



BROTSoLL – Annex 5

NGLL Technical Specifications

Communicated to the Belgian Institute for Postal services and Telecommunications on
13/06/2024

Table of Contents

1. Information about this document.....	3
1.1 Scope of this document.....	3
1.2 Abbreviations.....	4
2. NGLL building blocks.....	5
2.1 Connectivity services.....	5
2.1.1 Configuration of the NGLL Connectivity Service.....	5
2.1.2 Access Technology.....	7
2.1.3 Overall shaping of traffic.....	8
2.1.4 VLAN limitations.....	8
2.1.5 MAC learning.....	8
2.1.6 Special configurations.....	9
2.2 CPE (Customer Premises Equipment).....	12
2.2.1 Demarcation point.....	12
2.2.2 CPE types.....	12
2.2.3 MTU size.....	13
2.2.4 Protocol limitations.....	14
2.2.5 Multichassis LAG.....	15
2.2.6 CPE synchronization.....	15
2.2.7 Other CPE functionalities.....	15
2.3 Quality of Service (QoS).....	16
2.3.1 P-bit mapping.....	16
2.3.2 Data QoS.....	16
2.3.3 Voice QoS.....	17
2.3.4 Shaping and prioritization.....	17
3. NGLL Light Service.....	18
3.1 Concept.....	18

1. Information about this document

1.1 Scope of this document

The purpose of this document is to describe the technical specifications of the Next Generation Leased Line service (NGLL and NGLL Light).

To allow the Beneficiary to set up a service based on this service from Proximus, this document describes the interfaces.

The Beneficiary willing to offer features which require other technical characteristics than those implemented and supported by Proximus and described in the present reference offer, can implement them but without commitment of Proximus on its correct functioning. Examples of such features are:

- Non-tested protocols or protocols not supported by Proximus network equipment.
- Burst sizes, delay or jitter requirements, beyond the Proximus retail applications.

Proximus cannot guarantee, deliver support & be held liable regarding

- the correct functioning of such features in its network (at whatever time);
- the future evolution of other (supported or non-supported) features than those implemented and supported by Proximus for its own services.

Any enumeration of supported or non-supported protocols or features listed in this document are not exhaustive and are based on the Proximus best knowledge available at this moment.

Note that new EFM lines cannot be ordered anymore as of 01/01/2025. After that date existing lines will remain active and bandwidth updates will continue to be possible (when technically possible). Moving an existing EFM line to a new address will not be possible anymore.

1.2 Abbreviations

<i>Abbreviation</i>	<i>Description</i>
OAP	OLO Aggregation Point
APAL	(OLO) Aggregation Point Access Line
CBWFQ	Class-Based Weighted Fair Queuing
DS	Downstream
Gbps	Giga bit per sec (=1000Mbps)
HSCC	High Speed on Copper Concentrator
HSCR	High Speed on Copper Remote
Kbps	Kilo bit per sec
LAN	Local Area Network
LAG	Link Aggregation Group
LEX	Local Exchange
LTE	Line Termination Equipment
MAC@	MAC address
Mbps	Mega bit per sec (=1000Kbps)
MC-LAG	Multichassis LAG
MTU	Maximum Transmission Unit
NGLL	Next Generation Leased Line
OLO	Other Licensed Operator
p-bit	Priority bit
PoP	Point of Presence
QoS	Quality of Service
SR	Service Router
UNI	User Network Interface
US	Upstream
VLAN	Virtual LAN

2. NGLL building blocks

The NGLL service is based on the following service building blocks:

1. Connectivity services
2. CPE (Customer Premises Equipment) services
3. Class of Service/Quality of Service

Each of these building blocks is described in further detail in the following chapters.

2.1 Connectivity services

NGLL is an Ethernet (Layer2) connectivity service, based on Ethernet over MPLS. Proximus proposes the access technology Ethernet over copper (EFM) and fiber-based Gigabit Ethernet and 10 Gigabit Ethernet.

Note that no new EFM lines can be ordered anymore as of 01/01/2025.

The NGLL service connects 1 or more End-User Sites to 1 OLO Aggregation Point or a LAG of 2 OLO Aggregation Points. The connectivity service is transported over the Proximus TITAN network. Figure 1 schematically depicts the NGLL service.




Figure 1: Schematic of the NGLL Service, connecting several End-User sites over fiber or EFM to one OLO Aggregation Point via the APAL

Within the NGLL solution, traffic will always flow between an End-User site and an OLO Aggregation Point, never between an End-User Site and another End-User Site (no any-to-any functionality).

2.1.1 Configuration of the NGLL Connectivity Service

In the following paragraphs, a distinction will be made between OLO Aggregation Point (OAP) and End-User Sites.

A start configuration consists of an OLO Aggregation Point and one or more End-User Sites. Afterwards, adding End-User Sites to this configuration is relatively straightforward, with only the additional sites to be ordered and provisioned. The Beneficiary defines if and when the bandwidth of the Aggregation Point needs to be adapted, taking into account the required service quality and the estimated implementation timeframe.

2.1.1.1 OLO Aggregation Point

The OLO Aggregation Point is connected to a Proximus Service Node using a Gigabit Ethernet or 10Gigabit Ethernet access line over fiber optic, called Aggregation Point Access Line (APAL).

The traffic of all End-User Sites will be delivered on the OLO Aggregation Point Access Line and is demarcated by a Proximus managed CPE switch, which makes entirely part of the service. The physical interface to the Beneficiary, on the CPE switch, can be copper or fiber and must work in QinQ.

The Aggregation Point device of the OLO must work in QinQ.

2.1.1.2 End-User Sites

The End-User Sites can be connected via Ethernet over copper (EFM) or Ethernet over fiber. Both Access Technologies are demarcated on a Proximus managed CPE switch, which makes entirely part of the service. The physical interface to the End-User, on the CPE switch, can be copper or fiber and must work in dot1Q.

Note that no new EFM lines can be ordered anymore as of 01/01/2025.

2.1.1.3 NGLL connections

NGLL presents an Ethernet interface to Beneficiaries, simplifying the LAN/WAN boundary for Service Providers and Beneficiaries, and enabling rapid and flexible service provisioning, as the service bandwidth is not directly tied to the physical interface.

A unique VLAN-id identifies the End-User Site at the Aggregation Point level. This is the so-called S-tag or outer tag at the Aggregation Point level.

The End-User Site works in dot1Q and the VLAN tags (VLAN ids and p-bit) sent by a specific End-User Site are received unchanged at the OLO Aggregation Point switch as C-tags, under the S-tag, identifying the specific End-User Site. Both S-tags and C-tags can be freely chosen by the Beneficiary, within limitations mentioned further under section 2.1.4.

Figure 2 shows an example set-up with two End-User Sites.

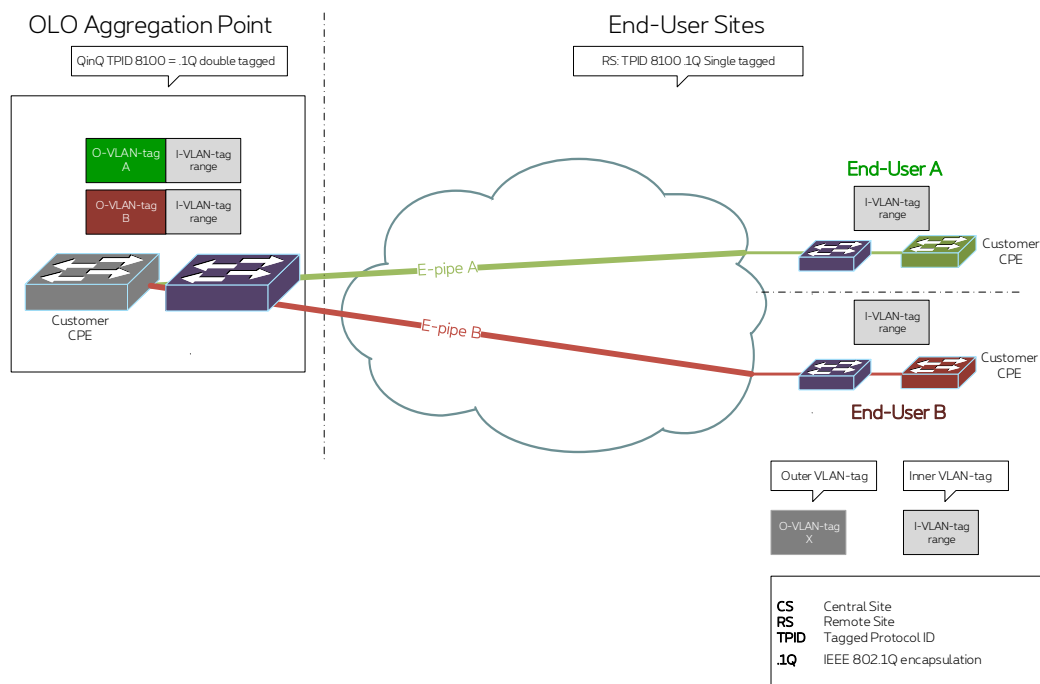


Figure 2: VLAN tagging within the NGLL solution

2.1.2 Access Technology

The Access Technology is the technology used between the Proximus equipment in the LEX up to the LTE at the NGLL Site.

For the OLO Aggregation Point, an Ethernet on Fiber link is always used. For End-User Sites the Beneficiary has the choice between Ethernet on Fiber or EFM. The choice of the Access Technology depends on the ordered bandwidth, the availability in the network, and the location in the network.

As from 01/01/2025, EFM will not be available anymore for new installations on End-User Sites.

2.1.2.1 Ethernet on fiber access technology

The Ethernet on fiber access is a high-speed data transfer service, offering interface speed options from 10Mbps to 10Gbps between the site and the TITAN network, regardless of the distance between them. Traffic shaping increases granularity, allowing for bandwidth profiles from 2Mbps up to 10Gbps in smaller bandwidth steps.

The Ethernet access service is based on the Proximus fiber optic infrastructure. This is a comprehensive service which includes the network infrastructure.

2.1.2.2 Ethernet in the First Mile (EFM) access technology

At the End-User Site, a High Speed on Copper End-User (HSCR) device is connected through one or more copper pair(s) (maximum 8) to a High Speed on Copper Concentrator (HSCC) in the closest LEX. From here, traffic is transported to the TITAN MPLS platform.

The availability of such service is determined by:

- the attenuation of the copper pairs (linked to the distance)
- the number of copper pairs available on the site

Available bandwidth profiles are 2, 4, 6, 8, 10 or 20 Mbps.

Note that no new EFM lines can be ordered anymore as of 01/01/2025

2.1.3 Overall shaping of traffic

The overall Ethernet flows on End-User and Aggregation Point access technology are shaped to a value, ordered by the Beneficiary, and further called the bandwidth profile.

- 2Mbps to 10Mbps: 2Mbps steps (2,4,6,8,10)
- 10Mbps to 100Mbps: 10Mbps steps (10,...90, 100)
- 100Mbps to 1Gbps: steps of 100Mbps
- 1Gbps to 10Gbps: steps of 1Gbps

There is no limitation to the data volume that can be transferred.

2.1.4 VLAN limitations

Max 1000 S-tags @ Aggregation Point in range [2,4000], to be allocated by the Beneficiary and configured by Proximus and Beneficiary. VLAN id 999, 1002, 1003, 1004, 1005 and > 4000 are reserved and cannot be used by the Aggregation Point CPE.

Max 1000 C-tags @ End-User Site in range [2,4000], to be allocated by the Beneficiary and configured by Proximus (in CPE switch) and Beneficiary. VLAN id 999, 1002, 1003, 1004, 1005 and > 4000 are reserved and cannot be used by the End-User.

The Beneficiary can request, under his own responsibility, to use S-tag or C-tag 1, but this is not recommended by Proximus and this is against the good usage rules recommended by Cisco.

In case of problems due to the usage of VLAN 1 by a Beneficiary, Proximus cannot be held responsible for the outage.

2.1.5 MAC learning

Because of MAC learning in several equipments in the NGLL configuration, the number of MAC addresses shall be limited to 500 per End-User Site and to 8000 per Aggregation point.

2.1.6 Special configurations

2.1.6.1 Distant Aggregation Point Access Line

The Service Nodes 03CEN and 91GKK respectively moved to 03WOM and 91WON.

The “Distant Intro offer” consists in providing a limited number of Distant Intros to the Beneficiary to allow him to reconnect its fibers spliced in different Beneficiary-manholes around the former Service Nodes O3CEN and/or 91GKK via Distant Intros of Proximus towards the Beneficiary colocation rooms in the respective new Service Node buildings O3WOM and/or 91WON.

For a detailed description, reference is made to the Appendix A of the Main Body of the Bitstream xDSL reference offer.

2.1.6.2 Inter-area NGLL connection

For all NGLL End-User Sites an interconnection to an APAL or LAG APAL in another Service Area is also possible. A list of all Proximus Service Areas and corresponding local nets is made available to the Beneficiary via the Personal Page.

2.1.6.3 Multichassis LAG APAL

It is possible to automatically bypass a failing Service Node by deviating all traffic via another Service Node in one of the 5 Service Areas, by means of the LACP protocol, aka “multichassis LAG”. In case of a multichassis LAG APAL (MC-LAG), this LAG will consist of two fibre links, each to a different Service Node router in the same or a different Service Area. One of these links will be working, the other will be stand-by.

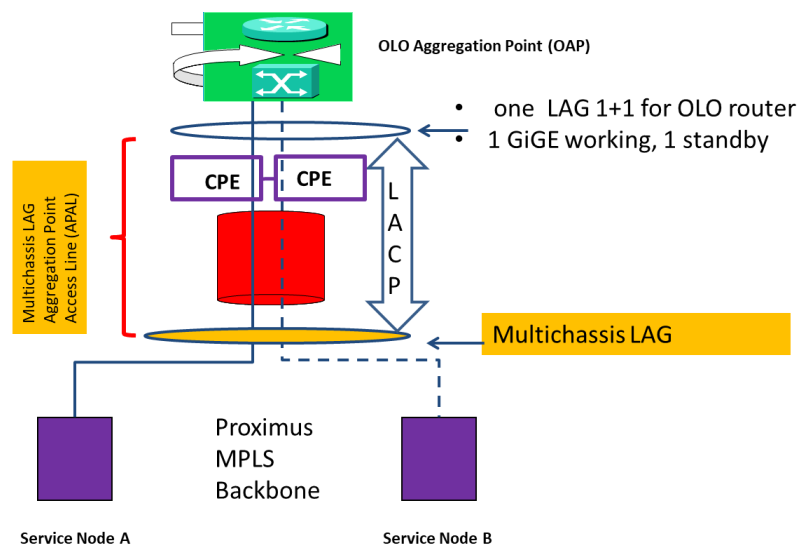


Figure 3: schematic of the NGLL MC-LAG APAL solution

The LACP protocol, to be configured on the LAG at OLO Aggregation Point device, is peered by both Service Node routers, one in each of the two Service Nodes in the Proximus MPLS Backbone. This LACP

protocol will make one link “working” and the other link “stand-by” and switches the connections, coming from the End-User Sites, to the working link.

MC-LAG adds node-level redundancy to the normal link-level redundancy that a LAG provides.

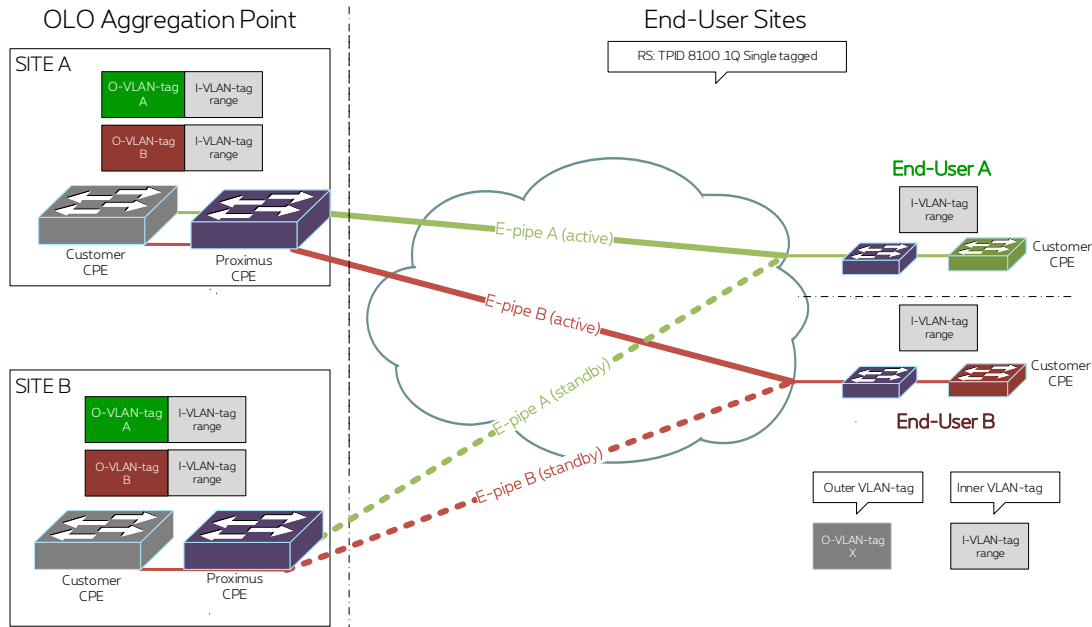


Figure 4: VLAN tagging in an MC-LAG set-up. The Customer CPE can be either one device or two separate devices

2.1.6.4 NGLL Standard+

The NGLL Standard+ solution¹ is a service consisting of a dual NGLL or NGLL Light (cfr. chapter 3 in this Annex for a description of the NGLL Light Service) fiber access between a single CPE on the End-User site and a TITAN L-Edge, in an Active/Standby set-up. On top of this physical setup a LAG will be configured between the CPE on the End-User site and the TITAN L-Edge to support an Active/Standby Layer 2 connection to the Proximus network. Under normal conditions the active Ethernet Access Line is in use. In case of service interruption on the active Ethernet Access Line, the standby Ethernet Access Line can take over the connectivity after a short service interruption (LAG convergence time).

¹ The NGLL Standard+ solution will be developed within a reasonable timeframe in case a concrete request is received from an OLO.

The figures below present the Standard+ solution on NGLL and NGLL Light.



Figure 5: schematic of the NGLL Standard+ solution (with NGLL or NGLL Light)

The NGLL Standard+ solution requires dual fiber SFPs in the CPE and L-Edge as well as dual NGLL or NGLL Light fiber access lines. Both NGLL or NGLL Light fiber access lines on an End-User site with an NGLL Standard+ solution have the same bandwidth profile. For the available NGLL and NGLL Light bandwidth profiles, reference is made to the “Main Body”. Based on the current CPE portfolio limited options are available on the CPE LAN interface².

The protection as provided by the NGLL Standard+ solution has a positive effect on the uptime of the service. Failure of power, fiber cuts including both Ethernet Access lines, etc. can still lead to a service interruption of the NGLL/NGLL Light End-User site.

² Note that today we can't support on NGLL a 1+1 configuration on 10G access, our current L2 CPE does not support 3*10Gbps ports.

2.2 CPE (Customer Premises Equipment)

2.2.1 Demarcation point

The connectivity access lines are installed & configured by Proximus, including the “demarcation” CPE, typically a switch. The Beneficiary can install and configure his own CPE (e.g.: an IPVPN router) behind the Proximus switch.

- The current demarcation CPE on the Aggregation Point and End-User Sites is a Proximus-managed switch.
- Traffic is delivered to the End-User via one customer-facing port (10BaseT RJ45, or 100BaseT RJ45, or 1000BaseT RJ45), or optical LC connector.

2.2.2 CPE types

As availability of CPEs can change in the future the currently proposed model can be replaced in the future with an equivalent model. These CPEs will be used today :

- OLO aggregation point Customer-sited or Proximus-sited :
 - o Up to 1Gbps:
 - OneAccess 1646: Gigabit RJ45 port or SFP Gigabit Ethernet (multimode or single-mode); AC power
 - o For 10Gbps:
 - ASR-920-12CZ-A/D - LAN: SFP 10Gigabit Ethernet (multimode or single-mode); AC or DC power
- OLO aggregation point Proximus-sited :
 - o Up to 1Gbps:
 - OneAccess 1646: Gigabit RJ45 port or SFP Gigabit Ethernet (single-mode); AC power
 - o For 10Gbps:
 - ASR-920-12CZ-A/D - LAN: SFP 10Gigabit Ethernet (single-mode); AC or DC power
- End-User sites :
 - o OneAccess 1646: Gigabit RJ45 port or SFP Gigabit Ethernet (multimode or single-mode); AC power
 - o For 10Gbps:
 - ASR-920-12CZ-A/D - LAN: SFP 10Gigabit Ethernet (multimode or single-mode); AC or DC power

Disclaimer: further tests are required to validate that the used CPEs support extreme configurations such as e.g. 50% Voice QoS on a 1 Gbps or more than 750Mbps Voice QoS on a 10 Gbps NGLL-line.

2.2.3 MTU size

In this document, the MTU size is defined as the size of the "MAC Client Data" field in a tagged MAC Frame. [Figure 5](#) and [Table 1a](#) and [Table 1b](#) list the fields in such a frame.

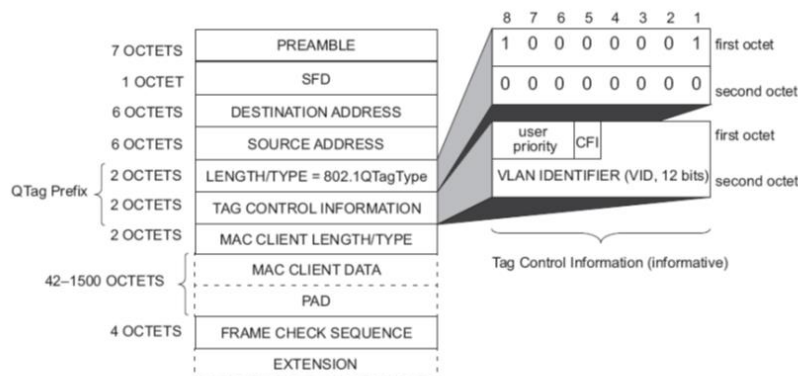


Figure 5: tagged MAC frame format (from IEEE 802.3-2005)

Table 1a: overview of fields in a tagged MAC Frame – End-User Sites on copper (using EFM technology)

Field	Field size
Destination Address	6 octets
Source Address	6 octets
Dot1Q inner VLAN tag (C-tag)	4 octets
MAC Client Length/Type	2 octets
MAC Client Data +PAD	1542 octets
Frame Check Sequence	4 octets
Total	1564 octets

Table 2b: overview of fields in a tagged MAC Frame -- End-User Sites and APALs on P2P dedicated Fiber (Ethernet over Fiber)

Field	Field size
Destination Address	6 octets
Source Address	6 octets
Dot1Q inner VLAN tag (C-tag)	4 octets
MAC Client Length/Type	2 octets
MAC Client Data +PAD	9000 octets
Frame Check Sequence	4 octets
Total	9022 octets

Note that one C-tag at End-User demarcation point is mandatory for indicating the user priority (also for pO). On the OLO Aggregation Point the MAC Frame Size also includes the additional 4 octets of the dot1Q outer VLAN-tag (S-tag).

Table 3

[Table 3](#) lists the MTU size and maximum Frame Size for both the End-User Site and the OLO Aggregation Point.

Table 3: MTU sizes for End-User Site and OLO Aggregation Point on P2P dedicated Fiber (Ethernet over Fiber)

Access Type	MTU	Max Ethernet Frame Size
End-User Site	9000 octets	9022 octets
OLO Aggregation Point	9000 octets	9026 octets

Note: on Beneficiary's request and in specific set-ups the MTU size can be increased to 1800 octets

2.2.4 Protocol limitations

The service is intended for transport of IP packets by Ethernet frames (IPoE). Some protocols outside of this IPoE framework do not transparently pass. Table 3 summarizes the transparency test results for the different NGLL CPEs.

The lists provided are indicative and non-exhaustive.

Transparency of the types of L2CP frames as listed in the table has been validated upon Destination MAC address of the L2CP frame, but not on any other field of the L2CP frame¹.

Beneficiary has the opportunity to request on a project mode basis for ad hoc testing to check the transparency of any specific protocol in the context of NGLL.

¹ Lab validation has been performed on the following CPE image versions:

- Cisco ME3400: Cisco IOS 12.2.60 EZ 11 – Metro Access (release date 05/06/2017)
- Cisco ASR920: Cisco ASR 920 Series IOS XE UNIVERSAL -NO PAYLOAD ENCRYPTION 3.18.0.SP – Metro Access Service License (release date 29/06/2016)
- OneAccess 1646: TDRE14.002.793 (release date 08/09/2017)

Table 4: summary of transparency results for NGLL CPEs

	<i>DA Ethernet / DA IP</i>	<i>Cisco ME3400</i>	<i>Cisco ASR920</i>	<i>OneAccess 1646</i>
Control frames				
IGMP		Discard	Forward	Forward
IEEE L2CP Frames				
STP IEEE	01-80-c2-00-00-00	Policed ²	Policed ²	Forward
Pause Frames	01-80-c2-00-00-01	Discard	Discard	Discard
LACP/LAMP	01-80-c2-00-00-02	Discard	Discard	Forward
802.3 ah	01-80-c2-00-00-02	Discard	Discard	Discard
Port authentication	01-80-c2-00-00-03	Discard	Policed ²	Forward
E-LMI	01-80-c2-00-00-07	Discard	Policed ²	Forward
LLDP	01-80-c2-00-00-0E	Policed ²	Policed ²	Forward
MMRP	01-80-c2-00-00-20	Forward	Forward	Forward
MVRP	01-80-c2-00-00-21	Forward	Forward	Forward
Cisco L2CP frames				
CDP	01-00-0c-cc-cc-cc	Policed ²	Policed ²	Forward
PVSTP+	01-00-0c-cc-cc-cd	Policed ²	Policed ²	Forward

2.2.5 Multichassis LAG

In the case of a multichassis LAG (MC-LAG) set-up two CPEs make up the demarcation point between the Beneficiary and Proximus.

Both CPEs are :

- configured transparent for LACP, so that LACP is peered by the SRs and not by the CPEs.
- interconnected by a trunk, allowing Proximus to manage also the CPE on the link in stand-by mode, via the working link and corresponding CPE.

2.2.6 CPE synchronization

No clock reference is provided by Proximus.

2.2.7 Other CPE functionalities

The Beneficiary has the opportunity to request on a project mode basis for ad hoc testing to check other CPE functionalities.

² For security reasons, parameters will not be publicly shared. Beneficiary shall contact Proximus if he has the need to receive this information.

2.3 Quality of Service (QoS)

The Proximus NGLL Service offers two types of Quality of Service, based on 4 Classes:

- 1) Voice QoS: a service based on the Voice QoS class, dedicated to offer better performance for jitter and delay sensitive traffic
- 2) Data: a service dedicated to data prioritization, split over 3 QoS classes
 - o The prioritisation of **Business Critical Data** is based on the p-bit value 3.
 - o The prioritisation of **Business Data** is based on the p-bit value 1.
 - o **Best Effort**: Ethernet frames marked with p=0 or non-ordered QoS services are treated as “Best Effort” in the network except Ethernet frames marked with p=5 when only Data QoS is ordered

2.3.1 P-bit mapping

Table 5: P-bit mapping of the QoS classes for Proximus NGLL

QoS class	Ethernet frame P-bit
Voice	5
Business Critical Data (BC)	3
Business Data (BD)	1
Best Effort (BE)	0

Traffic sent with p-bit value 2,4, 6 or 7 is transported as “Best Effort” and retagged to p-bit value 0.

If no “Data” QoS is ordered then all traffic with p-bit value 1 or 3 is transported as “Best Effort” but keeps the p-bit value 1 or 3.

If no “Voice” QoS is ordered then:

- If no Data QoS ordered: all traffic with p-bit value 5 is transported as “Best Effort” but keeps the p-bit value 5.
- If Data QoS ordered: most traffic with p-bit value 5 is dropped. The small part that is not dropped is transported as “Voice” and keeps the p-bit value 5.

Proximus considers the switch interface to the Beneficiary as a Trusted Boundary, meaning that it is up to the Beneficiary to take care of the correct P-bit marking and that e.g. only voice will be sent on accesses configured with Voice QoS. Proximus will not check the Layer 3 or upper layers.

2.3.2 Data QoS

The three different data classes are configured with a 80/16/4 distribution (i.e. 80% BC, 16% BD & 4% BE)

- o 80% of the bandwidth is assigned to BC traffic
- o 16% of the bandwidth is assigned to BD traffic
- o 4% of the bandwidth is assigned to BE traffic

The Data QoS (BC, BD & BE) classes are class-based weighted fair queuing classes: if bandwidth is not used in one of the data classes, it can be re-used in the other data classes.

2.3.3 Voice QoS

The Voice QoS class has the highest priority in the network and is also designed to offer better performance for jitter and delay sensitive traffic (e.g. voice and real-time traffic). The Voice QoS class is not intended for bursts of data. It is advised to send traffic with an appropriate shaping and small packet size to avoid packet loss.

The Voice QoS class is a strict priority class: the configured voice bandwidth cannot be exceeded. Any traffic sent tagged with Voice P-bit exceeding the configured Voice QoS bandwidth will be dropped, even if there is bandwidth available in other classes.

Voice QoS can be ordered as 10%, 25%, 50% or 75% of the access bandwidth profile, with a maximum of 750 Mbps.

2.3.4 Shaping and prioritization

Shaping and prioritisation of the overall traffic per site (End-User Site and OLO Aggregation Point) is performed in 3 steps per End-User Site and per OLO Aggregation Point:

- 1) Shaping of the site to the ordered "Bandwidth Profile", e.g. 500Mbps
- 2) The optional prioritisation of the "Voice" (p5 marked) traffic
- 3) The optional prioritisation of the "Data" traffic

If at any time the reserved bandwidth of a certain type of traffic is not completely used, then another type of data traffic can use the remaining bandwidth. This rule is applicable on the "Voice" bandwidth as well, but not for the Voice traffic which, as stated in paragraph 2.3.3, is limited to the ordered "Voice" QoS bandwidth.

If the total inflow of traffic surpasses the Access Line bandwidth profile then 80% of the bandwidth reduced with the actual "Voice" traffic is used to forward "Business Critical data" traffic and 20% is used to forward other traffic. The bandwidth used to forward other traffic in turn is split into 80% Business Data and 20% Best Effort.

3. NGLL Light Service

3.1 Concept

The NGLL light solution is only available on P2P Fiber and is based on "QoS resource management". In upstream a global Shaping for all E-pipes (i.e. Dedicated VLANs) will be applied with a higher oversubscription factor. All E-pipes will also share the same buffers per Pbit. In downstream, global Policing will be applied with a higher oversubscription factor for all E-pipes instead of Shaping, so there will be no buffers.

The drawing below represents this principle graphically.




Figure 6: Schematic of the NGLL Light solution

A start configuration consists of an OLO Aggregation Point and one or more End-User Sites. Afterwards, adding End-User Sites to this configuration is relatively straightforward, with only the additional sites to be ordered and provisioned. The Beneficiary defines if and when the bandwidth of the Aggregation Point needs to be adapted, taking into account the required service quality and the estimated implementation timeframe.

Proximus reserves the right to slightly change and optimize the concept as described above should this prove to be necessary during the development phase of the NGLL Light solution.

-----End of document-----