging:

- vbridge service
- bridge service

The bridge service provides an Ethernet transport where all received frames are carried transparently without filtering end to end. The vbridge service provides filtering options to allow sophisticated QoS to be applied to the Ethernet frames.

VBRIDGE SERVICE

The vbridge service at the CPE behaves like a hub rather than a bridge. It copies Ethernet frames, with the matching VLAN tag or no VLAN tag, between the Ethernet interface and the air link. The vbridge service uses a simple hub totally isolated from the IP stack. The vbridge service cannot be used in conjunction with the bridge service.

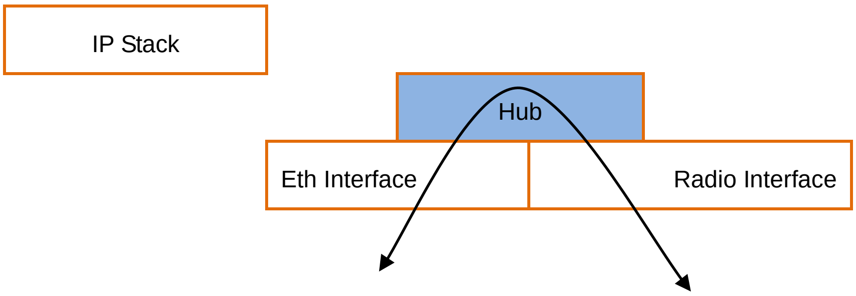

vbridge service treatment of Ethernet traffic	
CPE with Wallbox	The two hub interfaces are the Radio and the Ethernet interface on the CODU. The maximum Ethernet frame size is 1586 bytes.
CPE with 19" IDU	The two hub interfaces are the Radio and a specific IDU Ethernet interface. Multiple VLAN bridges (with different VLAN tags) per port are supported, but multiple ports per vbridge service are not supported. The Ethernet interface can support non-VLAN tagged frames. Each service maps to a single service. VLAN tags from 1–4090 are supported. VLAN tags from 4091–4094 inclusive are reserved. Creating a vbridge service will activate a disabled switch port (if that port is previously unused). The maximum Ethernet frame size is 1536 bytes if the frame is tagged on the IDU interface and 1532 bytes if the frame is untagged on the IDU interface.
Base Station (MPA)	At the MPA, the vbridge service creates one hub per service. Each service must have a unique set of filter criteria, e.g. VLAN, priority, etc. In R4.1 and early R4.2 (pre-4.2.50-xx), each VLAN can therefore only be connected to a single CPE.
Base Station (APC)	At the APC the vbridge service creates a true learning bridge. The vbridge will switch all frames between the APC's Ethernet interface and air link attached to virtual bridge groups which have Eth1 as a member. Multiple virtual bridges can be connected to the APC's Ethernet interface using different filtering criteria. See QoS Filtering and the vbridge Service on page 71 for more information. The maximum Ethernet frame size via the APC's Ethernet interface is 1586 bytes.

BRIDGE SERVICE

The bridge service can terminate as Ethernet on the CPE and/or the Base Station as well as being switched out of the Base Station as ATM VCs (RFC1483) into the core network and/or third-party termination equipment (e.g. Cisco router).

This service will be used when the customer wants to bridge all Ethernet traffic (including multiple VLANs) transparently over a single bridge. Multiple services can operate and these create multiple services into the same bridge. Spanning Tree Protocol (STP) is configurable.

The bridge service uses a fully-featured MAC address learning bridge, supporting STP, underneath the IP stack. As the bridge is connected to the IP stack, CLIP services, where supported, can route via an Ethernet interface that is already in use by a bridge service. The bridge service cannot be used on an Ethernet interface at the CPE if a vbridge service is already using that interface.

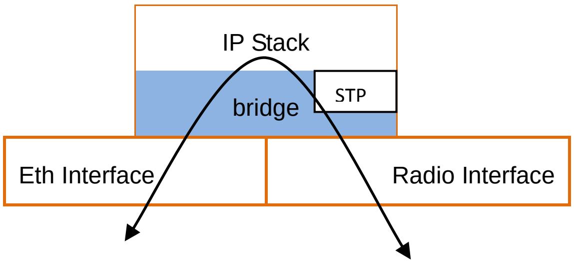

Bridge service treatment of Ethernet traffic	
CPE with Wallbox	Bridges Ethernet traffic entering the Wallbox Ethernet interface to a service. CPE Acts as a learning switch. The maximum Ethernet frame size is 1586 bytes.
CPE with 19" IDU	The bridge service will configure the IDU to act as a learning switch, i.e. it will enable all ports, forward local traffic between ports, and send everything else on to the CODU which will bridge all traffic onto a single service. The maximum Ethernet frame size is 1536 bytes.
Base Station (MPA)	The vbridge service is generally used rather than the bridge service on the MPA.

VLAN TRUNKING USING THE VBRIDGE SERVICE (Q IN Q)

From Software Release 3.5, an untagged vbridge service forwards all frames received on the specified CPE port, except the following:

- Those filtered off by other services
- Frames with reserved VLAN IDs 4091 to 4094

This functionality is called VLAN trunking.

Unique traffic streams within the radio network that egress the same port at the base station must be separated by adding an outer VLAN tag (Q in Q).

E1 SERVICES

VectaStar can carry E1 traffic using five different services:

- E1CES
- optimised E1CES
- sequentially optimised E1CES
- G.804
- IMA

The E1CES service can always be used to carry any E1 traffic (unstructured or fractional). However, in some cases it is more efficient (i.e. it uses less VectaStar sector capacity) to carry the traffic using one of the other services. For example:

- Optimised E1CES is more efficient for voice (PBX) transport and 2G (Abis) transport
- G804 and IMA are more efficient for carrying ATM traffic where each E1 is an independent G804 circuit (or a member of an IMA group).

E1 interfaces: E1 interfaces on IDUs and E1 Concentrators are presented on 120 Ohm balanced differential pairs via RJ48 sockets.

STANDARD E1CES – CIRCUIT EMULATION SERVICE (CES)

VectaStar is an ATM transport-based system. Each service is carried over a unique Virtual Service, which is switched from node to node through VectaStar. E1s are encapsulated using ATM Forum Circuit Emulation Service (af-vtoa-0078.000) on an AAL1 service. The interface can either be configured as an unstructured bit stream, or, it can have an ITU-T G.704 framing structure. These frames carry 32 TDM timeslots of one byte each at a rate of 8000 frames/s. The timeslots are numbered 0 through 31, and 0 is reserved for framing information.

Note: In the structured service, Sa bits within Timeslot 0 are not propagated across VectaStar.

E1 TYPES SUPPORTED

The following E1 types are supported:

- Unframed (also known as unstructured) E1 (G.703) The AAL1 service maps onto all 32 timeslots and uses 2048 Kbps plus 64 Kbps overhead (Total 2112 Kbps)
- Framed (also known as structured) E1 (G.704, 5a) Basic framing; The AAL1 service maps onto an ordered subset of timeslots {1…31} and uses 1984 Kbps plus 72 Kbps overhead. (Total 2056 Kpbs)
- Framed E1 CRC (G.704, 5b): As for Basic, plus CRC multi-framing to provide framing robustness and a way of monitoring the performance of the G.703 (G.704) link.

If you choose to configure a CES service as unstructured then you have the choice of whether to run a structured E1 test/service (G.704) or an unstructured E1 test/service (G.703) over the service. In this sense unstructured CES is more flexible, more generic. However, if you actually want to use/test G.704 structured E1 then you should get better error performance, and use 3% less bandwidth, by configuring the CES for a full structured E1.

VectaStar can maintain G.704 frame synchronisation on its E1 interfaces. With an unstructured CES service, a burst of lost or excessively delayed cells can cause G.704 frame synchronisation loss on the connected E1 equipment (frowned upon by all E1 testers). In a structured CES configuration this cell loss/delay will only affect frame payload.

For long-haul or error prone links, choose CRC4 on both the CPE and the E1 connected equipment for improved frame synchronisation and error detection.

The CRC only protects the E1 span, not the whole service. So, whereas the two interfaces connected by an E1 span must agree on use of CRC, the Base Station and CPE E1s can be different.

E1 TIMESLOTS

For framed E1 services, with and without CRC, you must specify which timeslots to use. A full E1 frame has 32 timeslots numbered from 0 to 31, of which slot 0 is reserved for framing information. A single timeslot corresponds to a 64Kbps connection. In the VNMS, when specifying timeslots, you must enter a comma-separated list of individual timeslots, or hyphenated ranges. If you want to use all timeslots, enter 1–31. Timeslots are always carried in ascending order. Note that both CES ends must agree in the number of timeslots being carried, although their precise positions need not. For example, the 3 timeslots 4, 5, 7 may be delivered as timeslots 12, 17 and 19.

Note: When using fractional E1 CES, up to 10 fractional CES can terminate on the same physical interface. Where multiple E1 CES terminate on the same physical interface, there can be no overlap of timeslots between the services i.e. each timeslot is uniquely associated with a service

CDV TOLERANCE

Default is 15ms.

The CDV Tolerance is the maximum cell arrival jitter in 10 microsecond increments that the reassembly processor will tolerate in the cell stream without producing errors on the CBR service interface. If this value is too small then underflow events may occur, which result in dummy data being played out. If it is too large then unnecessary delay is introduced. The maximum value is 50ms (5000).

MAX BUFFER SIZE

Default is 25ms.

This defines the maximum size in 10 microsecond increments of the reassembly buffer. If this value is too small then overflow events will occur, and data is discarded. If too large then unnecessary delay may be introduced into the circuit. In all cases it must be larger than the CDV tolerance parameter. The maximum value is 100 ms (10,000).

CHANNEL ASSOCIATED SIGNALLING (CAS)

The structured CES standard includes an optional feature for carrying CAS particularly efficiently for fractional E1. VectaStar currently does not support this, so it is essential that the equipment providing the upstream CES termination is configured to use BASIC AAL1 format rather than CAS AAL1 format, otherwise the service will fail. Where only one VC is associated with an E1 interface, any signalling can be carried over VectaStar by simply including TS16 in the fraction.

ABIS OPTIMISED E1CES

The E1CES service has an optimisation option “GSM Abis” which can be used for the following traffic types:

- GSM Abis interface
- PBX voice trunks

Only a full (31TS) structured service is supported.

E1CES optimisation is a transparent, lossless process that reduces the VectaStar radio bandwidth used through the following efficiency savings:

- Unused or idle 64 kb/s time slots are compressed to zero bandwidth (dynamic grooming)
- 16kbit/s TCH carrying TRAU/PCU frames are compressed depending upon their activity (dynamic grooming of GSM full rate codec voice channels)
- HDLC signalling channels which are compressed depending upon their activity (optimisation of GSM / PBX signalling channels)

It is important to note that as soon as an inactive or compressed timeslot (or TCH) becomes active, the data is passed across VectaStar with no loss or excess delay. The end equipment cannot detect that optimisation is taking place and neither can E1 test equipment.

Optimisation is lossless if the sector is not overbooked with optimised E1 services. In the event that it is overbooked with optimised E1CES services then there is a finite probability (depending on traffic distributions of each service) that the system will be overloaded, in which case the optimisation will cease to be lossless. This may affect other E1 services in the same sector. VectaStar has support for graceful degradation of GSM Abis traffic to reduce bandwidth demands. Speech quality may degrade marginally, by dropping occasional speech frames.

You can use regular pseudo-random binary (PRBS) test equipment, but with only a limited number of timeslots, as the PRBS creates an unrealistic load on the optimiser (quite unlike realistic TDM/Abis traffic). Eight timeslots should be fine.

The GSM Abis Optimisation service has the following constraints applied to this service:

- Partial cell fill in not available
- E1 type must be framed G704, with or without CRC
- Service is structured with all timeslots potentially carried (1-31)

G804 SERVICE

The g804 service is designed for carrying the following traffic:

- single E1 ITU G.804 (ATM over E1)
- single E1 within an IMA group “af-phy-0086, Inverse Multiplexing for ATM (IMA) Specification Version 1.1” or version 1.0

g804 is a transparent, lossless process which reduces the bandwidth used in VectaStar to the minimum possible, by removing Idle/Unassigned cells.

The following E1 types are supported:

- Framed (also known as structured) E1 (G.704, 5a) Basic framing; The AAL1 service maps onto an ordered subset of timeslots {1…31}
- Framed E1 CRC (G.704, 5b): As for Basic, plus CRC multi-framing to provide framing robustness and a way of monitoring the performance of the G.703 link

For long-haul or error prone links, choose CRC4 on both the CPE and the E1 connected equipment for improved frame synchronisation and error detection. The CRC only protects the E1 span, not the whole service. So, whereas the two interfaces connected by an E1 span must agree on use of CRC, the Base Station and CPE E1s can be different

Relative cell timing is preserved over VectaStar.

Optimisation is lossless if the sector is not overbooked with optimised E1 services. In the event that it is overbooked with optimised g804 services then there is a finite probability (depending on traffic distributions of each service) that the system will be overloaded, in which case the optimisation will cease to be lossless. This may affect other E1 services in the same sector.

You cannot use regular pseudo-random binary (PRBS) test equipment.

Typically 3G services utilise CRC G.704 framing and timeslots 1-15, 17-31, with payload scrambling enabled.

E1 TIMESLOTS

For framed E1 services, with and without CRC, you must specify which timeslots to use. A full E1 frame has 32 timeslots numbered from 0 to 31, of which slot 0 is reserved for framing information. A single timeslot corresponds to a 64Kbps connection. In the VNMS, when specifying timeslots, you must enter a comma-separated list of individual timeslots, or hyphenated ranges. If you want to use all timeslots, enter 1–31. Timeslots are always carried in ascending order. Note that both g804 ends need not agree on number or position of timeslots, unless the service is operating near capacity. The default set of timeslots is 1–15, 17–31, as suggested in ITU-T G.804.

CELL FLUSH TIME

ATM Cells extracted from the E1 are tunnelled over the provisioned service. Cell flush time is how long a partially generated cell is held waiting for more data before it is flushed to the tunnel i.e. a compromise between latency and bandwidth at low data rates.

Default 500uS

PLAYOUT BUFFER

VectaStar can maintain relative ATM cell timing from end to end. To achieve this buffering is necessary to remove any jitter introduced. This parameter is the target level of buffering. Set it to zero for low latency but non-jitter free operation. Underflows only cause cell timing errors, not data errors.

Default 15000uS

OVERFLOW

This is the upper limit of latency buffering. Overflows, which reduce the latency towards the target, cause cell timing errors only, not data errors.

Default 25000uS

SCRAMBLING

This is whether or not the cell payloads are scrambled as per ITU-T I.432 on the E1 interface.

IMA SERVICE

IMA, or Inverse Multiplexing for ATM, is based on a standard defined by the ATM Forum, AF-PHY-0086.001. The purpose of IMA is to be able to bind together a group of E1 interfaces to provide a single logical interface whose bandwidth is greater than a single E1, through which ATM cells can be carried. The IMA service is designed for carrying the following traffic:

- single E1 IMA groups requiring service prioritisation, otherwise the g804 service is as efficient.
- multiple E1 IMA groups where service prioritisation is required, otherwise the multiple instanced of the g804 service are as efficient.

Several VectaStar IMA service services can be defined on the same IMA interface, each corresponding to a single service flow on the IMA interface. In VectaStar, all interfaces in an IMA group must be on the same backplane (an IDU or E1 Concentrator).

The IMA specification provides for differential latency on the E1s, and for a varying number of E1s, as they fail or come back into service. G.704 framing with CRC4 is assumed on all IMA E1 interfaces; timeslots 1-15 and 17-31 are assumed for carrying the cells. Scrambling is assumed on all cell payloads on E1. Cell headers are automatically switched between external VPI/VCI and internal VCI by IMA termination.

Configuration parameters for CPE or BS termination:

- The set of E1 interfaces comprising the IMA group
- VCI/VPI (unrestricted range) of cells for this particular service.

Multiple IMA groups may be configured on a single VS BS/CPE backplane, commensurate with the number of E1 interfaces available of course.

There are some restrictions with VectaStar's IMA service:

- The IMA service is currently optimised for up to 64 services per IMA group, although this is not a hard constraint.
- CPE redundancy is not supported for IMA within VectaStar (as it is for a CES service terminated at BS E1CONC, for example).
- Optional features in the IMA specification are not supported.
- ITC (Independent Transmit Clock mode) is not supported, due to VectaStar hardware constraints.
- The IMA MIB (Appendix A to the IMA spec) is not currently supported.
- You cannot use regular pseudo-random binary (PRBS) test equipment.
- IMA interfaces must be framed with CRC (G.704, 5b).

E1 INTERFACES

Any number of E1 interfaces can be used as an IMA group providing they all reside on the same backplane. Non-contiguous E1 interfaces are supported e.g. interfaces 1, 3 & 5 on an E1 Concentrator.

VPI/VCI

The VPI and VCI full range (0-65535 for VCI, 0-4095 for VPI) is supported. Each active service must be configured separately.

IMA VERSUS G804

There are important differences and similarities between the IMA service and the current g804 service. They are both designed specifically to carry ATM cells through an E1 interface. The proprietary g804 service provides a tunnel to carry all ATM cells, irrespective of their VPI/VCI, from one E1 interface through to another. Each E1 may be independent or one of an IMA group, in which case that group is extended across both E1 interfaces using one service per E1 interface. With the IMA service, however, the IMA group is terminated and VCs extracted may be routed from this logical ATM interface independently; cells arriving on the interface which have no corresponding VC are dropped. Therefore IMA has the advantage of exposing individual VCs which can be separately prioritised within VectaStar e.g. to prioritise voice over data.

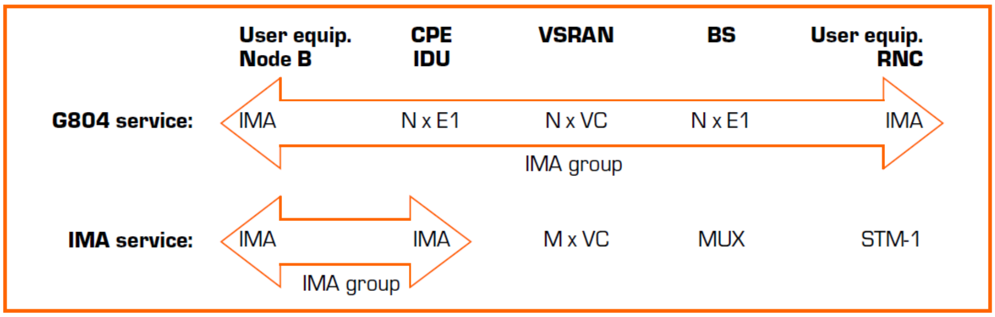

FIGURE 24: G.804 AND IMA COMPARISON

SYNCHRONISATION

TDM services are inherently synchronous. This adds a layer of complexity to the system in that not only does the data have to pass through the system with integrity, but the synchronisation must be maintained between end-points. In order to maintain flexibility of architecture, VectaStar uses a synchronisation scheme where the synchronisation information, used to keep STM-1 and E1 interfaces frequency locked, is passed through the network using dedicated synchronisation management services.

The VNMS has GUI support for configuring multiple synchronisation domains, where each domain:

- can use any synchronous interface as a Primary or Secondary clock source, whether on a CPE or Base Station interface
- can span multiple Base Stations
- supports automatic fail-over between Primary and Secondary, in the event the Primary clock source fails
- supports configurable (automatic or manual) fail-back, from Secondary to Primary.

The only constraint is that the synchronisation domains are mutually exclusive (i.e. they cannot share clock sources or clock syncs).

Synchronisation domain timing is propagated from each Access Point to all the CPEs within the sector via a broadcast synchronisation management service or a domain-specific synchronisation management service. The process which manages these services is called *syncd*. *syncd* is responsible for physical layer synchronisation and can act as a source or sink for timing information. This process can run on a CPE or APC.

Most E1 configuration problems are caused by third-party equipment not being correctly configured to:

- offer a synchronous ATM interface which is phase locked to the E1 source clock
- have appropriate buffer sizes set to tolerate the uplink CDV.

Note: For IMA, g804 and e1hdlc, it is not strictly necessary to have clocks locked at both E1 interfaces, although this may be required for other reasons (i.e. both ends could run on local or loop clocks independently).

Points to bear in mind:

- A MUX automatically becomes a target of a domain when it is used as either a primary or secondary reference
- An E1 Concentrator automatically becomes a target of a domain when it is used as either a primary or secondary reference
- An E1 Concentrator can also be specified as part of a domain from vse1conc (this option may be greyed out if it is a domain reference)
- A CPE E1 backplane can only be made a target of a domain by creating an E1 service that uses synchronisation. The user can specify which domain using the Synchronisation tab, but this will be greyed out if the backplane is already a target or reference for a domain.

SYNCHRONISATION TOPOLOGIES

CPE E1 SYNCHRONIZED TO BASE STATION STM-1 INTERFACE

Most networks require a synchronisation scheme in which the CPE E1 is synchronized to Base Station STM-1 interface. This topology uses E1/CES/ATM from a Network interface of the MPA through to the E1 interface at the CPE with the synchronisation derived from the MPA port.

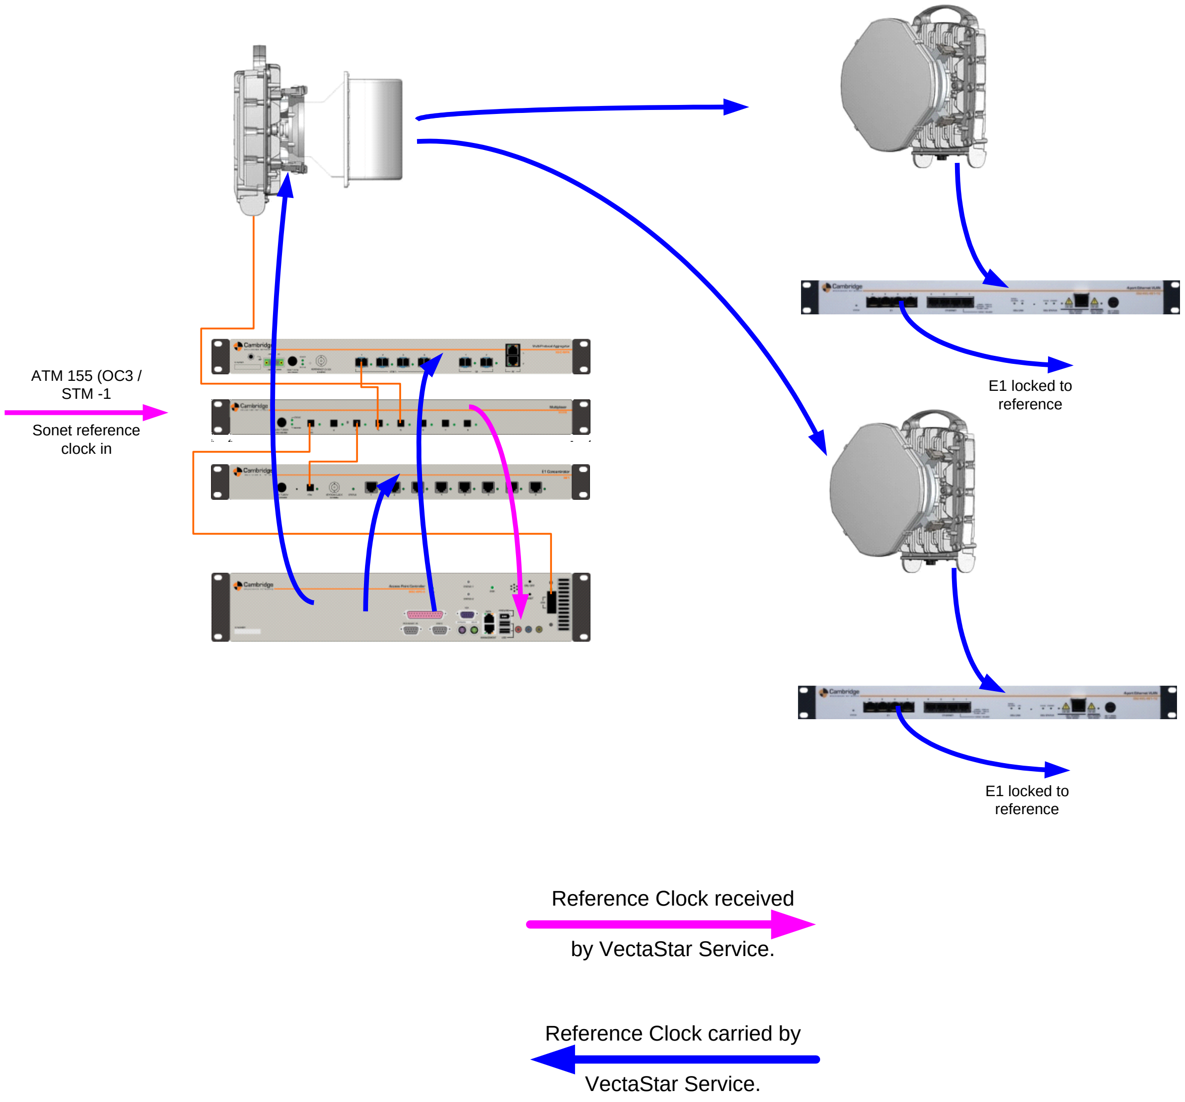

FIGURE 25 – CPE E1 SYNCHRONIZED TO BASE STATION STM-1 INTERFACE

CPE E1 SYNCHRONIZED TO BASE STATION E1 INTERFACE

This topology uses an E1 Concentrator at the Base Station to derive network timing and propagates it to all CPEs within the sectors:

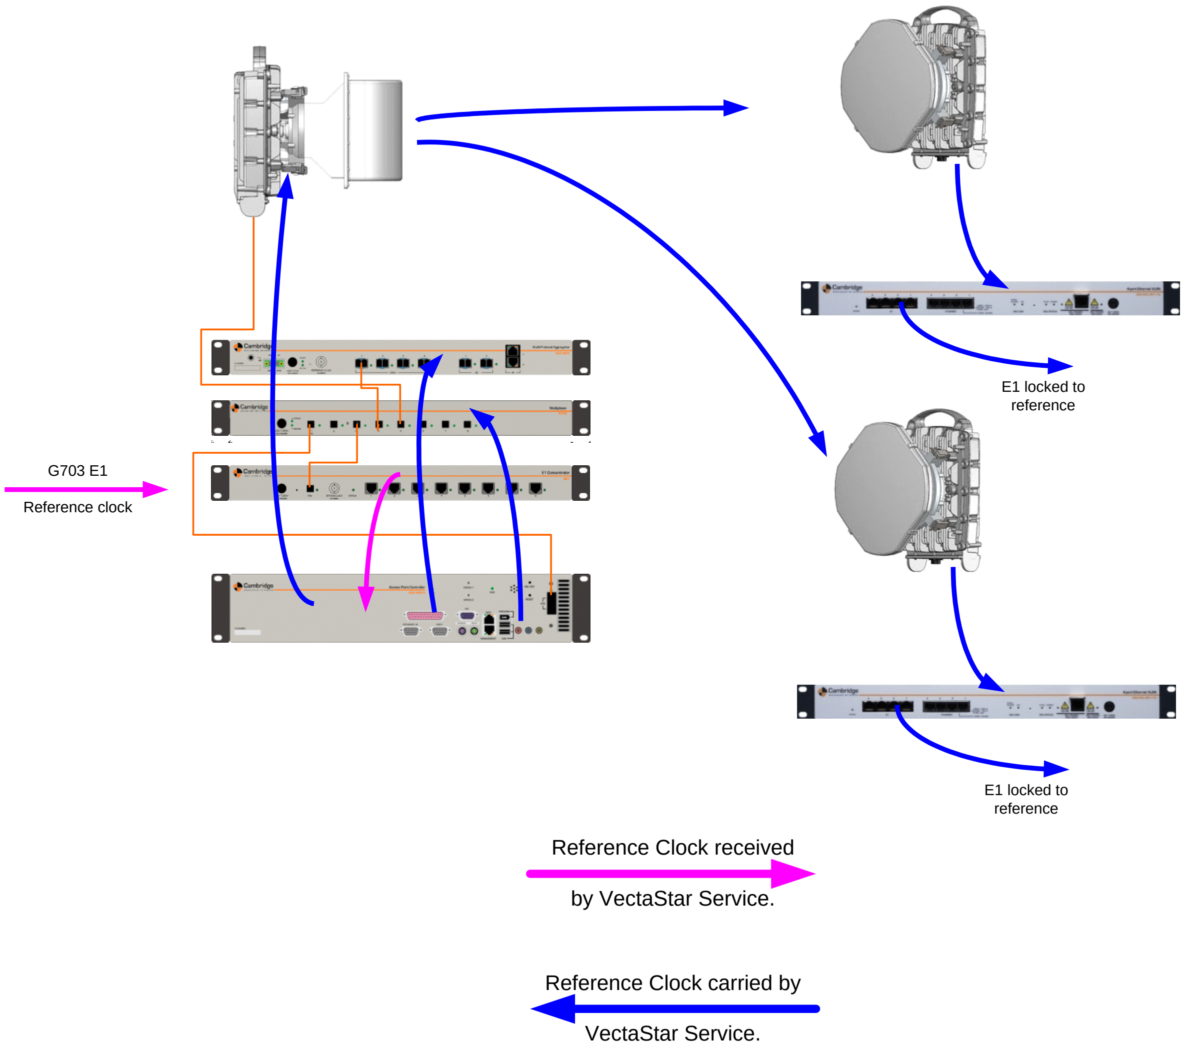

FIGURE 26 – CPE E1 SYNCHRONIZED TO BASE STATION E1 INTERFACE

The above diagram depicts a single synchronization domain. If there are other E1 CPEs connected to this Base Station, and other E1 Concentrators, VectaStar allows these to be synchronized independently, as part of a separate synchronization domain. However, note that each E1 backplane and each Multiplexer can only serve as a reference for only one domain, or as a target of only one domain.

BASE STATION SYNCHRONIZED TO CPE E1 INTERFACE

This topology uses an E1 interface at the CPE to derive network timing which is then propagated to the base station and to all other network elements in this synchronization domain:

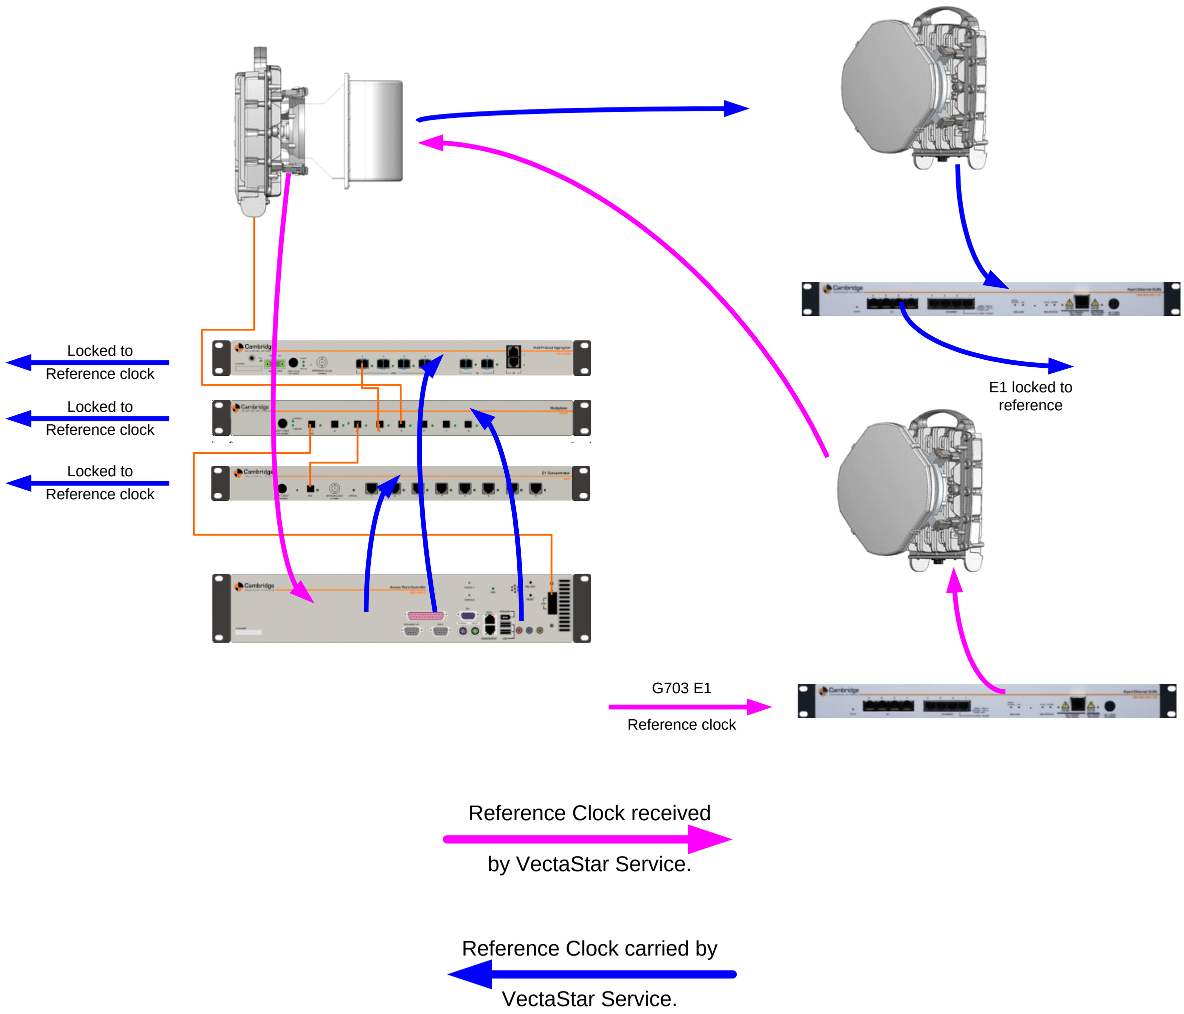

FIGURE 27 – BASE STATION E1 CONCENTRATOR SYNCHRONIZED TO CPE E1

Note that multiple E1 interfaces on a single VectaStar device are all timed from the same source.

SYNCHRONOUS CLOCKING

For a CES service to conform to the relevant standards, there is a requirement that the input and output E1 clocks be locked together. This could be achieved by using external clocks at each end, e.g. GPS or a method whereby the clock is carried with the data. Some applications, e.g. early GSM BTS, require network timing and recover this by averaging the incoming E1's clock. Synchronous clocking, where the timing is passed through the system with the data, is the default E1 timing configuration for VectaStar.

One of the disadvantages of synchronous clocking is that the physical layer clock must be configured at each node to ensure that the clock is passed to the next node and in particular the STM-1 connected to the MUX must have its physical layer clock phase locked to the source E1 clock. With some third-party equipment this is not always possible. Another complication is that as the E1 CES termination service is pacing data out of the E1 interface, locked to the source E1 clock, there is a finite tolerance for CDV within the complete system. This tolerance is determined by the size of the buffers in the VectaStar CPE CES service and also the E1 Concentrator or the core network CES termination service. If these buffers are too small, then CDV may cause data errors within the system. Larger buffers will increase the tolerance to CDV at the expense of increased latency.

Note: It is important to understand the role of buffering in Circuit Emulation. It is only relevant when data is received from AAL1 cells and then played out of the E1 interface. The E1 interface is a Constant Bit Rate (CBR) interface whereas the timing of the arrival of the Radio cells will (most probably) exhibit some randomness, even if they are being generated at exactly a constant and the correct rate, because there will be some variation in their individual propagation times through the radio network. Suppose a radio cell is late arriving and there is some data from previous cells still to be played out of the E1 interface, then, as long as the next cell arrives before all this 'slack' is used up, the E1 output data stream is uninterrupted and the lateness is acceptable. CES deliberately builds some slack into this data path by initially delaying the E1 play out when cells first arrive. If a buffer becomes exhausted, arbitrary data must be transmitted to maintain the constant bit rate. This is called a buffer underflow.

When using synchronous clocking it is important to remember that the external equipment, connected to VectaStar, must be correctly configured. Normally, this means:

- Equipment connected to a slave E1 backplane (IDU, ICU or E1 Concentrator) should be set to use the recovered clock (i.e. loop timing)
- Equipment connected to a master E1 backplane should be set to present a clock
- Equipment connected to the MUX should be set to present a clock to VectaStar if it is the clock master.

There are some cases where synchronous operation is not possible or required, for example where the E1 is not required to be synchronous.

In these cases it is possible to use adaptive clocking in the CPE.

ADAPTIVE CLOCKING

For situations where strict physical layer synchronisation is not required, adaptive locking may be suitable.

VectaStar's adaptive clocking algorithm controls the E1 transmit clock such that it matches the average rate of arrival of AAL1 cells i.e. each CPE adjusts its own internal 2.048MHz clock independently of the rest of the network.

If the cells arrive at the CPE with no CDV, then in theory, this will operate identically to a synchronous service, however there are several issues to be aware of:

- In reality, there will be CDV and this will cause the E1 clock to wander. This may break jitter and wander restrictions on the clock (ITU-T G.823).
- When an adaptive clocking E1 initially starts, there may be a short period of rapid variation in the E1 clock (up to 150ppm) as the adaptive algorithm settles.
- If a burst of n>7 cells is lost in the ATM network, then the CPE Tx clock will permanently shift by approximately 47 x 8 x n cycles compared to the remote E1 reference source clock.
- VectaStar units with multiple E1 interfaces have a common clock, so this cannot adapt each interface to its own VC. In practice, the clock is most affected by the VC with the slowest reference source.

This said, adaptive clocking is very useful when fault finding synchronisation problems as it allows the operator to separate the problem of passing the data from the problem of passing the synchronisation.

SECURITY

The standard VectaStar system has a number of security features, which can be used to prevent a CPE registering on a sector where it is not permitted. In the case of two commercial operators using VectaStar in the same town, it is actually useful to know that both systems can see each other and to be aware of any attempted cross-registrations as this implies that either frequency planning or configuration management needs to be addressed. A CPE can detect which operator it is receiving data from by reading the broadcast operator descriptor field sent to it by an Access Point. If a CPE attempts and fails to register on another operator's network, that operator will see an SNMP trap informing him of the attempt.

The VectaStar system uses separate virtual circuits to carry the data for each service. This ensures that another user or service cannot unintentionally intercept the data. Data from one subscriber is only available on the user interface of other subscribers if the system has been configured with a service that permits this. If desired, the system can be configured such that users can only see their own data.

VectaStar uses a proprietary and complex MAC layer, which cannot be reproduced without access to both VectaStar AP and CPE hardware and the VectaStar development software. Although CBNL works with the open source community, CBNL does not release the source code for these proprietary components.

ENCRYPTED MANAGEMENT

VectaStar is fully SNMP compliant and all management operations are achieved using SNMP e.g. registration and service creation uses SNMP. VectaStar supports SNMP v1, v2c and v3. SNMP v3 supports encryption and authentication i.e. the user at the management application is authenticated by the SNMP v3 entity before the entity allows the user to access any of the values in the MIB objects on the agent. In addition, all of the requests and responses from the management application to the SNMPv3 entity are encrypted, so that all the data is completely secure.

It is important to note that authentication and privacy are bi-directional. In other words, the SNMPv3 entity asking for the information is authenticated by the SNMPv3 entity providing the information. When the SNMPv3 entity receives the information, it authenticates the provider of the information. An alternative way of thinking about this is that no SNMPv3 entity trusts another SNMPv3 entity. VectaStar uses SNMP v3 for authenticating users at registration.

END TO END ENCRYPTION

Cambridge Broadband expects all security conscious operators, to use end-to-end data encryption across their entire transport network - as this is the ONLY secure model for encryption. Relying on the transport network for encryption is fundamentally flawed, as any line of attack will target the weakest link, which is often the transport network. End to end encryption ensures that the weakest link is no longer the transport network which means that the operator, who is concerned about security, need no longer worry about the vulnerabilities of the transport network.

The level of security provided by VectaStar can be compared to that of a good wired telephone system. Callers can be relatively certain that their call is private and between the caller and the recipient. The only interception possible is with the support of the telephone operator therefore complete security is provided by encryption of the data by both parties. Where end-to-end encryption is implemented, the VectaStar transport does not interfere with the security implementation.

Standard encryption hardware can be used with VectaStar for both CBR and packet-based transport schemes. With CBR systems, an E1 interface is encrypted using standard serial data stream crypto hardware. With IP connections, standard crypto VPN technology can be used to provide a secure IP tunnel across the transport network.

SECURITY AGAINST MISCONFIGURATION AND HACKING

All VectaStar CPEs and Base Stations are managed by IP connections. The use of a secure SSH2 connection guarantees authentication and encryption of every remote management connection. SSH access combined with layer 3 firewalls (iptables and tcpwrappers under Linux) ensures that strict access control can be enforced across the network. When combined with good network design this greatly reduces the possibility of hacking. In the event a CPE is compromised (captured) the user will be unable to alter their configuration or hack into the control network to change their connection (unless that unit is already authorised to do so). Use of an obfuscated password and regularly changing the access passwords will ensure that even if a unit is compromised, without knowledge of the keys over time, its usefulness will be temporally constrained.

RADIO FREQUENCY PLANNING

This chapter gives an overview of the VectaStar radio interface and how to manage the radio aspect of a VectaStar network.

Every radio link should be planned to ensure that the desired availability is achieved. The choice of availability may be determined by an SLA (Service Level Agreement) or it could be a pragmatic compromise between cost and range.

In order to simplify the planning process, Cambridge Broadband Networks provides an Excel-based Link Planning tool that enables the operator to specify the VectaStar equipment and ensure that it achieves the desired availability for any given LOS (Line of Sight) radio path. The planning tool uses ITU rain and fade data to calculate the required fade margin.

There are two aspects to link planning:

1. Customer specific requirements:
    i. Desired link availability (e.g. 99.995% per annum)
    ii. Frequency (e.g. 10.5GHz or 26GHz)
    iii. Link range, which determines Free Space Loss
    iv. Path inclination (i.e. relative height of CPE to AP)
    v. Geographic area (i.e. Latitude and Longitude)
    vi. Type of polarisation (e.g. vertical or horizontal)
    vii. Fade margin, which is a function of link reliability, frequency, range, path inclination, polarisation and geographic region
2. Parameters that affect the available link budget[2]:
    i. Channel size (as narrower channels have lower noise than wider channels)
    ii. Antenna gains at each radio (AP and CPE)
    iii. Transmit power, which depends on frequency variant and modulation
    iv. Losses in the radio system (Rx Noise figure, Tx filter losses, etc.)
    v. Modulation required – this determines the minimum level of Carrier to Noise Ratio (e.g. 64 QAM requires a higher CNR than 16 QAM)

Typical Link Ranges[3]	Cambridge, UK	Johannesburg, RSA	Kuala Lumpur, ML
10.5GHz	22.2km	17.4km	12.8km
26 GHz	6.1km	3.9km	2.6km
28 GHz	5.6km	3.6km	2.4km

Link ranges determined on 99.995% reliability, 1E-9 BER, 60cm CPE antennas, 14 MHz channel. AP 100m and CPE 20m above sea level

[2] Link budget is defined as the maximum attenuation that can be sustained at an acceptable Bit Error Rate (BER) between two radios before communication is no longer possible.

[3] A minimum range of 100m applies to VectaStar radio links.

FREQUENCY DIVISION DUPLEX (FDD)

VectaStar is a Frequency Division Duplex (FDD) system, which means that separate downstream and upstream radio channels can be used simultaneously by all radios. In a Point-to-Multipoint (PMP) system the downstream channel is broadcast – all CPEs listen to the same transmission from the Access Point. The downstream broadcast contains some information that all CPEs are expected to use, for example:

- Operation and Maintenance information
- Synchronisation information
- Downlink and uplink frame description

Each downstream frame may contain cells in any of the support modulation states and not all CPEs will be able to demodulate all cells. For example, a CPE at range using QPSK may not have the Carrier-to-Noise Ratio (CNR) to demodulate 64 QAM cells. Each CPE uses the information in the downlink frame descriptor to identify the data in the downstream that are intended for it and also to find out when it is scheduled to transmit upstream information. The frame description is always transmitted in the lowest modulation, QPSK, as all CPEs must be able to demodulate this.

MODULATION SCHEMES

VectaStar supports two methods of defining the modulation: static or Hitless Adaptive Modulation. Both schemes are able to use any/all of the supported modulations states.

The modulation is set per service so each CPE can simultaneously support services on any of the available modulations. Both Static and Adaptive modulations can be mixed within a sector and within a CPE. The wide range of modulation states allows a trade off between spectral efficiency and link budget.

CHOOSING A MODULATION SCHEME

The default modulation setting is QPSK. Before changing to another setting, it is important to understand fully the implications (especially before choosing Adaptive Modulation, which should only be used under certain circumstances).

In terms of capacity, the higher the modulation used, the more data can be transferred per radio symbol. The spectral efficiency of each modulation state is shown in the following table.

VectaStar 2 Spectral Efficiency b/s/Hz			
Modulation	**7 MHz Channel**	**14 MHz Channel**	**28 MHz Channel**
256QAM+T	5.00	4.94	4.68
128QAM+T	4.29	4.36	4.10
64QAM	3.86	3.99	3.69
64QAM+T	3.43	3.66	3.31
16QAM	2.43	2.59	2.34
16QAM+T	2.14	2.27	2.01
QPSK	1.43	1.42	1.30

VectaStar 2 Sector Ethernet capacity per channel size (Mb/s)			
Modulation	**7 MHz Channel**	**14 MHz Channel**	**28 MHz Channel**
256QAM+T	35	65	130
128QAM+T	30	60	120
64QAM	27	55	110
64QAM+T	24	50	100
16QAM	17	35	70
16QAM+T	15	30	60
QPSK	10	20	35

Note: when testing a VectaStar network for throughput performance, results may vary according to test configuration and traffic profiles. Please contact CBNL support for more information on test methodology.

In an ideal world, all service VCs would use the highest modulation available, as this is most spectrally efficient. Unfortunately, the choice of modulation cannot be made on capacity alone since it will also affect the Transmit power and Minimum Receive Carrier to Noise ratio: as the modulation type is increased, the transmit power is reduced. At the same time, higher modulation schemes require higher CNR levels. The result is that the link budgets of higher modulations are always lower than those at higher ones. What this means is that for a given SLA reliability, higher modulations have a shorter range than lower modulations. In most deployments there are CPEs where the range does not permit the use of the highest modulation scheme and a lower modulation is used instead.

ADAPTIVE MODULATION

If a VC has its modulation configured to adaptive modulation, VectaStar will start the VC at the highest modulation supported and, if necessary, reduce the modulation until a stable service can be established. The CPE transmit power is fixed to that which is required for the highest modulation scheme, regardless of the actual modulation in use.

The change to a lower modulation is hitless i.e. it occurs fast enough for error correction to mask any on-air errors (e.g. within a few ms of a fade occurring). The change to a higher modulation is also hitless but takes place on a longer timescale with a minimum wait of approximately 60 seconds, to ensure stability and prevent rapid oscillations between modulation states.

Careful link planning must be undertaken with due consideration for different traffic flows (OAM, voice, data) and link availability to ensure that each sector is operates within the desired parameters These may defined as no overbooking or to limit a sector overbooking by a desired factor.

ADAPTIVE MODULATION AND FADING

All radio links will suffer from fading. Typically, the longer the link and the higher the frequency, the more severe and frequent fading will be. Thus, radio links are planned to take account of fading and ensure reliable operation in the presence of a fading environment. This is done by allowing for a Fade Margin in the link budget i.e. spare power which is kept in reserve until needed to maintain the link in a severe fade. However, when the link is not fading, the link appears to be running conservatively may be able to sustain a higher modulation. In such a case, it is true that the link could operate at a higher modulation, but only during the fade-free period. For example, using a simplistic example of a link: 99.999% availability may be achievable at QPSK, 99.99% at 16 QAM and 99.9% at 64 QAM.

Given that the most likely cause of fading at higher frequencies, rain, will probably affect more than one link in a sector at the same time, the fading will not be uncorrelated within a sector and many links may simultaneously drop down to a lower modulation. However, if the capacity is planned taking account of the modulation recommended by the link planning process rather than the best case modulation, severe overbooking caused by fading can be avoided.

One advantage of adaptive modulation is that even if the links have been planned carefully and the capacity managed properly, then in the event of an unexpectedly severe fading event which would normally cause loss of service to static links, adaptive modulation may be able to preserve services, albeit with a reduced capacity. Obviously if the fade event is so severe that QPSK links cannot be sustained, service will be lost regardless of whether the modulation is statically or dynamically controlled.

PLANNING A LINK IN ADVANCE

As there are many factors that can affect the link budget and therefore the available modulation, care should be taken to maximise spectral efficiency i.e. ensure as many CPEs as possible use a high order modulation.

CPE choices that affect link budget are:

- Antenna choice e.g. using a different size antenna

System choices which affect link budget are:

- Frequency of deployment[4] e.g. 10.5 GHz or 26 GHz
- Channel bandwidth[5]
- Link reliability[6]

To facilitate link planning, an Excel link planner tool is available from Cambridge Broadband, which can be used to plan each link to ensure that the required modulation and link-reliability can be achieved.

SELECTING THE BEST MODULATION FOR A DEPLOYED CPE

Once a CPE is deployed, the CNR should be checked to ensure that it can operate correctly at the desired modulation. Be aware, that the CNR may change over time due to the following causes:

- An obstruction e.g. foliage growth or new buildings
- Atmospheric fading/multi-path which affects propagation when it is not raining
- Rain fade (noticeable at 10.5GHz and particularly prevalent at 26 and 28 GHz)
- Interference (either from another VectaStar radio using the same frequency within range or another radio system). Note: This may only become apparent at a later stage, e.g. after another VectaStar Base Station has been installed.

To check the CNR currently achievable for a CPE, please refer to the System Monitoring Guide for details.

RAIN FADES, ATMOSPHERIC FADES AND AVAILABILITY

Typically each operator/customer will have a target link reliability, often in the range 99.9% to 99.999%. The parameter that will affect link reliability most significantly is the fade margin. The fade margin is a parameter determined by link reliability, frequency, range, path inclination, polarisation and geographic region:

- High link availability will increase the fade margin
- Higher frequencies require more fade margin (all other parameters being kept constant)
- Shallower path inclinations require higher fade margins
- With rain fades, vertical polarisation has lower loss than horizontal polarisation, however with paths over water, the converse may be true
- Both atmospheric fading and rain fading vary around the world by region

To determine the link margin, the ITU provides a statistical prediction method (detailed in ITU-R P.530), which is the basis for the VectaStar Link Planner Excel Tool.

---

[4] Free space loss increases with frequency, therefore longer ranges are achieved at lower frequencies for the same link budget.

[5] Wider channel bandwidths reduce the Transmit power spectral density and increase the noise bandwidth, therefore reducing the link budget. However, wider channel bandwidths increase sector capacity and allow for improved statistical multiplexing.

[6] Link reliability, e.g. 99.99%, will determine the required fade margin. Higher availabilities necessitate a higher fade margin, which reduces the link budget and therefore range.

POWER CONTROL

Power control on the downstream is fixed (set in the VNMS) and the Access Point transmits at a constant power level, usually full power. The CPEs use automatic gain control (AGC) to adjust their receiver gain such that they compensate for the path loss in the radio channel. As each CPE only listens to one Access Point the problem of tracking path loss fluctuations is quite simple and, if necessary, the CPE will automatically adjust the gain in its receiver to compensate for any power variations.

The upstream frame, as seen by the Access Point, consists of a number of transmissions from several different CPEs. Each of these transmissions has passed over a different radio path that may be varying with time. The upstream access control arbiter supports contention (for non-CBR services) and this means that the AP cannot predict which CPE is transmitting in the contention slots. Therefore the AP runs without Rx AGC and the upstream AGC is distributed between the CPE and the AP, with the CPE adjusting its transmit power and the AP providing feedback.

Cellular re-use modelling of Fixed Wireless Access (FWA) and mobile systems reveals that the dominant cause of performance-degrading interference is CPE upstream interference between cells reusing the same frequency. Therefore it is important to ensure that each CPE only transmits at the minimum power necessary to maintain a good radio link with the AP, as this minimises upstream interference. In the event that the pathloss between an AP and a CPE changes, the CPE adjusts its transmit power accordingly.

Minimum CPE transmit power is dictated by the power level received by the Access Point which typically uses a closed loop power control mechanism to feed back to the CPE the correct transmit level it should use. There is a different set point at the Access Point for Rx power from CPEs for each modulation scheme. As the Closed Loop power control in the upstream cannot react instantly, there is a built-in fade margin (configurable per sector), nominally 5dB, which ensures that an instantaneous increase in the path loss of up to 5dB can be tolerated with no loss of data, and gives power control the ability to increase the CPE’s transmit power before data is lost. The 5dB default fade margin means that the standard US CNR targets, at the AP, are:

Target CNR at AP	10.5, 26 and 28 GHz
256QAM+T	30.3 dB
128QAM+T	29.3 dB
64QAM	27.3 dB
64QAM+T	23.3 dB
16QAM	20.2 dB
16QAM+T	16.2 dB
QPSK	13.4 dB

DOWNSTREAM FADING

The CPE can accurately measure the downstream path loss by comparing the received downstream power with the AP's Tx power (the AP broadcasts the value of its actual Tx power to all CPEs).

The AP transmits at constant power and the CPEs use AGC to adjust for their individual path loss. There is an implicit fade margin which is the delta between the minimum required signal power for the modulation in use and the actual received power from the AP. There is 6dB of instantaneous headroom and 5dB of margin (assuming the implicit fade margin is at least 5dB) built in to the digital demodulator – fade rates of 100 dB/sec can be tolerated within this range. Outside this range, AGC is required and the CPE can independently adjust its AGC at approximately 6Hz.

UPSTREAM FADING

For small fades, the digital demodulator in the AP can cope without requiring the CPE to adjust its Tx power. There is 3dB to 15dB (3dB for 64QAM, 9dB for 16 QAM, 15dB for QPSK) of instantaneous headroom built in to the digital demodulator, and power control is set to default with 5dB of fade margin (i.e. fading rates of 100 dB/sec can be tolerated within this range). Outside this range, ATPC is required.

For ATPC to operate, the CPE must know the upstream path loss and as such must wait until it is polled by the AP. (Polling is used to schedule an upstream channel sounding.) The resultant Rx power at the AP is then sent back to the CPE and with this the CPE can correct its Tx power. The rate of polling is dynamic, based on the best judgement of the AP – it assigns more resources to CPEs with the most troublesome radio channels. The range for upstream polling for power control is approximately 0.2Hz to 8Hz.

For dynamic channels, the problem is accurate prediction of upstream path loss over time. An exact calculation can only be guaranteed once every 5 seconds. Consequently there may be periods when the CPE does not posses any information about changes in the upstream path loss, especially if the AP is heavily loaded and cannot poll as quickly as it would like. To get around this, the CPE has two power control modes.

- Correlation mode – The CPE constantly correlates upstream and downstream path losses. If the correlation is high then the CPE will use the downstream path loss to predict the upstream path loss in between polls from the CPE. This allows the upstream power control algorithm to track large changes in the order of 20dB/sec in correlated channels with negligible performance penalty.
- Back-off mode – If the upstream and downstream path losses are poorly correlated, the CPE uses a 'safe' mode where the Tx power is backed off sufficiently to ensure that the received power at the AP does not exceed the target power (for example if the upstream channel is fading by 3dB then the CPE will set the transmit power such that the peak received power at the AP equals the AP target power, because the worst case headroom is 3dB for 64 QAM).

ASYMMETRIC CNR

A CPE quite close to the AP, running a QPSK VC would show a high downstream CNR, e.g. 30dB + but only 18dB CNR on the upstream. For example, the following screen capture from *vsstatus* shows CPE Damian with 30.6dB CNR downstream and 18.7dB CNR upstream.

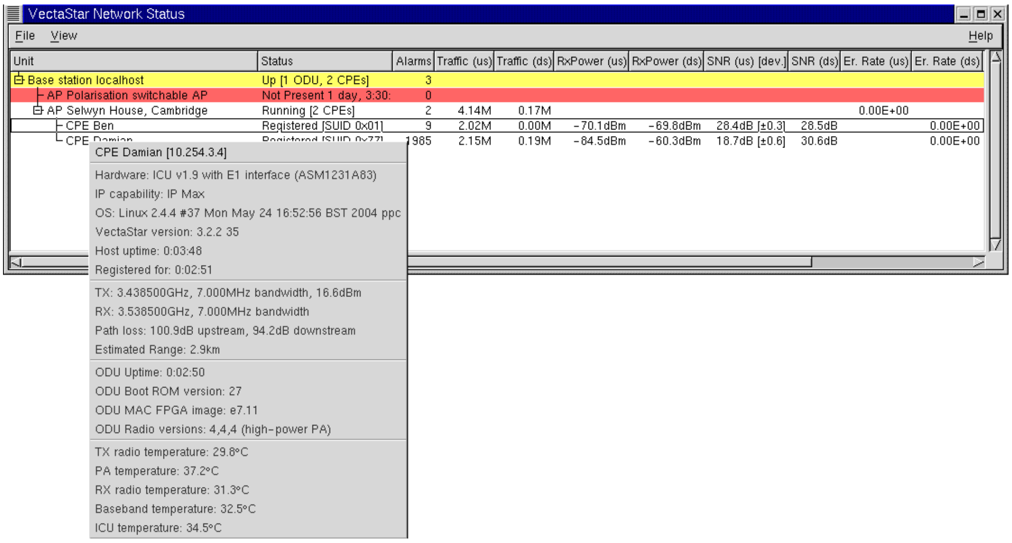

This is because the CPE has set its transmit power such that it meets the QPSK upstream target power at the AP (i.e. 18dB CNR). Right-clicking on the CPE name shows the Tx power 16.6dBm, significantly below the maximum possible for this CPE (which is 34dBm). If the CPE is reconfigured for a 64 QAM VC, the upstream Rx power and CNR will increase accordingly as the CPE increases its transmit power to meet the 64 QAM CNR target. In this case, the CPE can increase its Tx power by up to 17.4dB, which should increase the US CNR from 18.7dB to in excess of 30dB (i.e. good enough for 64 QAM operation).

Note: If a CPE is configured to use a high order modulation, then it will adjust its Tx power accordingly. However, the CPE only finds out which modulation to use once it has downloaded its VC configurations from the APC, which takes place after registration. So a step change may be seen in the uplink Rx power and CNR shortly after registration.

The only other thing that would cause this is asymmetric path loss, which is quite rare. The VNMS, in software release 3.2.2 and later, displays upstream and downstream path loss in *vsstatus*, and as long as these figures are fairly close (i.e. within a few dB), symmetric CNR performance should be achievable.

In the event that the US and DS path losses differ by more than 8dB (which is unusual and indicative of poor deployment or damaged radio or antenna) then an SNMP trap/alarm is generated automatically to alert the user to the potential problem.

INTERPRETING POWER AND CNR FIGURES

The different modulation schemes have different PA linearity requirements and cannot all be transmitted at the same power level through the PAs. Therefore a CPE or AP receiving for example QPSK, 16 QAM and 64 QAM data will see three distinct receive power levels with relative levels of 0dB, -2.7dB and -4.3dB for QPSK, 16 QAM and 64 QAM respectively.

To simplify the diagnostics and monitoring all power and CNR figures displayed by VectaStar are referenced to QPSK regardless of the modulations in use. Therefore a CNR figure of 30dB is only 30dB for QPSK which is equivalent to 27.3dB for 16 QAM and 25.7dB for 64 QAM.

RUNNING OUT OF POWER

If a CPE is transmitting at maximum power an SNMP trap/alarm is generated to inform the operator that the CPE is potentially power limited. A CPE will transmit at maximum power if the uplink power control mechanism determines that a transmit power at or above its maximum power is required to meet the target upstream power at the AP.

In the case where the path loss is very gradually increasing between the CPE and the AP (e.g. due to foliage growth), then at the point where the CPE first reaches maximum transmit power there is still a 5dB fade margin available – service has not been compromised. However, as the path loss increases further, the ability of the CPE to cope with fading will be diminished and, after a 5dB extra path loss, the link may become unreliable.

There are several methods of increasing the uplink link budget:

- Use a lower order modulation
- Use a narrower channel bandwidth
- Use a larger antenna (e.g. parabolic dish)
- Re-site the antenna to reduce any obstruction or diffraction

ERROR CORRECTION

VectaStar uses Automated Repeat Query (ARQ) for on-air error correction rather than Forward Error Correction (FEC). ARQ has several major advantages over FEC for Point to Multi Point (PMP) radio systems:

- ARQ can achieve a higher code rate than FEC
- ARQ can operate much closer to sensitivity than FEC
- ARQ dynamically changes the code rate to match the channel error rate
- ARQ achieves an individual code rate for each radio link

With Forward Error Correction the code rate is fixed and the Rx signal level can be adjusted to affect CNR and for each CNR level there will be a corresponding post FEC BER.

VectaStar uses Automated Repeat Query (ARQ) rather than FEC for error correction as it offers significant spectral efficiency savings in a PMP environment. This is because with PMP the downlink is broadcast but the radio link to each CPE is over an independent radio path. At any moment in time, each link requires a unique code rate, for optimum performance; where optimum performance is defined as the highest code rate (most efficient) for the required target BER. ARQ offers a better solution as the code rate is dynamically adjusted on the fly to each CPE to achieve a target post ARQ BER which is defined for all links. ARQ operates by sending each packet across the air interface with a small CRC. The CPE uses the CRC to determine if the packet was received with errors or not. If the packet is received error free it is acknowledged and no more data is sent regarding that packet. If it was received in error or not at all, a Negative Acknowledgement is sent and the packet resent until it is received error free or times out. This method always achieves the optimum code rate for each link at any moment in time.

ARQ operates by sending un-coded data in packets over the air, with each packet containing a CRC check-sum. If the packet is received in error, determined by the CRC check-sum, then the packet is either Negatively acknowledged (NACK) or simple not acknowledged at all (if it is so corrupted as to be unrecognisable).

In either event the sender will re-transmit that packet until it is positively acknowledged (ACK) or the packet lifetime has expired (e.g. if a CPE suddenly turns off then there is no point in endlessly sending the same cell hoping for an ACK).

VectaStar's ARQ implementation has been designed to achieve Quasi-error free end-to-end BER (< 10-9) with an on-air BER of 10-4. If the on-air BER is greater than 10-4 then the system may not be able to provide error free transmission.

ARQ VS FEC

With Forward Error Correction the code rate is fixed and the Rx signal level can be adjusted to affect CNR and for each CNR level there will be a corresponding post FEC BER.

VectaStar uses ARQ over FEC due to its superior security advantages. ARQ is fundamentally much harder to passively eavesdrop compared with FEC as with FEC all the parity information is sent with the data, so the end unit can independently reconstruct the original data sequence. With ARQ, the on-air error correction relies on the unique interaction between the CPE and AP for the appropriate data to be resent on demand. The passive eavesdropper cannot request data to be resent and as it will undoubtedly receive different packets in error to the valid CPE the eavesdropping CPE cannot gain any benefit from error correction.

In addition to traditional error detection and correction techniques, VectaStar now also has additional error detection and correction by using Trellis encoding. Trellis coding uses defined encoding data paths in order to allow error detection and correction without re-transmissions.

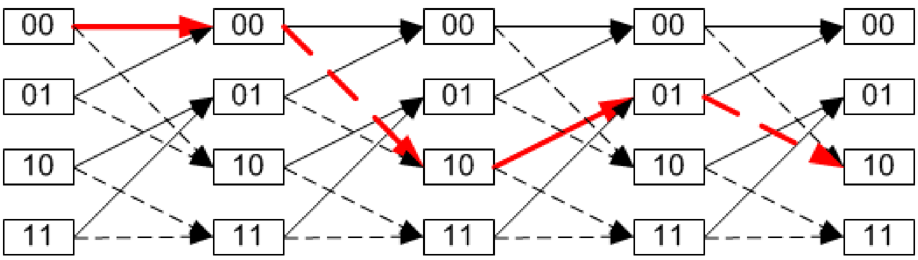

In order to use TCM there is a slight decrease (approximately .5 b/s/Hz) in spectral efficiency however this is amply compensated for by a large decrease in retransmissions. The use of TCM allows a lower SNR to be used as a higher radio BER can be accommodated.

TECHNICAL SPECIFICATIONS

SYSTEM SPECIFICATIONS

Backhaul Features

ATM/CES/IMA support

Dynamic IP and E1 bandwidth assignment (Statistical multiplexing)

Full QoS with PIR and CIR traffic shaping

Up to 4 E1s per CPE

Optimisation for E1, HDLC, Abis, and IuB services

BS and CPE redundancy

Access Features

Full configurable L2 and L3 QoS with 802.1p/q, MAC address filtering

E1/T1 full (G.703), fractional (G.704), unstructured services

Self backhaul services

Environmental

Outdoor equipment	Operational temperature range	-45°C to +55°C[7] (upper limit to be derated by 1°C for every 300m of altitude above sea level)
	Operational Humidity range	0 – 100%
Indoor equipment	Operational temperature range	+5°C to +40°C (upper limit to be derated by 1°C for every 300m of altitude above sea level)
	Operational Humidity range	0 – 85%

[7] According to ETSI EN 300 019-1-4 Class 4.1E, 4.2H. High temperature test conditions are at +70°C to simulate solar loading

Power Consumption		
Base Station (Redundant, 4 sectors)	-40.5V to –57V DC	280W typical
Terminal Station with IDU	-40.5V to –57V DC	45 W typical

STANDARDS COMPLIANCE

Type	Standard	
CE	All applicable directives including R&TTE	
Radio & Antenna	ETSI EN 302 326-2 and ETSI EN 302 326-3	
Earthing & Bonding	ETSI EN 300 253	
EMC	ETSI EN 301 489-1 and ETSI EN 301 489-4	
Operational environmental	Outdoor	ETSI EN 300 019-1-4 Class 4.1E, 4.2H
	Indoor	ETSI EN 300 019-1-3 Class 3.1
Mechanical Finish	All mechanical components and surface finishes will comply with EU RoHS Directive 2002/95/EC	
Safety	EN 60950-1 and 60950-22	
Storage environment	ETSI EN 300 019-1-1 class 1.3	
Transport environment	ETSI EN 300 019-1-2 class 2.3	
DC Power Supply	VectaStar operates within SELV limits and ETSI EN 300 132-2	

10.5 GHz Capacity	
Spectral efficiency	7.5 bit/sec/Hz (gross) without E1 optimisation
Throughput per AP	Up to 87.5Mbps @ 14MHz (gross) Up to 69Mbps @ 14MHz (Ethernet)
Throughput per CPE	Up to 87.5Mbps @ 14MHz (gross) Up to 69Mbps @ 14MHz (Ethernet)
Base Station Capacity	Up to five sectors (5 APs) and up to a total of 100 CPEs
Statistical multiplexing	Dynamic BW allocation for all services

26 and 28 GHz Capacity	
Spectral efficiency	Up to 7.5 bit/sec/Hz (gross) without E1 optimisation
Throughput per AP	175 Mbps @ 28MHz (gross) 132Mbps @ 28MHz (Ethernet)
Throughput per CPE	0 – 90 Mbps @ 28MHz (Note: rate is limited by CPE 100BaseT interface)
Base Station Capacity	Up to six sectors (6 APs)
Statistical multiplexing	Dynamic BW allocation for all services

APPENDIX A – 10.5GHZ SPECIFICATIONS

10.5 GHZ RADIO AND MODEM

Parameter	Value
Frequency of operation	10.15 to 10.65 GHz
Standards Conformance	ITU-R F.1568 and CEPT ERC T/R 12-05E
Duplex operation	FDD Full Duplex
Duplex separation	350 MHz
Duplex filter selection	AP 150MHz CPE 150MHz
Duplex filter loss	2 dB
Radio Access Method	FDD with TDMA in both uplink and downlink
Transmission technique	Single Carrier
Channel tuning	integer multiples of 250 kHz
Modulation	256QAM+T, 128QAM+T, 64QAM, 64QAM+T, 16QAM, 16QAM+T, QPSK uplink and downlink, fully adaptive Modulations with +T operate with Trellis Code Modulation. Adaptive Code Modulation uses all states.
RF Channel sizes	14 MHz, 7 MHz
Latency	IP Average < 0.6ms with 99.9% < 1.0ms at <=98% rated sector throughput E1 10ms, configurable jitter buffer
Jitter	Sub 1 ms
Power Control	DS: statically set via VNMS US: automatic transmit power control
Spectral mask	as per ETSI EN 302 326-2
Receive sensitivity	as per ETSI EN 302 326-2
Output power (pre duplexer)	AP 20 dBm CPE 20 dBm
Receiver sensitivity at antenna interface	7MHz channel QPSK -88.4 dBm 14MHz channel 256QAM+T -64.5 dBm

10.5 GHZ DUPLEXER SPACING

VectaStar is a FDD PMP system and therefore uses separate frequency domains to transmit and receive radio signals in a full duplex fashion. Every VectaStar ODU has a RF duplexer to isolate the receiver from the transmitter while permitting them to share the common antenna.

The figure below shows how VectaStar uses a single duplexer to cover the 10.5 GHz ETSI band. The standard configuration is shown with downlink signals on the high side of the duplexer. Downlink refers to AP → CPE direction whereas Uplink refers to CPE → AP. This duplexer configuration is referred to as Downlink High (DH). VectaStar ODUs can also be supplied as Downlink Low (DL) according to customer requirements.

The pricelist order codes listed below correspond as follows:

- DH: Downlink High
- DL: Downlink Low

These duplexer codes aim to simplify the ordering process such that the same duplexer code is used for AP and CPE.

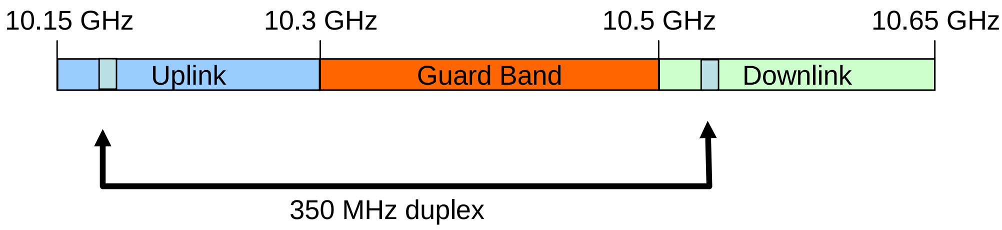

FIGURE 28: 10.5 GHZ DUPLEXER DESIGN

10.5 GHZ ODU DIMENSIONS

ODUs are designed with a "slip-fit" waveguide interface allowing the same ODU to be used with different polarisation antennas. Note: AP ODUs and CPE ODUs are NOT interchangeable.

Polarisation changes are performed by reconfiguring the slip-fit waveguide on the ODU along with using the appropriate antenna.

10.5 GHZ AP ODU DIMENSIONS

The AP ODU dimensions, not including slip-fit antenna are shown below:

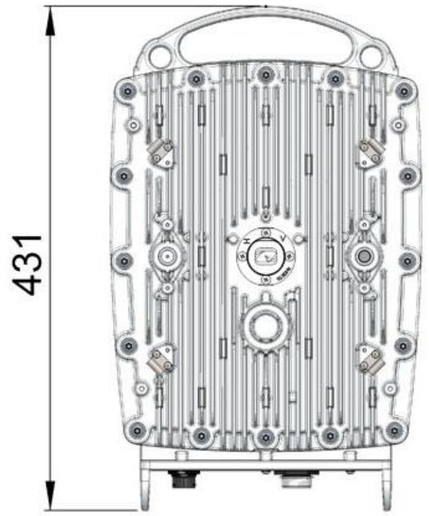

FIGURE 29: 10.5 GHZ AP ODU ANTENNA SIDE

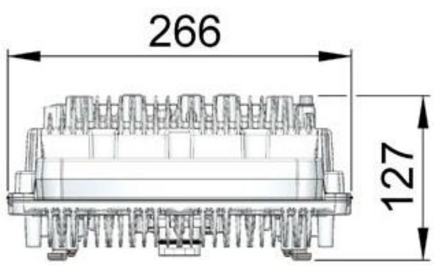

FIGURE 30: 10.5 GHZ AP ODU FROM TOP

10.5 GHZ CPE ODU DIMENSIONS

The CPE ODU dimensions, not including slip-fit antenna are shown below:

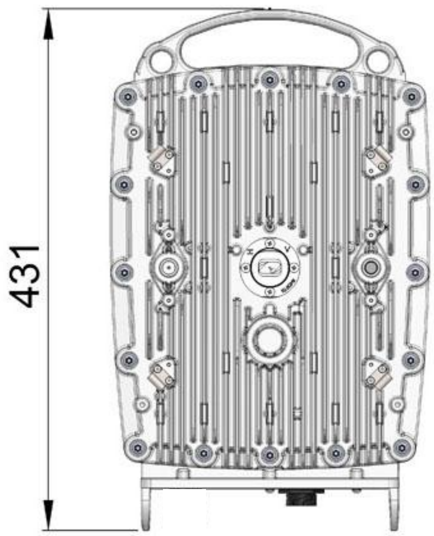

FIGURE 31: 10.5 GHZ CPE ODU ANTENNA SIDE

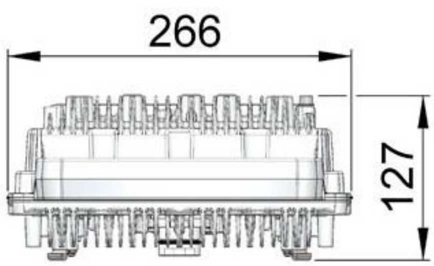

FIGURE 32: 10.5 GHZ CPE ODU FROM TOP

10.5 GHZ ANTENNAS

All 10.5 GHz AP antennas are waveguide horns and are linearly polarised of which either horizontally or vertically polarised horns are available. For 10.5 GHz CPEs, there are two types of slip-fit parabolic antenna: 1-foot (30cm) and 2-foot (60cm) dishes. Each antenna comes with pole mounting brackets to which the AP or CPE ODU attaches. Instructions for antenna mounting and adjustment are provided with the antenna as shipped.

Cambridge Broadband Networks Limited (CBNL) sources antennas from multiple suppliers to reduce antenna order intervals as well as to be able to make available antenna options such as new sizes and destination countries. Additional antenna options are available upon request – please contact your sales representative for more information.

AP ANTENNA SPECIFICATIONS.

Order Code		ANT-10G-AP-H	ANT-10G-AP-V
Frequency		10.15-10.65 GHz	10.15-10.65 GHz
Polarization		Horizontal	Vertical
Minimum Gain (dBi)		16	16
3dB BW, Typical (@centre frequency)		90° x 8°	90° x 8°
Minimum F/B Ratio (dB)		40	40
Waveguide interface		Slip-fit	Slip-fit
Dimensions (maximum)	Height (mm)	355	339
	Width (mm)	407	407
	Depth (mm)	302	272
Pole size range	Minimum	38mm (1.5")	38mm (1.5")
	Maximum	115mm (4.5")	115mm (4.5")
Maximum net weight (kg)		9.2	8.6
Elevation adjustment (min)		Fine +/-30°	Fine +/-30°
Azimuth adjustment (min)		Coarse 360°, Fine +/-30°	Coarse 360°, Fine +/-30°

PARABOLIC ANTENNA SPECIFICATIONS.

Order Code		ANT-10G-PSF-30	ANT-10G-PSF-60
Frequency		10.15-10.65 GHz	10.15-10.65 GHz
Polarization		V or H	V or H
Diameter		0.3m (1-ft)	0.6m (2-ft)
Minimum Gain (dBi)		26.8	33.7
3dB BW, Typical (@centre frequency)		6.5°	3.6°
Minimum F/B Ratio (dB)		42	54
Minimum Cross Polarisation Discrimination (dB)		25	30
VSWR, Max (R.L.)		1.50:1	1.50:1
Waveguide interface		Slip-fit	Slip-fit
Dimensions (maximum)	Height (mm)	353	663
	Width (mm)	517	689.5
	Depth (mm)	286	335
Pole size range	Minimum	38mm (1.5")	48mm (1.9")
	Maximum	115mm (4.5")	115mm (4.5")
Maximum net weight (kg)		8.5	14
Elevation adjustment (min)		±25°	±25°
Azimuth adjustment (min)		±10°	±10°
Wind velocity survival rating (km/h)		210	210

APPENDIX B – 26 GHZ SPECIFICATIONS

26 GHZ RADIO AND MODEM

Parameter	Value
Frequency of operation	24.5 to 26.5 GHz
Standards Conformance	ITU-R F.748-4 Annex 1 and CEPT ERC T/R 13-02E
Duplex operation	FDD Full Duplex
Duplex separation	1008 MHz
Duplex filter selection	Band specific according to customers allocation
Duplex filter loss	nominally 0.5 dB
Radio Access Method	FDD with TDMA in both uplink and downlink
Transmission technique	Single Carrier
Channel tuning	integer multiples of 250 kHz
Modulation	256QAM+T, 128QAM+T, 64QAM, 64QAM+T, 16QAM, 16QAM+T, QPSK uplink and downlink[9], fully adaptive (all states)
RF Channel sizes	28 MHz, 14 MHz, 7 MHz
Latency	IP Average < 0.6ms with 99.9% < 1.0ms at <=98% rated sector throughput E1 10ms, configurable jitter buffer
Jitter	Sub 1 ms
Power Control	DS: statically set via VNMS US: automatic transmit power control (ATPC)
Spectral mask	as per ETSI EN 302 326-2
Receive sensitivity	as per ETSI EN 302 326-2
Output power (pre duplexer)	AP 18 dBm CPE 18 dBm
Receiver sensitivity at antenna interface	7MHz channel QPSK -86.1 dBm 28MHz channel 256QAM+T -63.2 dBm

[9] Modulations with +T operate with Trellis Code Modulation.

26 GHZ DUPLEXER SPACING

VectaStar is a FDD PMP system and therefore uses separate frequency domains to transmit and receive radio signals in a full duplex fashion. Every VectaStar ODU has a RF duplexer to isolate the receiver from the transmitter while permitting them to share the common antenna.

The figure below shows how VectaStar uses two duplexers to cover the 26 GHz ETSI band. The pass bands of the high band and low band duplexers intentionally overlap as shown so that the majority of customer frequency allocations can be served with one type of VectaStar hardware. Customers should check their frequency allocations in order to select the correct VectaStar equipment for their use.

The pricelist order codes listed below correspond as follows (Downlink refers to AP → CPE direction):

- LBDH: Low Band, Downlink High
- LBDL: Low Band, Downlink Low
- HBDH: High Band, Downlink High
- HBDL: High Band, Downlink Low

These duplexer codes aim to simplify the ordering process such that the same duplexer code is used for AP and CPE.

Duplexer Code		Low Band	High Band		Low Band	High band	
LBDH		CPE Tx (AP Rx)			**AP Tx (CPE Rx)**		
LBDL		**AP Tx (CPE Rx)**			CPE Tx (AP Rx)		
HBDH			CPE Tx (AP Rx)			**AP Tx (CPE Rx)**	
HBDL			**AP Tx (CPE Rx)**			CPE Tx (AP Rx)	

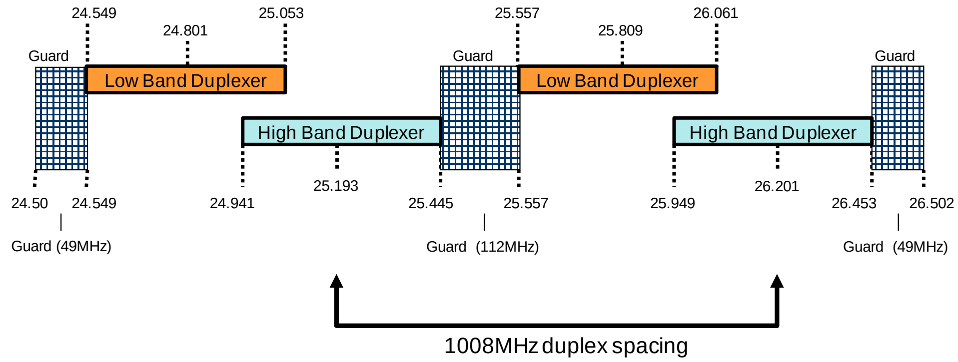

FIGURE 33: 26 GHZ DUPLEXER DESIGN

26 GHZ ODU DIMENSIONS

ODUs are designed with a “slip-fit” waveguide interface allowing the same ODU to be used with different polarisation antennas. Note: AP ODUs and CPE ODUs are NOT interchangeable.

Polarisation changes are performed by reconfiguring the slip-fit waveguide on the ODU along with using the appropriate antenna.

26 GHZ AP ODU DIMENSIONS

The AP ODU dimensions, not including slip-fit antenna are shown below.

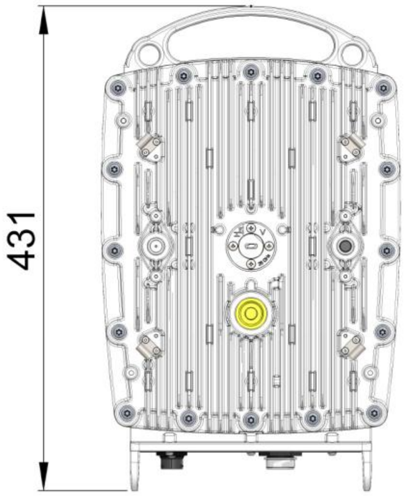

FIGURE 34: 26 GHZ AP ODU ANTENNA FRONT SIDE

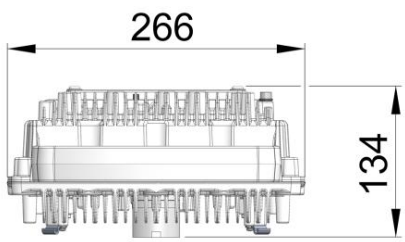

FIGURE 35: 26 GHZ AP ODU FROM TOP

26 GHZ CPE ODU DIMENSIONS

The CPE ODU dimensions, not including slip-fit antenna are shown below.

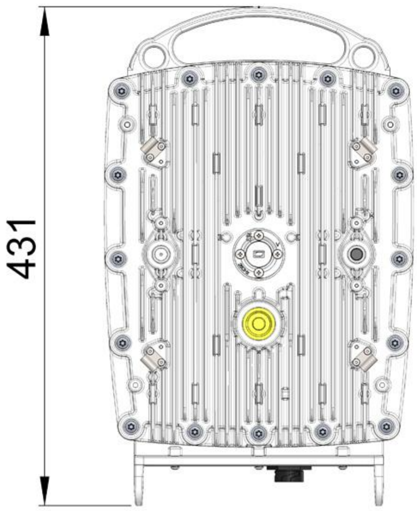

FIGURE 36: 26 GHZ CPE ODU ANTENNA SIDE

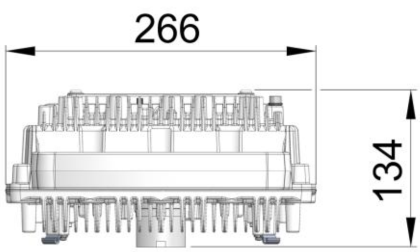

FIGURE 37: 26 GHZ CPE ODU FROM TOP

26 GHZ ANTENNAS

All 26 GHz AP antennas are waveguide horns and are linearly polarised of which either horizontally or vertically polarised horns are available. For 26 GHz CPEs, there are two types of slip-fit parabolic antenna: 1-foot (30cm) and 2-foot (60cm) dishes. Each antenna comes with pole mounting brackets to which the AP or CPE ODU attaches. Instructions for antenna mounting and adjustment are provided with the antenna as shipped.

Cambridge Broadband Networks Limited (CBNL) sources antennas from multiple suppliers to reduce antenna order intervals as well as to be able to make available antenna options such as new sizes and destination countries. Additional antenna options are available upon request – please contact your sales representative for more information.

AP ANTENNA SPECIFICATIONS.

Order Code		ANT-26G-AP-H	ANT-26G-AP-V
Frequency		24.25-26.50 GHz	24.25-26.50 GHz
Polarization		Horizontal	Vertical
Minimum Gain (dBi)		18	18
3dB BW, Typical (@centre frequency)		90° x 6°	90° x 6°
Minimum F/B Ratio (dB)		40	40
Waveguide interface		Slip-fit	Slip-fit
Dimensions (maximum)	Height (mm)	293	293
	Width (mm)	420	420
	Depth (mm)	333	333
Pole size range	Minimum	38mm (1.5")	38mm (1.5")
	Maximum	115mm (4.5")	115mm (4.5")
Maximum net weight (kg)		11	11
Elevation adjustment (min)		Fine +/-30°	Fine +/-30°
Azimuth adjustment (min)		Coarse 360°, Fine +/-30°	Coarse 360°, Fine +/-30°

PARABOLIC ANTENNA SPECIFICATIONS.

Order Code		ANT-26G-PSF-30	ANT-26G-PSF-60
Frequency		24.25-26.50 GHz	24.25-26.50 GHz
Polarization		V or H	V or H
Diameter		0.3m (1-ft)	0.6m (2-ft)
Minimum Gain (dBi)		35.3	40.7
3dB BW, Typical (@centre frequency)		2.5°	1.4°
Minimum F/B Ratio (dB)		62	68
Minimum Cross Polarisation Discrimination (dB)		30	30
VSWR, Max (R.L.)		1.50:1	1.50:1
Waveguide interface		Slip-fit	Slip-fit
Dimensions (maximum)	Height (mm)	389	663
	Width (mm)	552.5	689.5
	Depth (mm)	220	456
Pole size range	Minimum	50mm (2.0")	48mm (1.9")
	Maximum	115mm (4.5")	115mm (4.5")
Maximum net weight (kg)		6.8	14
Elevation adjustment (min)		±25°	±25°
Azimuth adjustment (min)		±10°	±10°
Wind velocity survival rating (km/h)		210	210

APPENDIX C – 28 GHZ SPECIFICATIONS

28 GHZ RADIO AND MODEM

Parameter	Value
Frequency of operation	27.5 to 29.5 GHz
Standards Conformance	ITU-R F.748-4 Annex 2 and CEPT ERC T/R 13-02E
Duplex operation	FDD Full Duplex
Duplex separation	1008 MHz
Duplex filter selection	Band specific according to customers allocation
Duplex filter loss	nominally 0.5 dB
Radio Access Method	FDD with TDMA in both uplink and downlink
Transmission technique	Single Carrier
Channel tuning	integer multiples of 250 kHz
Modulation	256QAM+T, 128QAM+T, 64QAM, 64QAM+T, 16QAM, 16QAM+T, QPSK uplink and downlink [10], fully adaptive (all states)
RF Channel sizes	28 MHz, 14 MHz, 7 MHz
Latency	IP Average < 0.6ms with 99.9% < 1.0ms at <=98% rated sector throughput E1 10ms, configurable jitter buffer
Jitter	Sub 1 ms
Power Control	DS: statically set via VNMS US: automatic transmit power control (ATPC)
Spectral mask	as per ETSI EN 302 326-2
Receive sensitivity	as per ETSI EN 302 326-2
Output power (pre duplexer)	AP 18 dBm CPE 18 dBm
Receiver sensitivity at antenna interface	7MHz channel QPSK -86.1 dBm 28MHz channel 256QAM+T -63.2 dBm

[10] Modulations with +T operate with Trellis Code Modulation.

28 GHZ DUPLEXER SPACING

VectaStar is a FDD PMP system and therefore uses separate frequency domains to transmit and receive radio signals in a full duplex fashion. Every VectaStar ODU has a RF duplexer to isolate the receiver from the transmitter while permitting them to share the common antenna.

The figure below shows how VectaStar uses two duplexers to cover the 28GHz ETSI band. The pass bands of the high band and low band duplexers intentionally overlap as shown so that the majority of customer frequency allocations can be served with one type of VectaStar hardware. Customers should check their frequency allocations in order to select the correct VectaStar equipment for their use.

The pricelist order codes listed below correspond as follows (Downlink refers to AP → CPE direction):

- LBDH: Low Band, Downlink High
- LBDL: Low Band, Downlink Low
- HBDH: High Band, Downlink High
- HBDL: High Band, Downlink Low

These duplexer codes aim to simplify the ordering process such that the same duplexer code is used for AP and CPE.

Order Code	Low Band	High Band	Low Band	High band
LBDH	CPE Tx (AP Rx)		**AP Tx (CPE Rx)**	
LBDL	**AP Tx (CPE Rx)**		CPE Tx (AP Rx)	
HBDH		CPE Tx (AP Rx)		**AP Tx (CPE Rx)**
HBDL		**AP Tx (CPE Rx)**		CPE Tx (AP Rx)

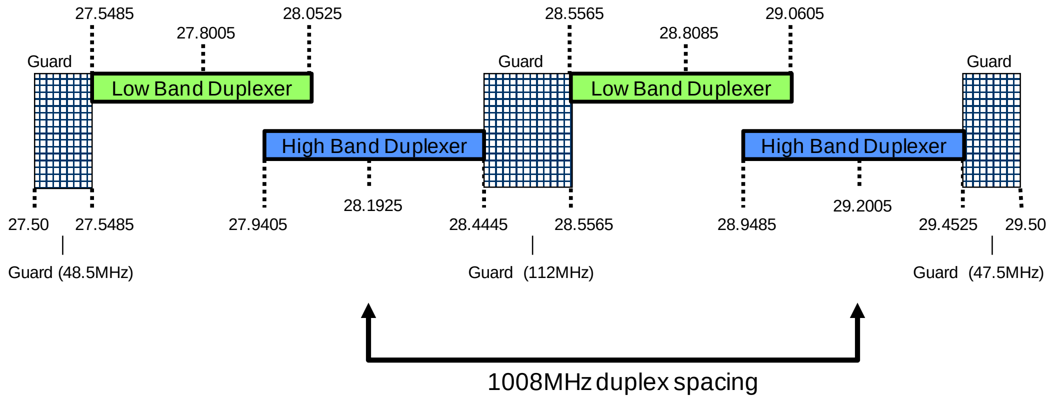

FIGURE 38: 28 GHZ DUPLEXER DESIGN

28 GHZ ODU DIMENSIONS

ODUs are designed with a "slip-fit" waveguide interface allowing the same ODU to be used with different polarisation antennas. Note: AP ODUs and CPE ODUs are NOT interchangeable.

Polarisation changes are performed by reconfiguring the slip-fit waveguide on the ODU along with using the appropriate antenna.

28 GHZ AP ODU DIMENSIONS

The AP ODU dimensions, not including slip-fit antenna are shown below:

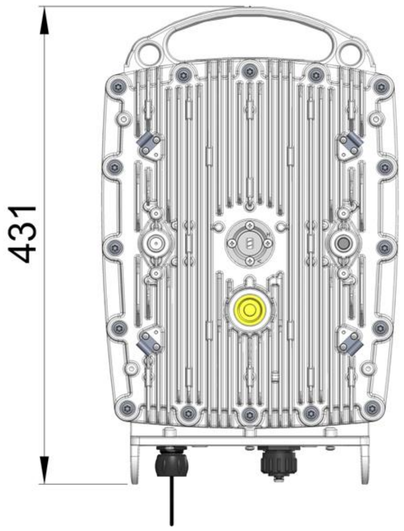

FIGURE 39: 28 GHZ AP ODU ANTENNA SIDE

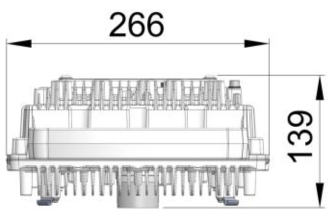

FIGURE 40: 28 GHZ AP ODU FROM TOP

28 GHZ CPE ODU DIMENSIONS

The CPE ODU dimensions, not including slip-fit antenna are shown below:

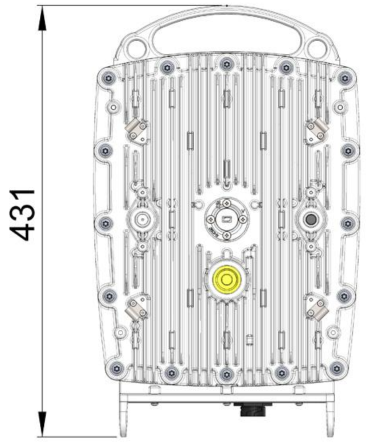

FIGURE 41: 28 GHZ CPE ODU ANTENNA SIDE

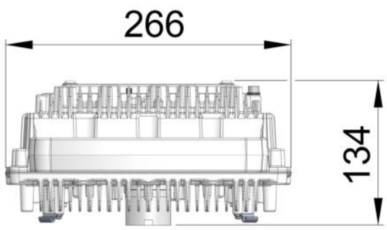

FIGURE 42: 28 GHZ CPE ODU FROM TOP

28 GHZ ANTENNAS

All 28 GHz AP antennas are waveguide horns and are linearly polarised of which either horizontally or vertically polarised horns are available. For 28 GHz CPEs, there are two types of slip-fit parabolic antenna: 1-foot (30cm) and 2-foot (60cm) dishes. Each antenna comes with pole mounting brackets to which the AP or CPE ODU attaches. Instructions for antenna mounting and adjustment are provided with the antenna as shipped.

Cambridge Broadband Networks Limited (CBNL) sources antennas from multiple suppliers to reduce antenna order intervals as well as to be able to make available antenna options such as new sizes and destination countries. Additional antenna options are available upon request – please contact your sales representative for more information.

1.1.1 AP ANTENNA SPECIFICATIONS.

Order Code		ANT-28G-AP-H	ANT-28G-AP-V
Frequency		27.3-29.5 GHz	27.3-29.5 GHz
Polarization		Horizontal	Vertical
Minimum Gain (dBi)		18	18
3dB BW, Typical (@centre frequency)		90° x 6°	90° x 6°
Minimum F/B Ratio (dB)		40	40
Waveguide interface		Slip-fit	Slip-fit
Dimensions (maximum)	Height (mm)	293	293
	Width (mm)	420	420
	Depth (mm)	333	333
Pole size range	Minimum	38mm (1.5")	38mm (1.5")
	Maximum	115mm (4.5")	115mm (4.5")
Maximum net weight (kg)		11	11
Elevation adjustment (min)		Fine +/-30°	Fine +/-30°
Azimuth adjustment (min)		Coarse 360°, Fine +/-30°	Coarse 360°, Fine +/-30°

1.1.2 PARABOLIC ANTENNA SPECIFICATIONS.

Order Code		ANT-28G-PSF-30	ANT-28G-PSF-60
Frequency		27.30-29.50 GHz	27.30-29.50 GHz
Polarization		V or H	V or H
Diameter		0.3m (1-ft)	0.6m (2-ft)
Minimum Gain (dBi)		35.8	41.8
3dB BW, Typical (@centre frequency)		2.2°	1.3°
Minimum F/B Ratio (dB)		64	68
Minimum Cross Polarisation Discrimination (dB)		30	30
VSWR, Max (R.L.)		1.50:1	1.50:1
Waveguide interface		Slip-fit	Slip-fit
Dimensions (maximum)	Height (mm)	389	663
	Width (mm)	552.5	689.5
	Depth (mm)	220	456
Pole size range	Minimum	50mm (2.0")	48mm (1.9")
	Maximum	115mm (4.5")	115mm (4.5")
Maximum net weight (kg)		6.8	14
Elevation adjustment (min)		±25°	±25°
Azimuth adjustment (min)		±10°	±10°
Wind velocity survival rating (km/h)		210	210

APPENDIX D RECOMMENDED CABLE TYPES

CAT-5E CABLES.

All Cat-5e cables are wired as straight through connections.

All outdoor cables must be solid core at least 24AWG (.25mm^2) 8 wires in total and able to terminate into a standard RJ-45 connector.

2-CORE CABLES.

All 2-core cables must be outdoor grade a minimum of 20AWG (0.52 mm^2) and a maximum of 14 AWG (2.5mm^2) with a cable sheath diameter size of no less than 3mm and no more than 12mm.

FIBRE OPTIC CABLES.

All fibre optic cables are supplied by CBNL as standard parts.

VectaStar fibre optic cables are multimode (62.5/125µm) terminated into LC type connectors at the AP end which has an IP 67 rated plug ready attached and SC type connectors at the indoor end. The entire outdoor portion of the cable is coated in an external grade, waterblocking, UV stable outer jacket making it 5.3mm in diameter.

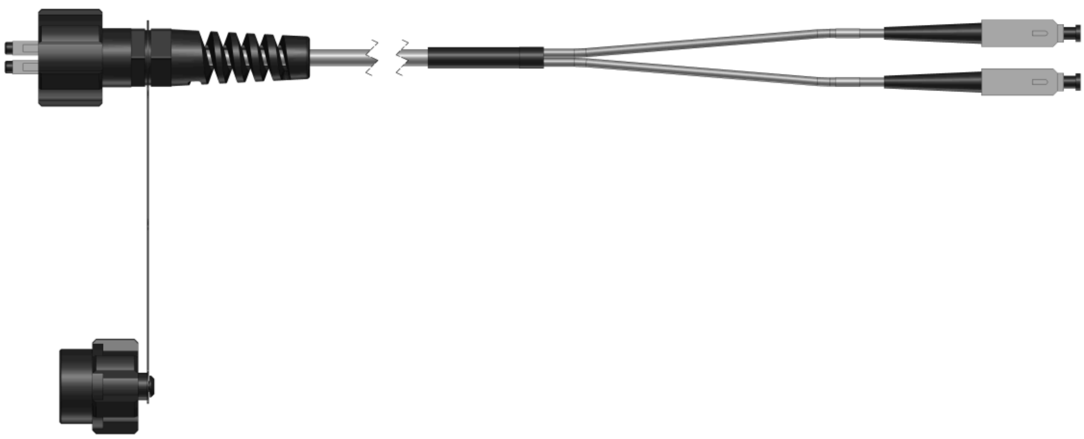

FIGURE 63: GRAPHICAL REPRESENTATION OF VS AP OPTICAL FIBRE CABLE.

N.B. CABLES MAY BE COLOURED.

APPENDIX E REFERENCE DESIGN GUIDE

INTRODUCTION

This section describes the most common (standard) VectaStar designs. In most cases, the network design deployed will be identical to those described herein and where the design differs, clarification should be sought from Network Architecture and / or Product Management to ensure that the proposed design is both supported and suitable for the intended network.

Using VectaStar Gigabit (VSG) RTs in place of VS2 CPEs.

VSG RTs have an improved functional specification over VS2 RTs, for example Gigabit Ethernet interface and voltage alignment facility. VSG RTs are compatible with VS2 hubs and 1U IDUs and provide an excellent means to migrate to the higher performance VSG platform. Therefore CBNL recommends that VSG RTs are purchased in preference to VS2 CPEs.

The following sections include details on how to install a VSG RT in a VS2 network and lists the common items that are required for a full installation. Where order codes and equipment lists show a # this indicates a variable according to the latest price list. For full details on VSG RTs and other VSG equipment, please refer to the VectaStar Gigabit Technical Specification Document.

REMOTE TERMINALS

This section describes VSG RT configurations. There are four permutations considered:

Ethernet only, 230v ac Mains power

Ethernet only, -48v dc power

Ethernet and E1s, 230v ac Mains power

Ethernet and E1s, -48v dc power

VSG RT ETHERNET ONLY (MAINS POWERED)

For an RT with lightning protection, the following equipment is required.

Note: This does not include installation materials e.g. feeder clamps, earth cables, nor does it include installation tools e.g. alignment tools or voltmeter etc.

DIAGRAM

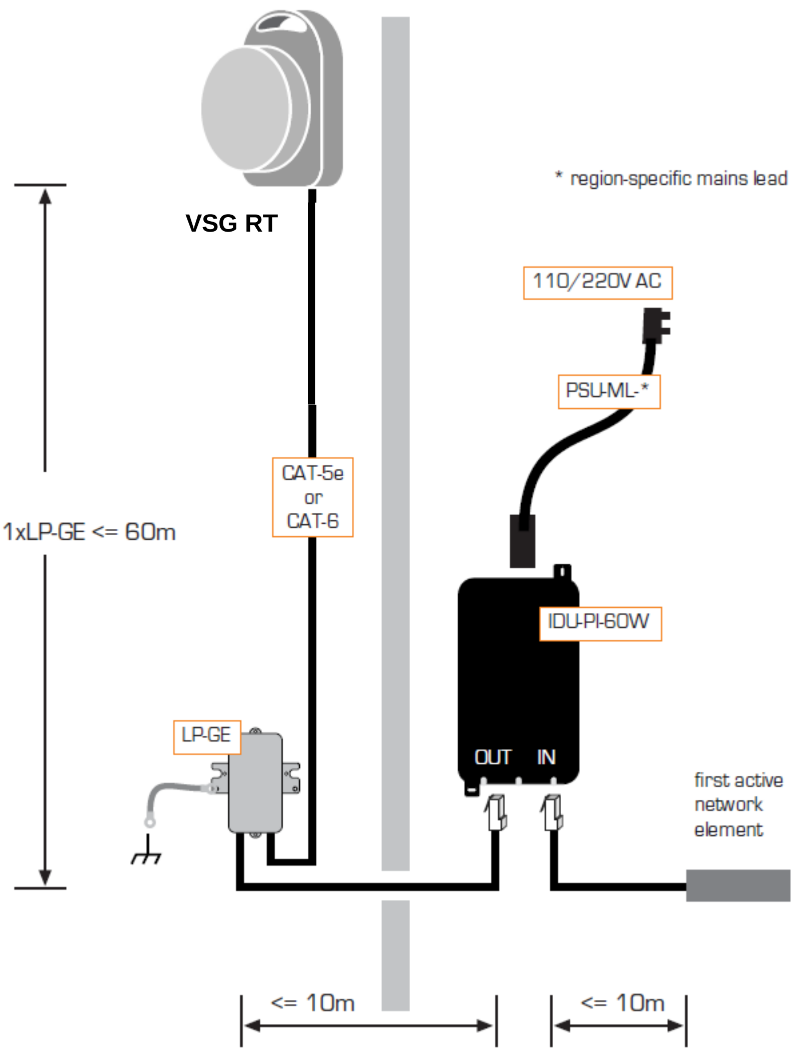

RT OUTDOOR EQUIPMENT

Order Code	Description	Number Required
CPE-##GE-####	RT outdoor unit without antenna.	1
ANT-##G-PSF-##	RT 30cm or 60cm slip-fit parabolic antenna	1
LP-GE	RT lightning surge protection box (CAT-5e only)	1
CABS-CAT-5	250m drum of external grade (UV protected) cat-5e cable	as required

RT INDOOR EQUIPMENT

Order Code	Description	Number Required
IDU-4VL-4E1	Indoor termination unit with 4xEthernet 4xE1 interfaces	0
VSG-DON-4VL	VSG to VS1/2 CAT-5 converter for use between VSG ODUs and IDU-4VL variants	0
IDU-GPI	Gigabit Power Injector Unit (supplied with input power connector)	0
IDU-PSU-70W	100 / 240V AC to -54V DC Power Supply (70W)	0
IDU-PI-60W	110 / 240V AC Power Injector with single Gigabit Ethernet Port (60W)	1
PSU-MIL-##	Mains Lead (select exact country variant)	1

Mains power is connected to the IDU-PI-60W, which in turn powers the VSG ODU.

VSG RT ETHERNET ONLY (-48V DC)

For an RT with lightning protection, the following equipment is required.

Note: This does not include installation materials e.g. feeder clamps, earth cables, nor does it include installation tools e.g. alignment tools or voltmeter etc.

DIAGRAM

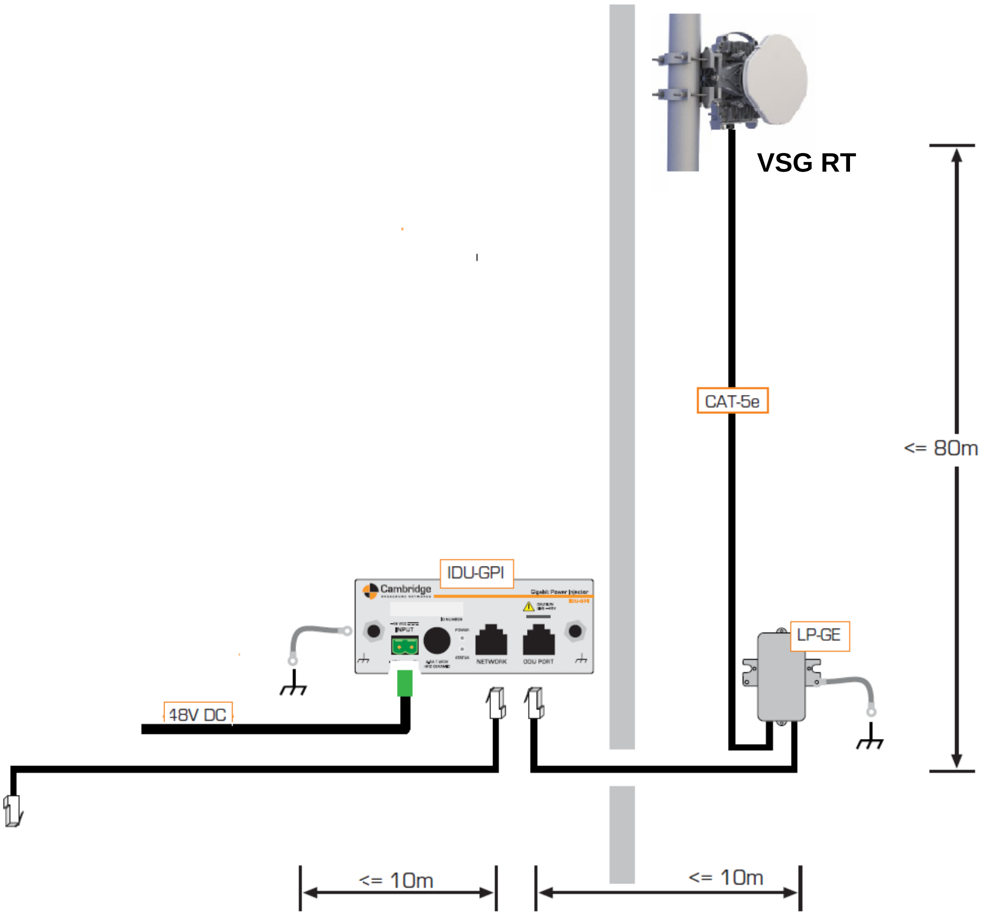

RT OUTDOOR EQUIPMENT

Order Code	Description	Number Required
CPE-##GE-dd	RT outdoor unit without antenna.	1
ANT-##G-PSF-##	RT 30cm or 60cm slip-fit parabolic antenna	1
LP-GE	RT lightning surge protection box (CAT-5e only)	1
CABS-CAT-5	250m drum of external grade (UV protected) cat-5e cable	as required

RT INDOOR EQUIPMENT

Order Code	Description	Number Required
IDU-4VL-4E1	Indoor termination unit with 4xEthernet 4xE1 interfaces	0
VSG-DON-4VL	VSG to VS1/2 CAT-5 converter for use between VSG ODUs and IDU-4VL variants	0
IDU-GPI	Gigabit Power Injector Unit (supplied with input power connector)	1
IDU-PSU-70W	100 / 240V AC to -54V DC Power Supply (70W)	0
IDU-PI-60W	110 / 240V AC Power Injector with single Gigabit Ethernet Port (60W)	0
PSU-MIL-##	Mains Lead (select exact country variant)	0

-48V DC power is connected to the IDU-GPI, which powers the VSG ODU.

VSG RT ETHERNET AND E1S (MAINS POWERED)

For an RT with lightning protection, the following equipment is required.

Note: This does not include installation materials e.g. feeder clamps, earth cables, nor does it include installation tools e.g. alignment tools or voltmeter etc.

DIAGRAM

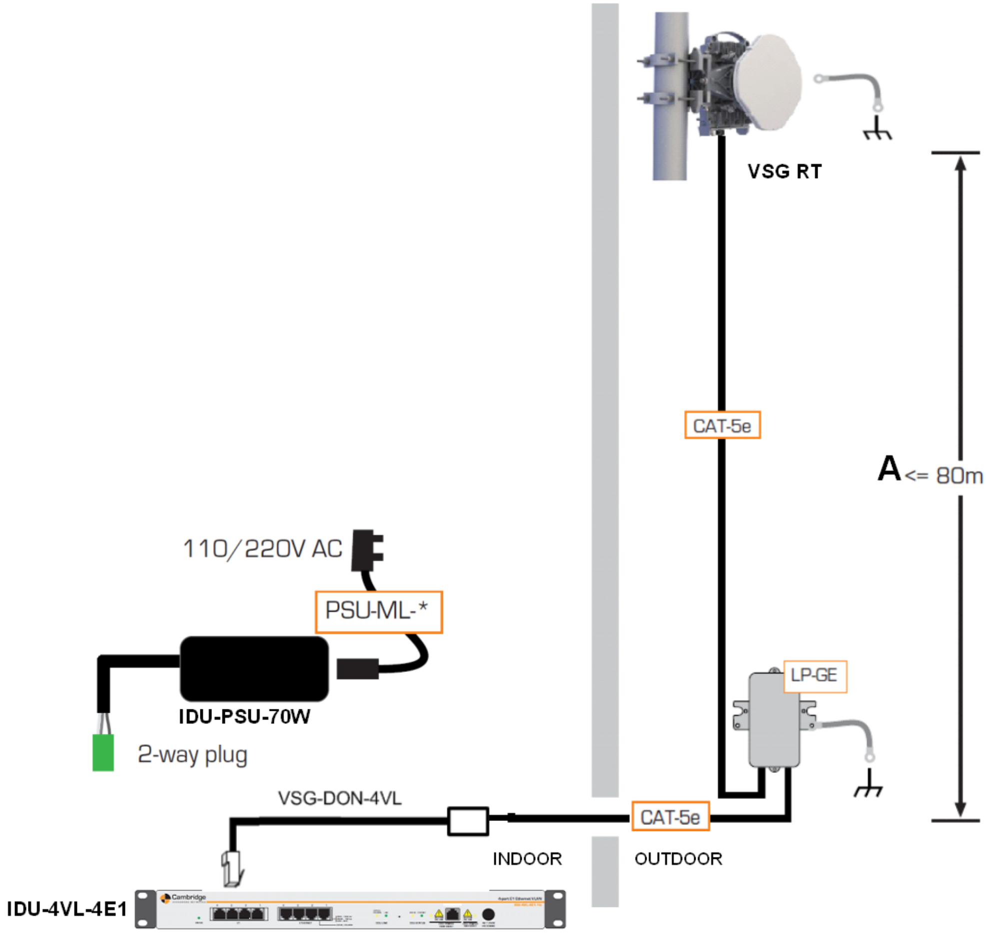

RT OUTDOOR EQUIPMENT

Order Code	Description	Number Required
CPE-##GE-dd	RT outdoor unit without antenna.	1
ANT-##G-PSF-##	RT 30cm or 60cm slip-fit parabolic antenna	1
LP-GE	RT lightning surge protection box (CAT-5e only)	1
CABS-CAT-5	250m drum of external grade (UV protected) cat-5e cable	as required

RT INDOOR EQUIPMENT

Order Code	Description	Number Required
IDU-4VL-4E1	Indoor termination unit with 4xEthernet 4xE1 interfaces	1
VSG-DON-4VL	VSG to VS1/2 CAT-5 converter for use between VSG ODUs and IDU-4VL variants	1
IDU-GPI	Gigabit Power Injector Unit (supplied with input power connector)	0
IDU-PSU-70W	100 / 240V AC to -54V DC Power Supply (70W)	1
IDU-PI-60W	110 / 240V AC Power Injector with single Gigabit Ethernet Port (60W)	0
PSU-MIL-##	Mains Lead (select exact country variant)	1

Mains power is connected to the IDU-PSU-70W, which powers the IDU-4VL-4E1.

The IDU-4VL-4E in turn powers the VSG ODU via the VSG-DON-4VL.

VSG RT ETHERNET AND E1S (-48V DC)

For an RT with lightning protection, the following equipment is required.

Note: This does not include installation materials e.g. feeder clamps, earth cables, nor does it include installation tools e.g. alignment tools or voltmeter etc.

DIAGRAM

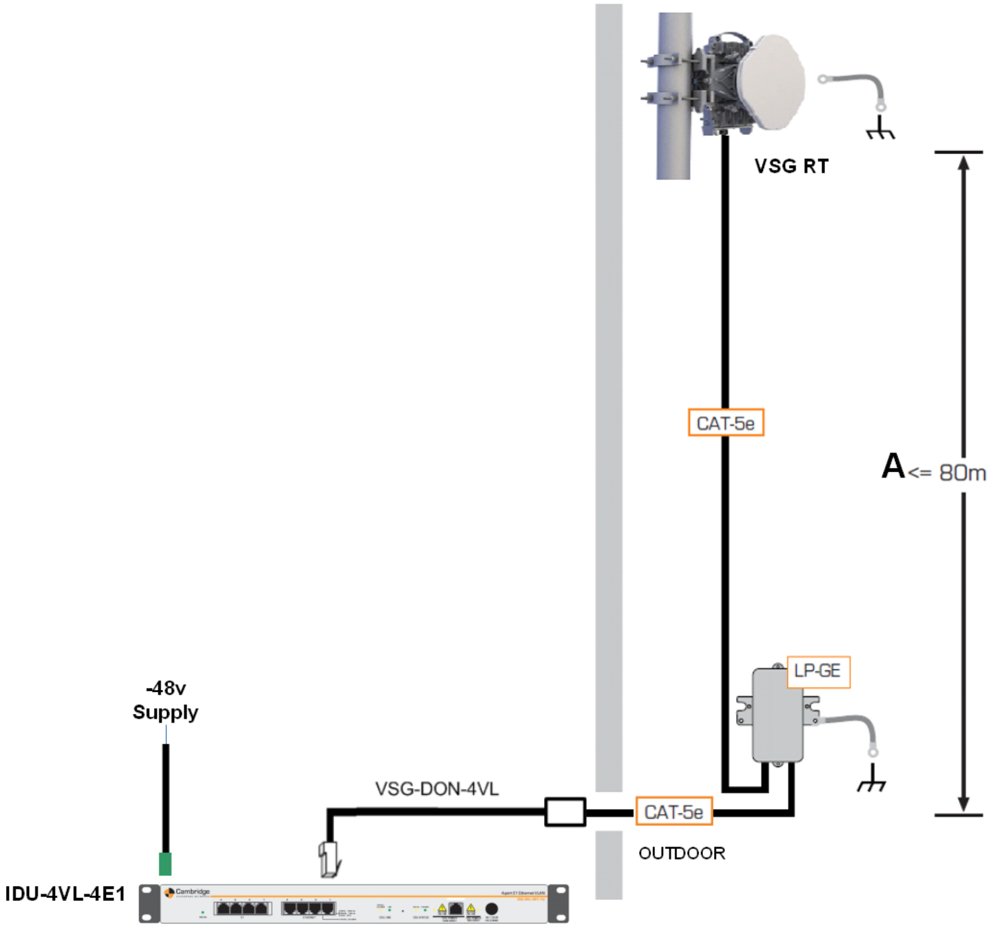

RT OUTDOOR EQUIPMENT

Order Code	Description	Number Required
CPE-##GE-dd	RT outdoor unit without antenna.	1
ANT-##G-PSF-##	RT 30cm or 60cm slip-fit parabolic antenna	1
LP-GE	RT lightning surge protection box (CAT-5e only)	1
CABS-CAT-5	250m drum of external grade (UV protected) cat-5e cable	as required

RT INDOOR EQUIPMENT

Order Code	Description	Number Required
IDU-4VL-4E1	Indoor termination unit with 4xEthernet 4xE1 interfaces	1
VSG-DON-4VL	VSG to VS1/2 CAT-5 converter for use between VSG ODUs and IDU-4VL variants	1
IDU-GPI	Gigabit Power Injector Unit (supplied with input power connector)	0
IDU-PSU-70W	100 / 240V AC to -54V DC Power Supply (70W)	0
IDU-PI-60W	110 / 240V AC Power Injector with single Gigabit Ethernet Port (60W)	0
PSU-MIL-##	Mains Lead (select exact country variant)	0

-48 V DC power is connected to the IDU-4VL-4E1 which in turn powers the VSG ODU via the VSG-DON-4VL.

VS2 HUB NON REDUNDANT ETHERNET AND SDH (VIA MPA)

For a non redundant Hub with lightning protection, the following equipment is required.

Note: This does not include installation materials e.g. feeder clamps, earth cables.

DIAGRAM

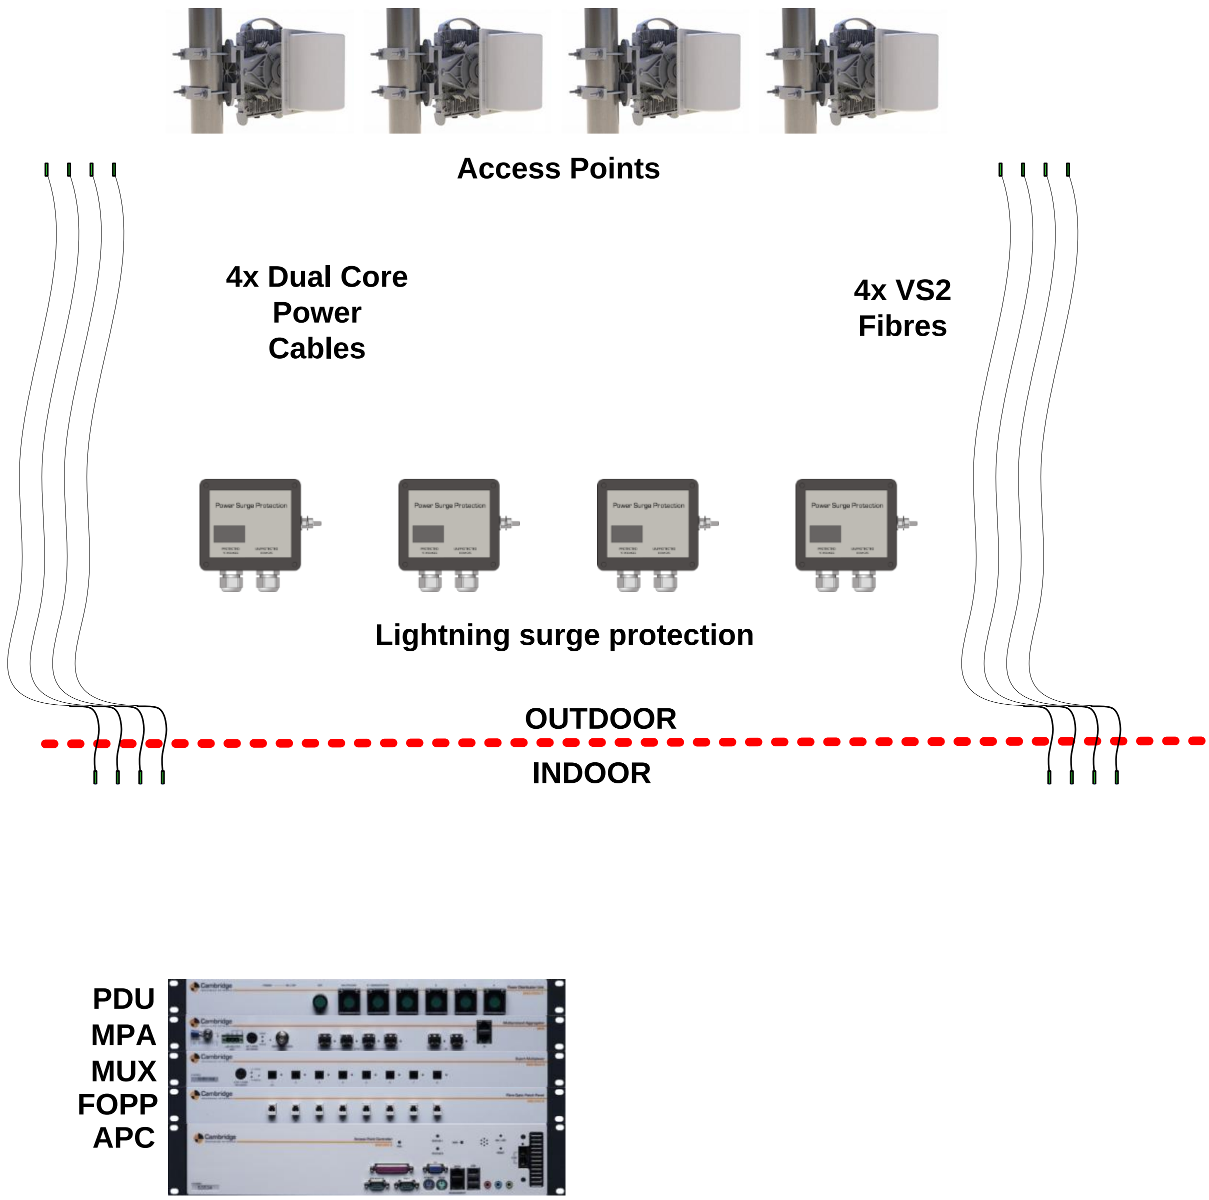

HUB OUTDOOR EQUIPMENT

Order Code	Description	Number Required
AP-##G-VS2-dd	Access Point outdoor unit without antenna	4
ANT-##G-AP-V	Access Point 90deg sector antenna – V polarisation	2
ANT-##G-AP-H	Access Point 90deg sector antenna – H polarisation	2
FIB-NG-##	## metre exterior grade fibre	4
LP-AP	2-core lightning surge protected power extension box	4
CABS-2C1.5	2-core 1.5mm2 copper cable, external grade (UV-protected). Priced per metre	as required

NB The exact polarisation of the antennas will depend on the RF plan and may not be a 50:50 split.

HUB INDOOR EQUIPMENT

Order Code	Description	Number Required
BSC-APC-C	Access Point Controller	1
BSC-MUX8-C	Access Point Multiplexer	1
BSC-MPA-C	MPA (includes 1x SDH SFP & Fibre to Mux)	1
BSC-E1CON8-C	E1 Concentrator	0
BSC-PAT8-C	Fibre Optic Patch Panel	1
BSC-PDU7-C	Power Distribution Unit	1
BSC-CABS-RD1	Cable and connector set for 1+1 BSC	0
MPA-CABS-RD1	Cable and connector set for 1+1 MPA	0
SFP-SDH-MM / SFP-SDH-SM-IR	SDH SFP connector STM-1 (1 off)	1
SFP-GE-1000LX / SFP-GE-1000T	Gigabit Ethernet SFP connector (1 off)	1

VS2 HUB REDUNDANT ETHERNET AND SDH (VIA MPA)

For a redundant Hub with lightning protection, the following equipment is required.

Note: This does not include installation materials e.g. feeder clamps, earth cables.

DIAGRAM

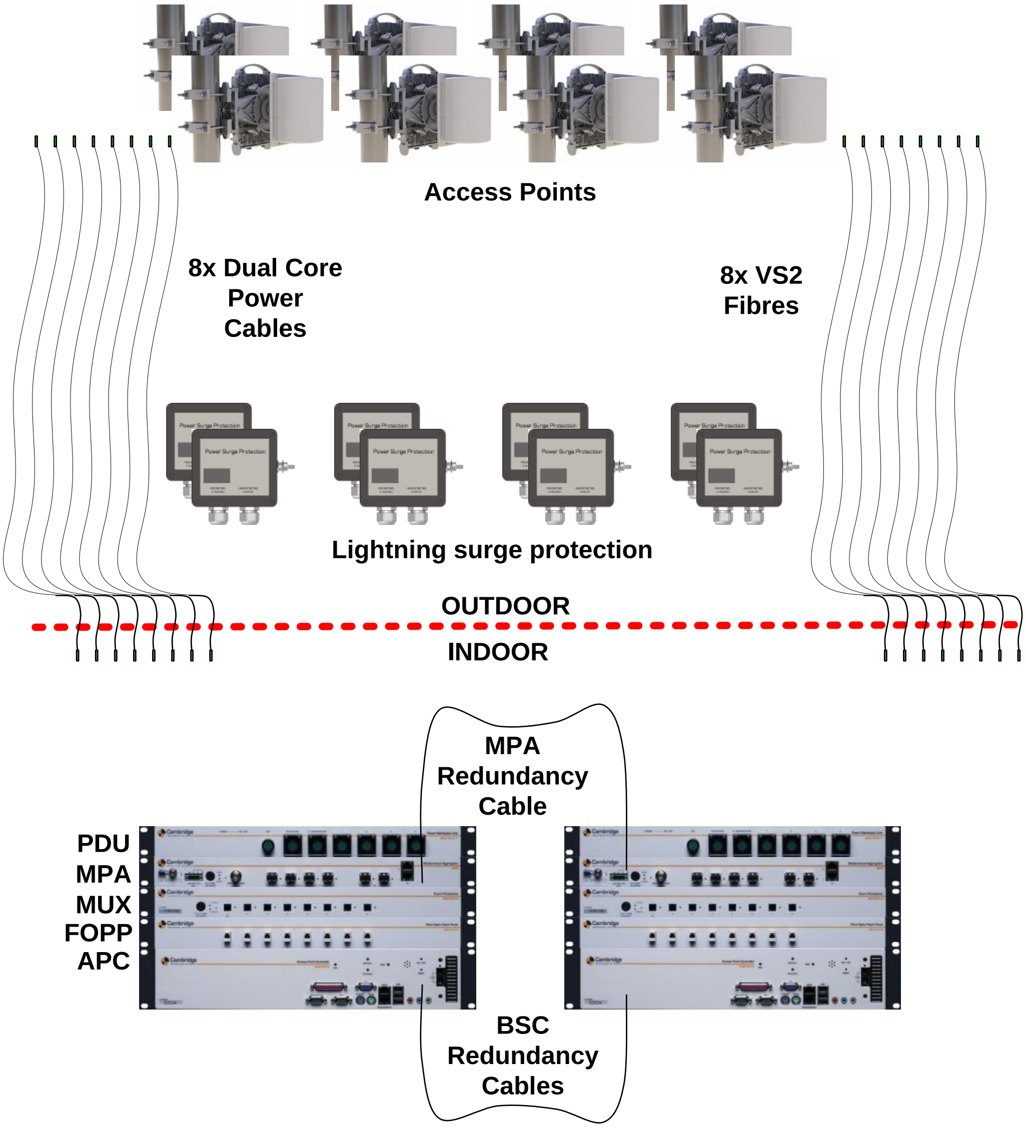

HUB OUTDOOR EQUIPMENT

Order Code	Description	Number Required
AP-##G-VS2-dd	Access Point outdoor unit without antenna	8
ANT-##G-AP-V	Access Point 90deg sector antenna – V polarisation	4
ANT-##G-AP-H	Access Point 90deg sector antenna – H polarisation	4
FIB-NG-##	## metre exterior grade fibre	8
LP-AP	2-core lightning surge protected power extension box	8
CABS-2C1.5	2-core 1.5mm2 copper cable, external grade (UV-protected). Priced per metre	as required

NB The exact polarisation of the antennas will depend on the RF plan and may not be a 50:50 split.

HUB INDOOR EQUIPMENT

Order Code	Description	Number Required
BSC-APC-C	Access Point Controller	2
BSC-MUX8-C	Access Point Multiplexer	2
BSC-MPA-C	MPA (includes 1x SDH SFP & Fibre to Mux)	2
BSC-E1CON8-C	E1 Concentrator	0
BSC-PAT8-C	Fibre Patch Panel	2
BSC-PDU7-C	Power Distribution Unit	2
BSC-CABS-RD1	Cable and connector set for 1+1 BSC	1
MPA-CABS-RD1	Cable and connector set for 1+1 MPA	1
SFP-SDH-MM / SFP-SDH-SM-IR	SDH SFP connector STM-1 (1 off)	2
SFP-GE-1000LX / SFP-GE-1000T	Gigabit Ethernet SFP connector (1 off)	2

REFERENCES

ETSI EN 300 019-1-1: Environmental Engineering (EE); Environmental conditions and environmental tests for telecommunications equipment; Part 1-1: Classification of environmental conditions; Storage

ETSI EN 300 019-1-2: Environmental Engineering (EE); Environmental conditions and environmental tests for telecommunications equipment; Part 1-2: Classification of environmental conditions; Transportation

ETSI EN 300 019-1-3: Environmental Engineering (EE); Environmental conditions and environmental tests for telecommunications equipment; Part 1-3: Classification of environmental conditions; Stationary use at weatherprotected Locations

ETSI EN 300 019-1-4: Environmental Engineering (EE); Environmental conditions and environmental tests for telecommunications equipment; Part 1-4: Classification of environmental conditions; Stationary use at non-weatherprotected locations

ETSI EN 300 132-2: Environmental Engineering (EE); Power supply interface at the input to telecommunications equipment; Part 2: Operated by direct current (dc)

ETSI EN 301 489-1: Electromagnetic compatibility and Radio spectrum Matters (ERM); Electromagnetic Compatibility (EMC) standard for radio equipment and services; Part 1: Common technical requirements

ETSI EN 301 489-4: Electromagnetic compatibility and Radio spectrum Matters (ERM); Electromagnetic Compatibility (EMC) standard for radio equipment and services; Part 4: Specific conditions for fixed radio links and ancillary equipment and services

ETSI EN 302 326-2: Fixed Radio Systems; Multipoint Equipment and Antennas; Part 2: Harmonized EN covering the essential requirements of article 3.2 of the R&TTE Directive for Digital Multipoint Radio Equipment

ETSI EN 302 326-3: Fixed Radio Systems; Multipoint Equipment and Antennas; Part 3: Harmonized EN covering the essential requirements of article 3.2 of the R&TTE Directive for Multipoint Radio Antennas

EN 60950-1: Information technology equipment. Safety. General requirements

EN 60950-22: Information technology equipment - Safety Part 22: Equipment installed outdoors

CEPT ERC T/R 12-05E: Harmonised radio frequency channel arrangements for digital terrestrial fixed systems operating in the band 10.0 - 10.68 GHz

CEPT ERC T/R 13-02E: Preferred Channel Arrangements for Fixed Services in the Range 22.0-29.5 GHz-Montreux 1993

ITU-R F.1568: Radio-frequency block arrangements for fixed wireless access systems in the range 10.15-10.3/10.5-10.65 GHz-Questions ITU-R 136/9 and ITU-R 229/9

ITU-R F.748-4: Radio-Frequency Arrangements for Systems of the Fixed Service Operating in the 25, 26 and 28 GHz Bands

ITU-T G.823: Control of Jitter and Wander within Digital Networks which are based on the 2048 KBIT/S hierarchy series G: Transmission Systems and Media, Digital Systems and Networks; Digital Networks - quality and availability targets

ITU-T G.704: Synchronous Frame Structures Used at 1544, 6312, 2048, 8448 and 44 736 kbit/s Hierarchical Levels - Series G: Transmission Systems and Media, Digital Systems and Networks Digital Transmission Systems - Terminal Equipments – General

ITU-T G.804: ATM cell mapping into plesiosynchronous digital hierarchy (PDH) Series G: Transmission systems and media, digital systems and networks.

ITU-T I.432: ISDN user-network interfaces – Layer 1Recommendations

GLOSSARY

Term	Definition
3G	3rd Generation GSM
ABIS	The interface within the GSM architecture, between the BTS and the BSC
ACM	Adaptive Code Modulation
AGC	Automatic Gain Control
AP	Access Point
APC	Access Point Controller
ARQ	Automatic Retransmission reQuest
ATM	Asynchronous Transfer Mode
ATPC	Automatic Transmit Power Control
BS	VectaStar Base Station
BSC	Base Station Controller (GSM)
BTS	Base Transceiver Station (GSM)
CBR	Constant Bit Rate
CDV	Cell Delay Variation
CES	Circuit Emulated Service
CIR	Committed Information Rate
CNR	Carrier to Noise Ratio
CPE	Customer Premises Equipment (complete installation of ODU and indoor equipment)
DSP	Digital Signal Processor
E1Conc	Base Station E1 Concentrator
EMS	Element Management System
EPD	Early Packet Discard
ETSI	European Telecommunications Standards Institute
FE	Fast Ethernet (100BaseT)
FEC	Forward Error Correction
FWA	Fixed Wireless Access
GE	Gigabit Ethernet (1000BaseX)
GSM	Global System for Mobile communications. Also referred to as 2G.
HSPA	High Speed Packet Access
IDU	InDoor Unit
IMA	Inverse Multiplexed ATM
IWF	InterWorking Function

LP-AP	Lightning Protection for an AP
LP-CODU	Lightning Protection for a CODU
LTE	Long Term Evolution
MAC	Medium Access Control
MPA	Multi Protocol Aggregator
MPOA	Multi Protocol Over ATM
MPLS	Multi-Protocol Label Switching
MSP	Multiplexer Section Protection
MUX	Base Station Multiplexer
NGN	Next Generation Networks
ODU	OutDoor Unit
PDU	Power Distribution Unit
PIR	Peak Information Rate
PMP	Point-to-Multipoint
PPC	PowerPC IBM CPU
PRBS	Pseudo-Random Bit Sequence
PtP	Point-to-Point
QAM	Quadrature Amplitude Modulation
QPSK	Quadrature Phase Shift Key
QoS	Quality of Service
RAN	Radio Access Network
RNC	Radio Network Controller (3G)
RoI	Return on Investment
RSL	Receive Sensitivity Level
SAR	Segmentation And Reassembly
SDH	Synchronous Digital Hierarchy
SFP	Small Form-factor Pluggable
SLA	Service Level Agreement
SNR	Signal to Noise Ratio
STP	Spanning Tree Protocol
TBC	To Be Confirmed
TBD	To Be Defined
TCM	Trellis Code Modulation
TDM	Time Division Multiplexing
TDMA	Time Division Multiple Access
TRAU	Transcoder and Rate Adaptation Unit

VC	Virtual Circuit
VNMS	VectaStar Network Management System
VS2	VectaStar2
WB	Wallbox

DOCUMENT PROPERTIES

Revision	Action	Reference
1.0	original text format	DKT
1.1	Minor change to CPE table	BRF
1.3	Change to 10.5GHz specification	DKT
1.4	Updated all sections to latest specs & features and sensitivity figures changes to be referenced to antenna interface	BRF
1.5	Company logo updated. Removed support for 5 and 10MHz channels for 10.5GHz variants	MDO
1.6	Updated with 28GHz specs	MDO
1.7	Updated with slip-fit variant and new MTBF values	MDO
1.9	Technical drawings added	MDO
2.1	VectaStar2 26 and 28 GHz specifications added	MDO
2.6	VectaStar2 10.5 GHz specifications added	MDO
2.6	Various changes to reflect updated product specifications	RS
3	Various changes added reference design guide	RS

DOCUMENT REVISION HISTORY

This document is based on template DOC-TMPL-0005-C00