

Architecture for Ethernet Active Line Access (ALA)

NICC Standards Limited

Michael Faraday House,
Six Hills Way,
Stevenage
SG1 2AY

Tel.: +44(0) 20 7036 3636

Registered in England and Wales under number 6613589

NOTICE OF COPYRIGHT AND LIABILITY

© 2010 NICC copyright

The present document may be made available in more than one electronic version or in print. In any case of existing or perceived difference in contents between such versions, the reference version is the Portable Document Format (PDF). In case of dispute, the reference shall be that printing on NICC printers of the PDF version kept on a specific network drive within the NICC.

Users of the present document should be aware that the document may be subject to revision or change of status. Information on the current status of this and other NICC documents is available at:

<http://www.niccstandards.org.uk/publications/>

If you find errors in the present document, please send your comments to:

<mailto:help@niccstandards.org.uk>

Copyright

All right, title and interest in this document are owned by NICC Standards Limited ("NICC") and/or the contributors to the document (unless otherwise indicated that copyright is owned or shared with a third party). Such title and interest is protected by United Kingdom copyright laws and international treaty provisions.

The contents of the document are believed to be accurate at the time of publishing, but no representation or warranty is given as to their accuracy, completeness or correctness. You may freely download, copy, store or distribute this document provided it is not modified in any way and it includes this copyright and liability statement.

You may not modify the contents of this document. You may produce a derived copyright work based on this document provided that you clearly indicate that it was created by yourself and that it was derived from this document and provided further that you ensure that any risk of confusion with this document is avoided.

Liability

Whilst every care has been taken in the preparation and publication of this document, neither NICC, nor any working group, committee, member, director, officer, agent, consultant or adviser of or to, or any person acting on behalf of NICC, nor any member of any such working group or committee, nor the companies, entities or organisations they represent, nor any other person contributing to the contents of this document (together the "Generators") accepts liability for any loss or damage whatsoever which may arise from the use of or reliance on the information contained in this document or from any errors or omissions, typographical or otherwise in the contents.

Nothing in this document constitutes advice. Nor does the transmission, downloading or sending of this document create any contractual relationship. In particular no licence is granted under any intellectual property right (including trade and service mark rights) save for the above licence to download copy, store and distribute this document and to produce derived copyright works.

The liability and responsibility for implementations based on this document rests with the implementer, and not with any of the Generators. If you implement any of the contents of this document, you agree to indemnify and hold harmless each Generator in any jurisdiction against any claims and legal proceedings alleging that the use of the contents by you or on your behalf infringes any legal or other right of any of the Generators or any third party.

None of the Generators accepts any liability whatsoever for any direct, indirect or consequential loss or damage arising in any way from any use of or reliance on the contents of this document for any purpose.

The NICC Standards Web site contains the definitive information on the [IPR Policy and Anti-trust Compliance Policy](#)

If you have any comments concerning the accuracy of the contents of this document, please write to:

The Technical Secretary, NICC Standards Ltd.,

Michael Faraday House,
Six Hills Way,
Stevenage
SG1 2AY

Contents

Intellectual Property Rights	5
Foreword	5
Introduction	5
1 Scope	6
2 References	7
2.1 Normative references	7
2.2 Informative references	7
3 Key Words, Definitions and Abbreviations	8
3.1 Key Words	8
3.2 Definitions	8
3.3 Abbreviations	8
4 Introduction to Active Line Access (ALA)	10
4.1 An overview of ALA documentation	10
4.2 ALA Business Entities	11
4.3 Network topologies	12
5 ALA Service Overview	14
5.1 ALA User Connections	14
5.2 End to End ALA Service Definitions	15
6 ALA Architecture and reference points	17
6.1 ALA for access networks using Ethernet aware NTUs	18
6.2 ALA for wires only interfaces to access networks	18
6.3 ALA and baseband voice services	18
6.4 Reference Point A-1	19
6.5 Reference Point A-2	19
6.6 Reference Point A-3	19
6.7 Reference point A-4	19
6.8 Support for multiple ALA users at reference points A-1 and A-2	20
6.8.1 Multiple ALA users at reference point A-1	20
6.8.2 Multiple ALA users at reference point A-2	20
7 Ethernet Transport Solutions	22
7.1 1:1 VLAN Ethernet transport solution	22
7.2 N:1 VLAN Ethernet transport solution	22
7.3 Ethernet transport solution for ALA	22
7.4 AUC Endpoint maps	23
7.4.1 S and C VLAN tags in ALA	23
7.4.2 AUC End point mappings at the UNI	23
7.4.3 AUC end point mappings at the NNI	26
8 Quality of Service	27
8.1 Classes of Service	27
8.2 Bandwidth Profiles	27
8.3 Bandwidth Profile Enforcement	28
8.4 Service Level Specification	29
8.5 Multiple ALA Users at a customer premises	29
9 Ethernet ALA multicast architecture	30
9.1 Supporting multicast services using multicast AUCs	30
9.2 Supporting multicast services over point to point AUCs	31
9.3 Multicast AUC End Point mappings	32
9.3.1 Multicast AUC end point mappings at the UNI	32
9.3.2 Multicast AUC end point mappings at the NNI	32
9.4 Multicast capability at the end user premises	32

9.5	Multicast capability in the ALA provider network	33
10	ALA CPE Management	35
Annex A (informative): ALA Quality of Service implementation options		36
A.1	QoS solution components	36
A.1.1	Jitter Control and Queue Allocation.....	36
A.1.2	Class A	36
A.1.3	Class B	37
A.1.4	Classes C and D	37
A.2	An example implementation	37
A.3	ALA Queues and Scalability	40
A.4	Queuing Implementations for GPON and Point to Point Ethernet	40
History		41

Intellectual Property Rights

IPRs essential or potentially essential to the present document may have been declared to NICC. Pursuant to the NICC IPR Policy, no investigation, including IPR searches, has been carried out by NICC. No guarantee can be given as to the existence of other IPRs which are, or may be, or may become, essential to the present document.

Foreword

This NICC Document (ND) has been produced by NICC Ethernet Working Group.

Introduction

This document describes the architecture for Ethernet Active Line Access.

1 Scope

This document specifies the architecture required for the support of Ethernet Access and Ethernet Interconnect in order to support the services defined within the context of Active Line Access. The architecture covers the solution within the ALA provider and backhaul domains. It does not cover end user networks or service provider networks.

Active Line Access is technology agnostic and for this reason the architecture is careful to avoid references to specific technologies, other than as examples. For instance the architecture does not seek to explicitly define what physical elements are deployed within an ALA provider's network as this will vary with technology and commercial considerations.

2 References

For the particular version of a document applicable to this release see ND1610 [1].

NOTE: While any hyperlinks included in this clause were valid at the time of publication NICC cannot guarantee their long term validity.

2.1 Normative references

The following referenced documents are indispensable for the application of this document. For dated references, only the edition cited applies. For non-specific references, the latest edition of the referenced document (including any amendments) applies.

- [1] ND1610 Next Generation Networks, Release Definition
- [2] ND1642 Requirements for Ethernet Interconnect and Ethernet ALA
- [3] ND1030 Ethernet ALA Service Definition
- [4] ND1031 ALA UNI Specification
- [5] ND1036 ALA NNI Specification

2.2 Informative references

- [i.1] IEEE 802.3™-2005: “IEEE Standard for Information technology- Telecommunications and information exchange between systems-Local and metropolitan area networks--Specific requirements Part 3: Carrier Sense Multiple Access with Collision Detection (CSMA/CD) Access Method and Physical Layer Specifications.” <http://standards.ieee.org/getieee802/802.3.html>
- [i.2] MEF 4: “Metro Ethernet Network Architecture Framework - Part 1: Generic Framework”. May 2004.
http://www.metroethernetforum.org/MSWord_Documents/MEF4.doc
- [i.3] Technical Report DSL Forum TR-101 Migration to Ethernet-Based DSL Aggregation <http://www.broadband-forum.org/technical/download/TR-101.pdf>
- [i.4] MEF 6.1 Ethernet Services Definition – Phase 2
http://metroethernetforum.org/PDF_Documents/MEF6-1.pdf
- [i.5] IEEE 802.1ad-2005 IEEE Standard for Local and Metropolitan Area Networks -- Virtual Bridged Local Area Networks-- Revision -- Amendment 4: Provider Bridges <http://standards.ieee.org/getieee802/download/802.1ad-2005.pdf>
- [i.6] IEEE Std 802.1Q™-2005: “IEEE Standard for Local and Metropolitan Area Networks—Virtual Bridged Local Area Networks”.
<http://standards.ieee.org/getieee802/download/802.1Q-2005.pdf>
- [i.7] MEF Technical Specification MEF 26 “External Network Network Interface (ENNI) – Phase 1” http://metroethernetforum.org/PDF_Documents/technical-specifications/MEF26.pdf
- [i.8] ETSI Directives: “ETSI drafting rules Clause 14a:- Verbal Forms For The Expression Of Provisions”. Version 26, July 2009.
http://portal.etsi.org/Directives/26_directives_july_2009.pdf

3 Key Words, Definitions and Abbreviations

3.1 Key Words

The key words “shall”, “shall not”, “must”, “must not”, “should”, “should not”, “may”, “need not”, “can” and “cannot” in this document are to be interpreted as defined in the ETSI Drafting Rules [i.8].

3.2 Definitions

For the purposes of the present document, the following terms and definitions apply:

ALA domain: span of control of the ALA-provider

ALA-provider: Operator of the access network segment supporting Ethernet ALA

ALA-user: Direct user of Ethernet ALA

ALA-user connection: connection between the UNI and NNI supported by Ethernet ALA for each ALA-user

Channel: A multicast channel is defined by IETF RFC 4607 as the multicast stream defined by the combination of an SSM (Source Specific Multicast) destination address and a specific source, e.g., an (S,G) pair. For ALA a multicast channel is also defined for a non SSM address as being the multicast stream defined by *,G i.e. the multicast destination address also known as the multicast group address.

Customer Premises Equipment: equipment provided and operated by the ALA-user or end-user

NOTE: The terms ‘ALA-user CPE’ and ‘end-user CPE’ are used within the text where it is necessary to distinguish between the two.

End-user: Ultimate recipient of services provided over ALA

NOTE: End-users include both residential consumers and business users.

Ethernet ALA: An ALA Ethernet service as defined by the NICC Ethernet Working Group between the serving exchange and the customer premises provided by the ALA-provider to the ALA-user

Network Termination Unit: device provided and operated by the ALA provider at the customer premises that terminates the network of the ALA provider and provides the UNI

User-Network Interface: interface between the ALA-provider and the ALA-user at the customer premises.

3.3 Abbreviations

For the purposes of the present document, the following abbreviations apply:

AIS	Alarm Indication Signal	(Y.1731)
ALA	Active Line Access	(Ofcom)
ADSL	Asymmetric Digital Subscriber Line	(G.992)
ASM	Any-Source Multicast	(IETF RFC3569)
AUC	ALA User Connection	
CBS	Committed Burst Size	(MEF 10)
CIR	Committed Information Rate	(MEF 10)
CPE	Customer Premises Equipment	
DHCP	Dynamic Host Configuration Protocol	(IETF RFC 2131)
EBS	Excess Burst Size	(MEF 10)
EIR	Excess Information Rate	(MEF 10)

ENNI	External Network-Network Interface	(MEF 10)
EPL	Ethernet Private Line	(MEF 6.1)
IGMP	Internet Group Management Protocol	(IETF RFC 3376)
L2CP	Layer 2 Control Protocol	(MEF 10)
MEG	Maintenance Entity Group	(ITU-T Y.1731)
MEP	MEG End Point	(ITU-T Y.1731)
MIP	MEG Intermediate Point	(ITU-T Y.1731)
NNI	Network-Network Interface	
NTU	Network Termination Unit	
OAM	Operation, Administration and Maintenance	(ITU-T Y.1731)
PPPoE	Point-to-Point Protocol over Ethernet	(IETF RFC 2516)
QoS	Quality of Service	
S-Tag	Service VLAN Tag	(IEEE 802.1Q)
SSM	Source-Specific Multicast	(IETF RFC3569)
TPID	Tag Protocol Identifier	(IEEE 802.1Q)
UNI	User-Network Interface	
VDSL2	Very high speed Digital Subscriber Line 2	(ITU-T G.993)
VLAN	Virtual Local Area Network	(IEEE 802.1Q)
xDSL	Generic DSL service (may be ADSL,VDSL or SDSL)	

4 Introduction to Active Line Access (ALA)

Active Line Access enables Next Generation Access (NGA) networks to provide connectivity between residential and business consumers and their Service Providers in an open and flexible way. It provides a technology agnostic connectivity solution, being applicable to DSL, PON and also Active Ethernet (point to point) access networks. It provides a solution that allows a tier one network provider to offer logically unbundled access and it can also be used by a small community network operator as an industry standard interconnect to allow their community to connect to any number of Service Providers.

ALA uses Ethernet transport to provide a logical connection between a residential or a business customer and their chosen service provider over a physical access network that is owned by a third party network infrastructure provider. This logical connection is defined at the Ethernet layer allowing the Service Provider maximum freedom in how they wish to build their service and in this way it differs markedly from previous Wholesale Broadband solutions which typically operated using PPP and L2TP or IP. ALA also supports Quality of Service which enables delay critical real time service such as voice and live video, and Multicast which enables cost effective deployment of broadcast services.

ALA has been defined by NICC to satisfy two sets of requirements based on the output of industry consultations. One set of requirements originated from the UK regulator Ofcom and the other from a UK industry body NGN UK. These requirements were combined to create the full set of ALA requirements as described in [2]

4.1 An overview of ALA documentation

ALA is fully defined in the following NICC documents.

- This Architecture document
- The ALA Service definition [3]
- The ALA UNI definition [4]
- The ALA NNI definition [5]

This document describes the concepts of ALA, the architecture and reference points and it also provides some examples of possible ALA solutions.

The Service definition describes the ALA service in detail and how it is realised using Ethernet transport technology. This includes definitions of Service Attributes that can be used by a network operator offering an ALA service.

The ALA architecture defines two interfaces between the network operator providing the ALA service and the service provider using that service. A UNI¹ which exists at the premises of the residential or business consumer and an NNI which exists at the boundary between the Service provider's network and the network operator who is providing the ALA service.

The UNI and NNI documents should be read in conjunction with the Service definition.

¹ The term UNI does not map to the UNI reference point in MEF 4 [i.2]

NICC ALA specifications describe the service and the interfaces between the network operator providing the ALA service and the service provider or network operator using the ALA service. They do not define the interface between the service provider and the consumer which is wholly in the remit of the service provider. Similarly the NICC ALA specifications do not define any interfaces that are internal to the ALA Provider's Network.

NICC ALA specifications leverage the industry standards defined by the Broadband Forum, the Metro Ethernet Forum and IEEE 802.

4.2 ALA Business Entities

The ALA service is defined in terms of three specific business entities each of which has a specific role with regard to the ALA service. These business entities are described here.

- The **ALA provider** is responsible for the provision of the active and passive infrastructure over which ALA is delivered. The ALA provider offers standardised interfaces to which the ALA user can connect, and delivers the ALA user's traffic between these interfaces, across the ALA domain. The ALA domain extends from the end user premises to an interconnect point further up the network. ALA providers may include both fixed and some types of wireless infrastructure providers. They may own or lease the passive and active parts of network, e.g. an ALA provider may own the active electronics, but lease the passive infrastructure.
- The **ALA user** purchases Ethernet transport to an end user from the ALA provider over which it delivers services such as voice, video and internet connectivity. The ALA user has a direct, contractual relationship with the ALA provider. The ALA user may also have a direct relationship with the end-user or with other communications providers on a wholesale basis. ALA users may include ISPs and triple-play operators.
- The **End user** is the ultimate recipient of services provided over ALA. End users include both residential consumers, and business users. The End user is unaware of the ALA service and of the interfaces between the ALA provider and the ALA user.

An End user is typically (but not always) served by a single ALA provider². The end user may buy services from multiple ALA users who would in turn buy service from the ALA provider(s) serving the end user. This means that at a given end user premises it is possible to have multiple ALA users consuming the service offered by a single ALA provider.

In addition to the above business entities, ALA also supports the concept of a **Backhaul provider**. A backhaul provider is a third party network operator offering connectivity between an ALA User's premises and a given ALA Provider's network. A network operator acting as a backhaul provider for an ALA service allows ALA users who lack large scale network infrastructure to get connectivity to the ALA providers who own the access networks that connect to the end users. It is not mandatory for an ALA provider or an ALA user to make use of the service of a backhaul provider. If the backhaul provider wishes to offer the ALA user a full ALA service (as described in this document) then they should give the ALA user an NNI that is equivalent to that offered by the ALA provider. It is possible that a backhaul provider may wish to offer alternative transport solutions, however these are out of scope of the ALA specifications.

² Where an end user is served by more than one ALA provider then each ALA provider will have their own separate access line to the end user.

4.3 Network topologies

Ethernet Active Line Access is designed to deliver services from service providers (the ALA users) in the core of a network to residential and business end users (the end users). It is specifically intended to provide an alternative to passive unbundling in those parts of the access network where this is not economically desirable.

The topology of access networks is such that they connect a large number of end users into the core network at a limited number of points at the edge of the core network. Typically all traffic from the end users in a given part of the access network will be passed into the core via one resilient edge node. Note this edge node may offer geographical resilience utilising Ethernet technologies such as Spanning Tree, VLAN protection, or Multi-Chassis Link Aggregation.

This allows the ALA service definition to assume a simpler connectivity solution than might be common in a typical meshed enterprise network. This topology can be seen in the following diagram, which shows the Network to Network interfaces (NNI) between the ALA users, ALA providers and backhaul providers.

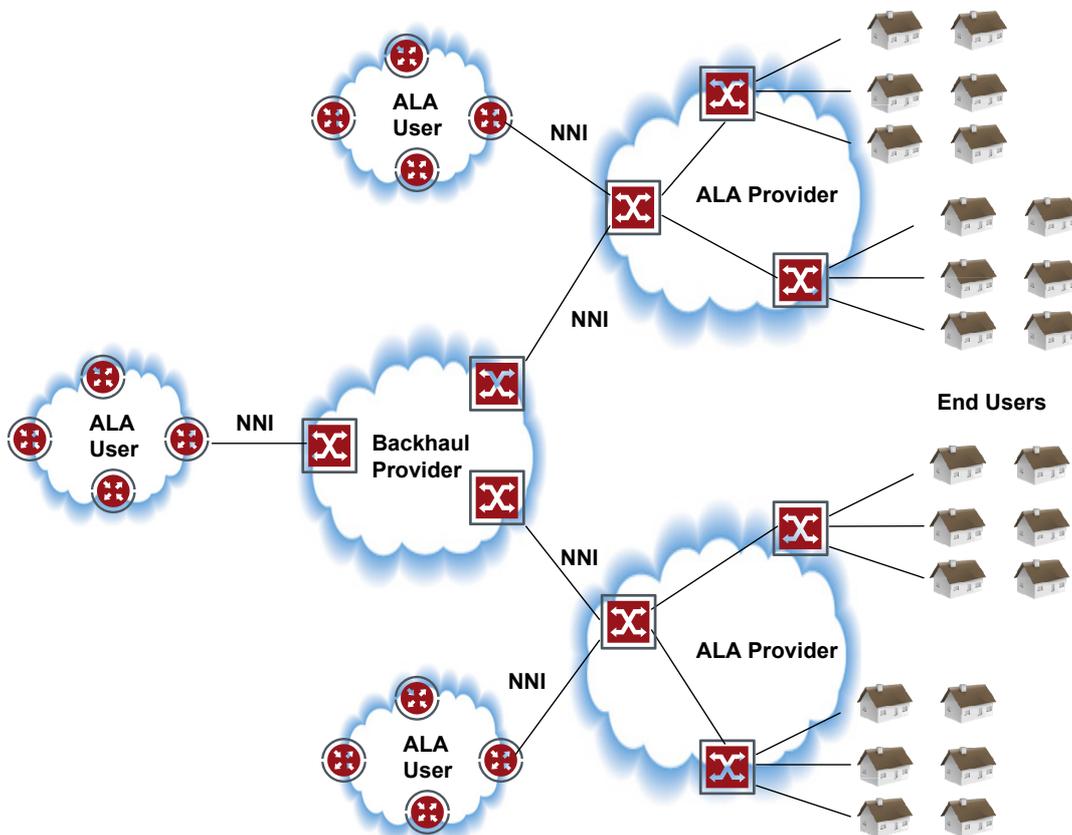


Figure 4.1 An example network topology

An ALA user may connect directly to an ALA provider or use the services of a third party backhaul provider, however they reach an individual end user via a single (possibly protected) NNI.³ The internal topology and switching within the ALA providers network, the ALA user's network or the backhaul provider's network is outside of the scope of ALA.

³ If the ALA user avails themselves of the services of a backhaul provider then an individual end user is reached from a single NNI at the ALA User to backhaul provider interface. The backhaul provider will themselves reach the end user via a single NNI at the backhaul provider to ALA provider interface, however the ALA user has no visibility of this interface.

The ALA architecture supports interconnect to an ALA provider at any number of locations, for example a street cabinet, a local exchange, a metro node or a core node. The interconnect points that are actually offered by an ALA provider as part of their service are however outside of the scope of the ALA specification.

5 ALA Service Overview

The ALA service is defined around the ALA User Connection (AUC), which provides an Ethernet connection between the ALA user network and the ALA user equipment at the end user premises.

5.1 ALA User Connections

The ALA service transports Service Frames across an ALA User Connection between an ALA UNI and an ALA NNI. The ALA User Connection isolates the traffic of different ALA users within the ALA provider network. Each Service Frame is an Ethernet MAC frame [i.1].

ALA supports two different service types, a point to point ALA service and a multicast ALA service and these are described by the service definition [3] in terms of two types of ALA User Connection, a point to point AUC and a multicast AUC. A point to point AUC provides a 1:1 connection between an ALA UNI and an ALA NNI, whereas a multicast AUC allows an ALA NNI to be connected to multiple ALA UNIs as shown in the following diagrams.

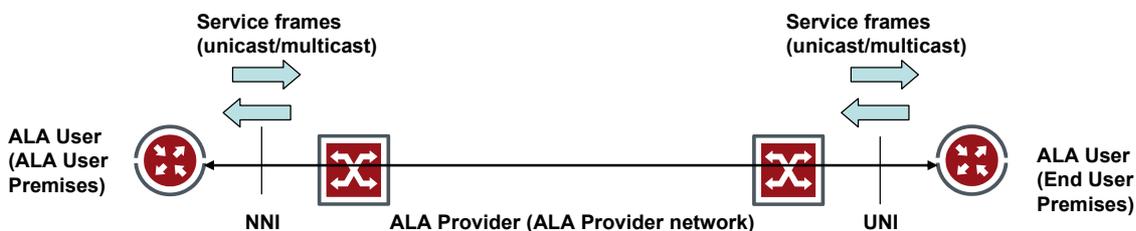


Figure 5.1 A point to point AUC

In the case of a multicast AUC service, the service frame is always a multicast frame and a higher layer protocol (IGMP) is used to determine which end users receive a given service frame.

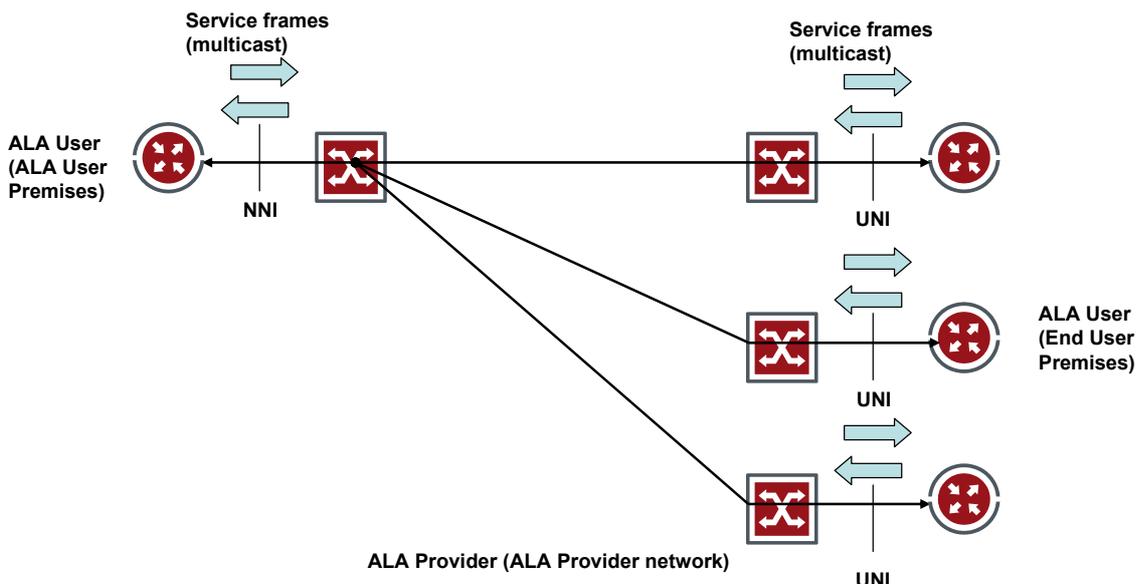


Figure 5.2 A Multicast AUC

A multicast ALA AUC does not permit service flows between UNIs.

The point-to-point ALA service is similar to the MEF E-Line service type [i.4] but is specified so that it could be implemented using the VLAN architectures specified in Broadband Forum TR-101. The multicast ALA service is based upon Broadband Forum TR-101 multicast requirements [i.3].

5.2 End to End ALA Service Definitions

The ALA service transports Ethernet frames from an ALA user to an end user over an ALA provider's network. Because ALA can be deployed over many different types of network, which each have different physical elements, these specifications do not describe ALA in terms of the nodal entities that make up the ALA provider's network. ALA is defined in terms of the ALA User Connection and the associated mapping functions at the UNI and NNI. This approach builds on the principles defined in MEF 26 [i.7]. This section provides an overview of the ALA service and its component parts.

A given end user's ALA service can be defined according to the AUCs that make up that service. An AUC that is providing connectivity between the ALA user and the end user over the ALA provider's network can be uniquely identified at the NNI and at the UNI. This allows service frames to be appropriately classified and forwarded at these points. This mechanism is described in ALA using the concepts of AUC end points and AUC end point maps and is shown in the following diagram.

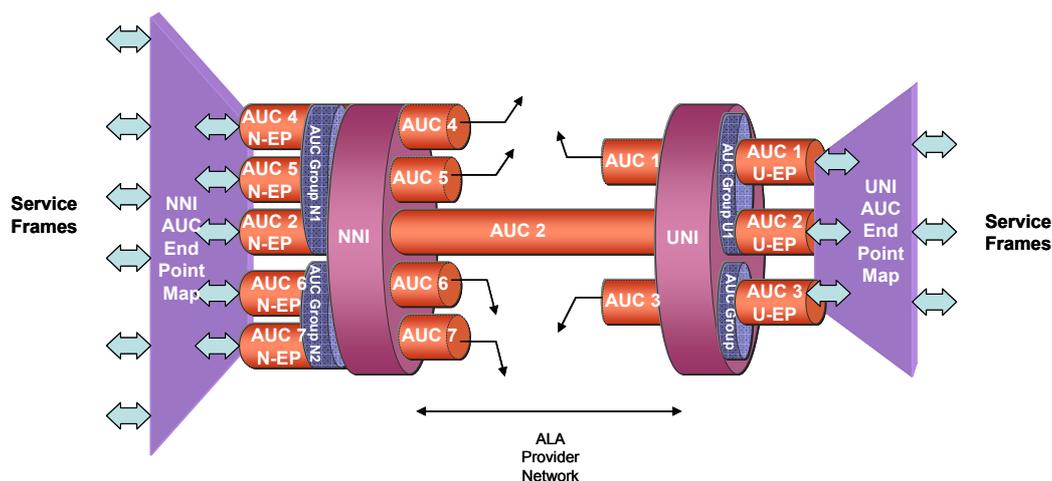


Figure 5.3 An overview of an ALA AUC and related objects

Service frames are classified and tagged appropriately at the interfaces by the UNI or the NNI AUC end point map. This maps them into or out of the appropriate AUC end point and into the AUC for transport over the provider's network.

An ALA AUC also supports quality of service as defined by a Service Level Specification and a bandwidth profile. Bandwidth profiles are created against an AUC group which is made up of one or more AUCs. AUCs that are members of the same AUC group share bandwidth, this can be used to create aggregate policers and shapers at the NNI and to allow multicast and unicast AUCs to share the same bandwidth at the UNI. It is important to note that membership of an AUC group does not imply any common tagging or path through the network or link the AUCs in any way other than by a shared bandwidth profile. An overview of ALA QoS is provided in section 8 of this document.

In the above diagram a single AUC (AUC 2) is shown end to end. Frames from the NNI for AUC2 are classified and tagged by the NNI end point map and have their class of service determined (see section 8); they enter the AUC via the NNI end point (AUC 2 N-EP) and are passed through the provider's network. At the UNI the frames "emerge" at the UNI end point (AUC 2 U-EP) and have their tagging modified according to the UNI end point map before being passed back into the ALA user's domain. Note at both the UNI and the NNI AUC 2 shares an AUC group with other AUCs but these are not destined to the same UNI. If AUC 2 is a point to point AUC then at the NNI AUC Group N1 might describe the total bandwidth to be shared by its member AUCs; while at the UNI AUC Group U1 might describe the total bandwidth available at the UNI to AUC 2 and a multicast AUC, AUC 1.

For a full definition of the Ethernet ALA service model, the AUC and its associated objects refer to the service definition [3].

6 ALA Architecture and reference points

Figures 6.1 and 6.2 show the end to end architecture for Active Line Access for two possible deployment cases. In the first case (Figure 6.1) the ALA user connects to the ALA provider network directly, in the second case (Figure 6.2) the ALA user connects to the ALA provider network using the services of a third party backhaul network provider.

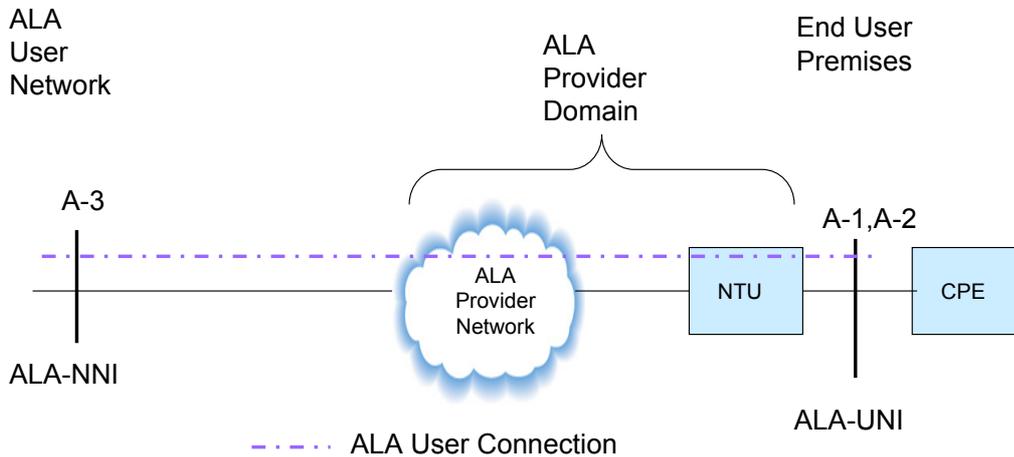


Figure 6.1 – End-to-end architecture ALA User directly connected to the ALA provider

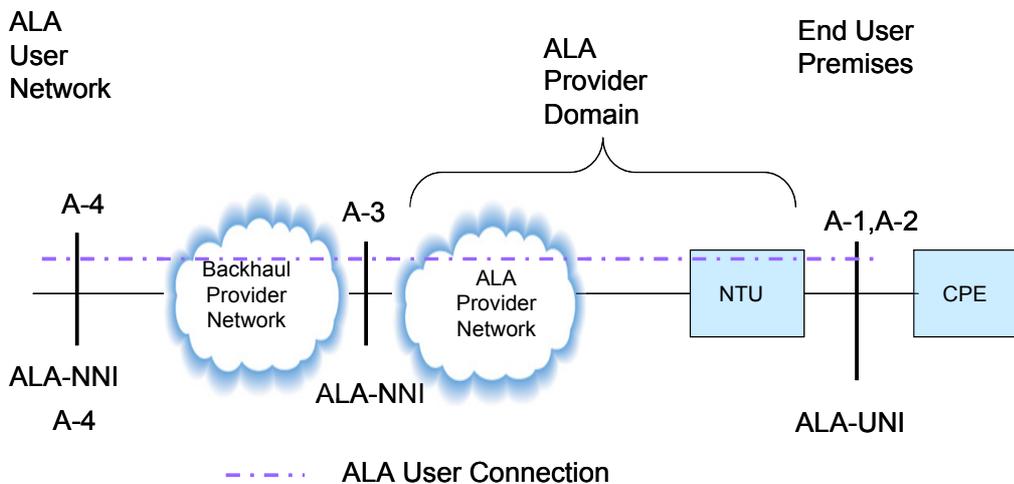


Figure 6.2 – End-to-end architecture ALA User connected to the ALA provider via a backhaul provider

The ALA provider network may use a number of technologies to transport the ALA service. For example the ALA provider network might be built from GPON, GEAPON, WDM-PON, VDSL2, ADSL2+, or Gigabit Ethernet direct fibre. Therefore although the Ethernet ALA architecture defines 4 reference points, the exact locations of functionality within the ALA provider network may vary according to the technology. Physical layer aspects of the interconnect are defined within the ALA UNI and NNI definitions [4],[5].

Next generation access networks may have different physical interfaces and fulfilment models depending on the underlying technology. Specifically a given access network may require an active NTU that is owned and managed by the access network operator, i.e. the ALA provider or it may support a wires only delivery at the customer premises. In both cases the demarcation between the

ALA provider and the ALA user in the end user premises is the NTU, however the characteristics of the NTU are very different in each case.

6.1 ALA for access networks using Ethernet aware NTUs

Where a wires only interface is not supported then the NTU contains active electronics, provides an IEEE 802.3 interface and is configured and managed by the ALA provider. The ALA provider UNI functions are supported on the NTU and the ALA User CPE may be a switch/router that connects to the NTU with an Ethernet interface. The most common example of this type of service is for GPON networks⁴ where the ALA provider owns and manages the GPON ONT. This device has a PON interface on the network side and a number of Ethernet ports on the end user side.

Note The CPE may be managed by the ALA user, or it may be a piece of un-managed CPE that is bought by the end user and configured by them according to the requirements set out by the ALA user.

6.2 ALA for wires only interfaces to access networks

For a wires only interface the NTU will be a passive device. This means that the ALA provider UNI functions will be supported from within the provider's network (at the closest piece of active electronics to the end user). The ALA user UNI functions will be supported at the CPE which is either provided by the ALA user to the end user or may be purchased by the end user and configured according to the instructions of the ALA user. The most common example of this is for existing ADSL networks where the access network is terminated at a passive device (in the UK an NTE-5) and the CPE is an ADSL modem that terminates the DSL interface and provides access to an Ethernet interface.

A wires only interface means that the ALA User and not the ALA provider is responsible for admitting frames onto the access link from the customer premises, and in particular for the upstream scheduling, and hence (upstream) quality of service on the access link. The earliest point at which the ALA provider can drop a frame is at the first piece of active networking equipment that the end customer's access link is connected to. For PON solutions the ALA provider retains some ability to control scheduling into the access link because the OLT effectively grants permission for the ONT to send traffic.

The support of a wires only interface also has an impact on the Service Level Specification, as described in [3]. The ALA provider is responsible for the performance of the ALA Service from the NNI up to NTU in the customer premises. However the ALA provider is not responsible for the performance of the ALA Users CPE (even though it is providing some ALA UNI functions). This means that where a performance issue is identified on the network between the ALA provider and the customer premises diagnostic effort is required to determine whether the problem is caused by the ALA Users CPE (an ALA User issue and not part of the SLS) or the actual transmission performance of the access link itself (an ALA provider issue and part of the SLS).

6.3 ALA and baseband voice services

Some copper access technologies, such as DSL, allow POTS/ISDN⁵ baseband access to be supported to the customer premises, these signals are filtered out from the broadband network with

⁴ It is noted that significant work is being done in the Broadband Forum on GPON interoperability which is a first step towards enabling wires only interfaces for GPONs. However the fact that GPON is a shared medium system complicates any solution for wires only presentation.

⁵ The UK Access Network Frequency Plan does not allow ISDN and DSL to co-exist on a single copper access, however other EU territories have alternative ANFPs which do permit co-existence.

the aid of a splitter at the NTU. This permits baseband POTS and ISDN services to exist independently of broadband data and voice services. This solution is out of the scope of the ALA architecture definition. However an ALA provider can support such a solution at the customer premises if they choose to do so and the support of an ALA service does not imply anything with regards to the availability of baseband.

6.4 Reference Point A-1

Reference point A-1 is at the NTU for the case where the ALA provider supplies an Ethernet NTU. The ALA user accesses the ALA service at this reference point using the ALA UNI.

The NTU is provided and operated by the ALA provider. The CPE is provided and operated by either the ALA user or the end user.

An Ethernet NTU can support multiple ALA UNIs, for multiple ALA user connections, at reference point A-1. Each point to point AUC uses no more than one UNI, and a UNI can support multiple AUCs. The UNI AUC endpoint map is used to identify the AUC for service frames passed over the UNI. The physical presentation of the UNI is Ethernet, see [4].

See section 6.8.1 for a description of physical connectivity at reference point A-1.

6.5 Reference Point A-2

Reference point A-2 is at the ALA NTU for the case where the ALA provider supplies a wires only interface. The ALA user accesses the ALA service at this reference point using the ALA UNI.

The (passive) NTU is provided and operated by the ALA provider. The CPE is provided and operated by either the ALA user or the end user. The physical presentation of the ALA UNI definition will be determined by the underlying ALA provider technology, see [4], which the CPE must terminate.

Reference point A-2 supports a single ALA UNI and the UNI supports multiple ALA user connections. The UNI AUC endpoint map is used to identify the AUC for service frames passed over the UNI.

See section 6.8.2 for a description of physical connectivity at reference point A-2.

6.6 Reference Point A-3

Reference point A-3 is between the ALA provider's network and either the ALA user's network or a network providing backhaul for the ALA service on behalf of the ALA user. (the ALA user may connect to the ALA provider directly or via a backhaul provider.) This interconnect uses the ALA NNI. The ALA NNI presents a number of ALA user connections over a physical Ethernet presentation as defined by the ALA NNI profile.

The ALA NNI may transport the traffic of a single ALA user, or it may transport the traffic of multiple ALA users over a single physical Ethernet interface. The NNI endpoint map is used to identify the AUC and the forwarding behaviour for service frames passed over the NNI.

6.7 Reference point A-4

Reference point A-4 is between a backhaul provider network and the ALA user network. The ALA user network interfaces with the backhaul provider's network using the ALA NNI. The ALA NNI

transports the traffic of a single ALA user, in all other respects it is identical to that defined at reference point A-3. This ensures that the introduction of an intervening backhaul provider's network does not have an impact on the capabilities of the ALA service.

6.8 Support for multiple ALA users at reference points A-1 and A-2

A key requirement of ALA is the ability to allow the end user to purchase service from at least two ALA users over a single ALA provider's infrastructure. The nature of the access network means that this solution is different depending on whether the ALA UNI is implemented on reference point A-1 or A-2.

6.8.1 Multiple ALA users at reference point A-1

Where the UNI is supported over reference point A-1 the CPE/NTU solution shown in figure 6.3 applies.

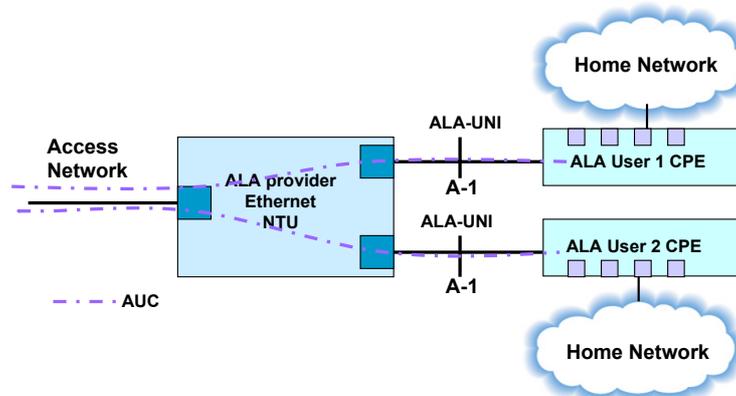


Figure 6.3 Multiple ALA users for an Ethernet aware NTU

In this case the ALA user connects to the ALA provider's network over a physical Ethernet presentation that corresponds to a physical port on the NTU. An NTU may have a number of ports (typically 2 or 4). In this solution each ALA user can connect to the ALA provider's network over their own dedicated port. The UNI AUC end point map can map service frames to AUC based on VLAN id, or in the case of a port based UNI all service frames on the interface are mapped to the same AUC.

It is possible for multiple ALA users to connect to the ALA provider over the same physical port because the VLAN id of the incoming service frames identify the AUC to which they belong.

6.8.2 Multiple ALA users at reference point A-2

Where the UNI is supported over reference point A-2 the solution is more complex because the ALA user is responsible for providing customer premises equipment to terminate the access networks physical layer (e.g. an x.DSL interface, a point to point Ethernet interface or a GPON interface). In this case one ALA user therefore controls access to the ALA UNI at reference point A-2.

In order to support multiple ALA users in this case then the ALA user who is responsible for the CPE equipment directly connected to the ALA provider needs to support the capability to transparently transport another ALA User's AUC through their CPE and present this at a physical Ethernet port. This effectively creates an instance of reference point A-1 on the end user side of their CPE as shown in figure 6.4

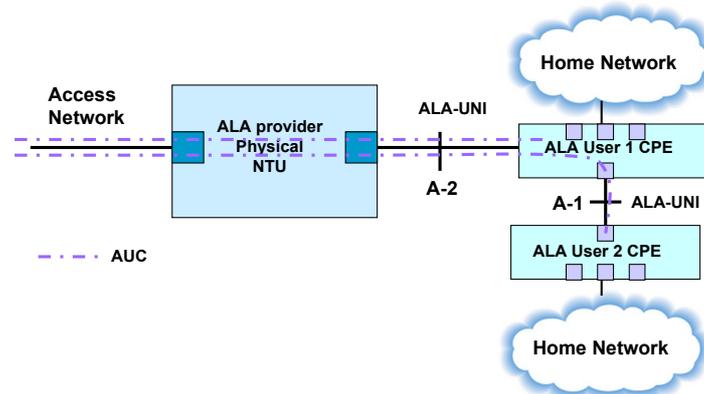


Figure 6.4 Multiple ALA users for a wires only NTU

This second ALA user connects to the network using reference point A-1 exactly as if they were directly connected to the ALA provider. While typically each additional ALA user would consume a single physical port on the user side of the first ALA user's CPE there is no reason (as described in section 7.4.2) why multiple additional ALA users could not share the same physical CPE port using the VLAN id to identify the AUC.

The UNI end point map at reference point A-2 will identify the AUC based on the VLAN tag of the ingress service frames. Note this does not imply the frames need to be tagged because the UNI end point map supports the concept of a default VLAN for an interface, see section 7.4.2

7 Ethernet Transport Solutions

ALA is defined to be implemented over an Ethernet transport network, however there are two possible architectures for providing an Ethernet transport solution as defined by the Broadband Forum [i.3]. These are the 1:1 VLAN architecture and the N:1 VLAN architecture.

7.1 1:1 VLAN Ethernet transport solution

The 1:1 VLAN architecture is similar to the MEF E-Line service [i.4].

In this architecture frames are forwarded to and from the end user based on the VLAN ID only. All traffic between an end user and the network flows within a dedicated VLAN, which may be identified by one or two VLAN tags depending on scalability requirements. Ethernet mechanisms such as broadcast, multicast and MAC learning are turned off in the network and operate only at the end user premises and the IP edge of the network. In order to support multicast services these networks typically use dedicated Multicast VLANs that each connect multiple end-users to multicast routers. Forwarding in these Multicast VLANs is configured using IGMP and end-user MAC address learning is not required.

Networks implemented using this solution are optimised to switch a large number of VLANs, do not support MAC learning and generally have a large-scale OSS to provision individual VLANs to end users.

7.2 N:1 VLAN Ethernet transport solution

The N:1 VLAN architecture is similar to the MEF E-tree service [i.4] but requires the security extensions defined by the Broadband Forum in TR-101 [i.3].

In this architecture frames are forwarded to and from the end user based on MAC address and VLAN ID. Many end users share a single VLAN. Ethernet mechanisms such as MAC learning and broadcast are enabled in the network. In order to restrict user to user communication and to protect against spoofing and denial of service attacks, these networks use functionality to limit the scope of Ethernet broadcast and to prevent duplicate or spoofed MAC addresses and IP addresses from interrupting service. In this network architecture multicast can simply be a special type of type of traffic within a shared VLAN, or separate Multicast VLANs can be used.

Networks implemented using this solution are optimised to handle a large number of MAC addresses and a relatively small number of VLANs. In addition, the Ethernet spanning-tree protocol can be used to provide resilience within the network. Typically they would have a smaller OSS and dynamic MAC address learning means less work is required to provision an end user.

7.3 Ethernet transport solution for ALA

ALA can be implemented over either Ethernet transport solution however there are differences in the implementation depending on the choice of transport solution. Therefore within the ALA specifications NICC has defined ALA for both 1:1 VLANs and N:1 VLANs.

For a given ALA deployment the choice of Ethernet transport solution is dependent on the capabilities of the ALA provider's network and OSS. It is the choice of the ALA provider which model they support and this will be explicitly stated as part of their published ALA service

definition. An ALA provider may chose to support only the 1:1 VLAN solution or only the N:1 VLAN solution or both solutions.

Interworking between 1:1 and N:1 VLAN solutions is not addressed by this document.

7.4 AUC Endpoint maps

The AUC endpoint maps exist at the UNI and the NNI and they define how the AUC is identified at the end point and any required VLAN tagging. The mapping of an AUC to VLAN at the NNI depends on whether the underlying Ethernet transport network is using a 1:1 VLAN architecture or an N:1 VLAN architecture.

7.4.1 S and C VLAN tags in ALA

The terms S-tag and C-tag are widely used within Ethernet networks as defined by IEEE 802.1ad [i.5], and also within broadband networks as defined by Broadband Forum TR-101 [i.3]; however the two standards are not consistent in their approach. Ethernet ALA follows the Broadband Forum architecture which means that an ALA provider may use either single tagging (an S-tag) or double tagging (an S-tag and a C-tag) to identify AUCs at the NNI.

IEEE802.1ad defines an S-tag as having a TPID of 0x88A8 and a C-tag as having a TPID of 0x8100, however there remain a significant number of network operators who use a technology known as Q-in-Q to support double tagging in their network. These networks use two tags with a TPID of 0x8100 where the first tag is treated as an S-tag with a TPID of 0x8100.

Ethernet ALA defines the S-tag as being the first (i.e. outermost) tag presented at a UNI or an NNI that is used by the AUC endpoint to map the frame to an AUC. The TPID of the S-tag may be either 0x88A8 or 0x8100 depending on the service attributes supported by the ALA provider. The ALA provider will explicitly define this as part of their service definition.

ALA defines the C-tag as being the second tag presented at an NNI that is used by the AUC endpoint to map the frame to an AUC. A C-tag, if present, will always have a TPID of 0x8100.

In ALA VLAN tags that are not significant to the ALA provider are passed transparently as traffic, for example at a single tagged ALA NNI the first tag in a frame would be identified as the S-tag, if a second tag was present it would be treated as user data.

Untagged frames are frames without VLAN tags received at the UNI or NNI. Priority tagged frames are frames which contain a VLAN tag value of zero but carry priority bits.

A frame where a VLAN tag has a TPID value that does match the TPID in force at the UNI/NNI will always be treated as an untagged frame. At the NNI untagged frames are not supported and will be dropped.

7.4.2 AUC End point mappings at the UNI

ALA allows three options for the UNI presentation.

- A port based UNI.
- An S-tagged UNI.
- A customer edge port based UNI

At reference point A-1 the ALA provider shall offer an S-tagged UNI and may offer a Port based UNI and/or a Customer edge port based UNI.

Note: The S-Tagged UNI requires the ALA User to identify the AUC of a frame using either an S-VLAN tag or the default S-VLAN on the port. This means that to use a VLAN-tagged presentation at the UNI in the S-Tagged interface, the ALA User needs to send two VLAN tags at the NNI and configure the S-VLAN to use at the UNI.

A customer edge port based UNI offers the same capability but in this case the C-VLAN tags used at the UNI are tunnelled over a point-to-point AUC. This means that to use a VLAN-tagged presentation at the UNI the ALA User will need to send three VLAN tags at the NNI; where the two outer VLAN tags are defined by the AUC end-point map at the NNI and the inner-most VLAN tag needs to match the VLAN tag being used by the ALA User at the UNI. This inner-most VLAN tag is not significant to the ALA provider at the NNI and is carried transparently over it.

At reference point A-2 the ALA provider shall offer an S-tagged UNI if there are multiple ALA AUCs.

The ALA end point mappings at the UNI behave the same way in the N:1 and 1:1 VLAN architectures.

7.4.2.1 Frame handling at a port based UNI

At a port based UNI the UNI identifier identifies the AUC. All service frames received on the UNI that have a TPID of 0x8100 shall be accepted regardless of any VLAN tagging and treated as untagged frames⁶. All service frames received on the UNI that have a TPID of 0x88A8 should be accepted regardless of any VLAN tagging and treated as untagged frames⁷. VLAN tagging is preserved through the ALA provider's network. In the downstream direction (towards the end user) any VLAN tags used by the ALA provider to identify the AUC will be stripped before the frame is passed over the UNI. Note that this mode only supports a single ALA Class of Service.

7.4.2.2 Frame handling at an S-tagged UNI

At an S-tagged UNI the AUC end point map uses the S-VID of the frame to identify the AUC. The UNI AUC end point map will therefore contain a list of S-VID to AUC mappings. Traffic arriving over the AUC will be passed out over the UNI with the appropriate S-tag. An S-VID will map to at most one AUC and an AUC will map to only one S-VID.

At each S-tagged UNI it is possible to define a default VLAN to which any untagged or priority tagged frames are mapped. This default VLAN can be mapped to an AUC (as for any other S-tagged VLAN) by the end point map. If no default VLAN is defined for the UNI then untagged and priority tagged frames will be dropped since they do not match a defined S-VID. This behaviour is summarised in the following table.

⁶ Priority tagged frames are also treated as untagged frames, i.e. the .p bits will not be used to determine the CoS of the traffic.

⁷ The MEF EPL service which the port based UNI is intended to support is only defined for C-tagged traffic (TPID 0x8100). Network equipment that conforms to IEEE 802.1ad may drop S-tagged (TPID 0x88A8) frames received on a port based interface. Therefore transparent transport of frames with a TPID of 0x88A8 may not be supported by an individual ALA Provider and this should be declared in their product definition.

Upstream frame	Action	Downstream frame	Action
S-tagged	Admitted to an AUC if it matches an AUC<->S-tag entry else dropped	Arriving at on an AUC that maps to an S-tag	Forward over the UNI with the S-tag
Untagged / priority tagged	Admitted to the AUC that maps to the default VLAN. If no such AUC exists then dropped.	Arriving on an AUC that maps to the default VLAN	Forwarded untagged over the UNI
Tagged with an incorrect TPID	Treated as an untagged frame	N/A	N/A

7.4.2.3 Frame handling at an customer edge port UNI

At a customer edge port UNI, one C-Tag value is used to identify the multicast AUC⁸. This value needs to be defined by the ALA Provider as part of their product description, All other frames, untagged, priority tagged or tagged with a value other than the multicast tag, are mapped to the point to point AUC. The priority within the AUC for frames entering the network at the UNI is mapped from the priority bits of C-tags received.

This behaviour is summarised in the table below.

Upstream frame	Action	Downstream frame	Action
C-tagged	Admitted to the multicast AUC if it matches the multicast AUC<->C-tag entry, all other C-tags are mapped to the point to point AUC. The C-tag received at the	Any frame arriving on an AUC present at the UNI.	For a Point-to-Point AUC, forwarded over the UNI with all ALA provider tags stripped (i.e. the ALA S-tag and any ALA C-tag used by the provider). Note this MAY mean that the frame is presented over the interface with a C-tag (if a third tag was provided at the NNI by

⁸ The Multicast AUC may be mapped to the default VLAN and in this case would be untagged at the UNI.

	UNI is preserved but is treated as user data NOT a C-tag at the NNI.		the ALA user). For a Multicast AUC, forwarded over the UNI with the configured C-Tag.
Untagged / priority tagged	Admitted to the AUC that maps to the default VLAN.	Arriving on an AUC that maps to the default VLAN	Forwarded untagged over the UNI

7.4.3 AUC end point mappings at the NNI

ALA allows two options for an AUC at the NNI.

- A Single tagged AUC where the AUC is identified by the S-tag
- A double tagged AUC where the AUC is identified by the combination of S-tag and C-tag.

Untagged and priority tagged frames are dropped at the NNI.

The AUC end point mappings in an N:1 VLAN architecture are different to those defined for a 1:1 VLAN architecture.

7.4.3.1 Frame handling at the NNI in a 1:1 VLAN architecture

In a 1:1 VLAN architecture the AUC can be identified at the AUC end point from the VLAN tags. For each single tagged AUC the NNI AUC end point map contains a mapping from AUC to S-tag values. For each double tagged AUC the NNI AUC end point map contains a mapping from AUC to S-tag and C-tag values. Frames arriving at an NNI with an incorrect TPID will be treated as untagged traffic and dropped. At most one AUC is identified by an S-tag (single tagged AUC) or S-tag, C-tag combination (double tagged AUC). Similarly an AUC maps to one S-tag (single tagged AUC) or one S-tag, C-tag combination (double tagged AUC).

7.4.3.2 Frame handling at the NNI in a N:1 VLAN architecture

In a N:1 VLAN architecture many AUCs may share the same S-tag. Double tagged AUCs are not supported in a N:1 VLAN architecture. In this architecture the ALA provider and ALA user must configure the S-tag to be used for a set of AUC end points at the NNI but the actual AUC is identified by a combination of the S-tag and the MAC address.

The MAC address for a given AUC is learnt dynamically at the NNI, end user isolation is achieved by adding the security mechanisms defined by the Broadband Forum in [i.3] to the standard Ethernet MAC learning mechanisms defined in [i.6]. In the upstream direction the AUC is identified by the combination of S-tag and source MAC address, and in the downstream the AUC is identified by the combination of S-tag and destination MAC address.

An AUC can map to only one S-tag value but a given S-tag value may map to many AUCs. An individual AUC is identified by S-tag and at least one MAC address. For a given S-tag multiple MAC addresses may map to the same AUC. Frames for which no AUC can be identified are dropped by the ALA provider network.

8 Quality of Service

Ethernet ALA supports quality of service mechanisms that enable ALA Users to offer services that require a defined Ethernet performance. On ingress to the ALA provider network at the UNI or NNI each service frame is mapped to a class of service. This class of service in combination with the AUC classification is used to map each service frame to a bandwidth profile. Each class of service has a service level specification that defines the performance objectives that must be met for packets of that class.

8.1 Classes of Service

Ethernet ALA supports four classes of service as shown in Table 8.1.

Table 8.1: ALA Classes of Service

Class	Typical Use
A	Realtime, delay sensitive applications e.g. voice
B	Streaming applications (video)
C	Internet data
D	Guest or 3rd party Access

Each of these classes has associated performance objectives that form a per-class Service Level Specification. These performance objectives are such that Class A will have better performance than Class B, which will have better performance than Classes C and D.

In order to meet the Service Level Specification, it is expected that the ALA Provider will need to implement strict priority scheduling at any congestion points in their network to prioritise transmission of Class A traffic over Class B traffic, which in turn would be prioritised over Class C and D traffic. Starvation of the lower priority queues can be avoided by the use of per Class Policers. Example implementations are described in Annex A of this document.

Classes A and B support only committed bandwidth. The bandwidth available for each of Classes A and B may need to be restricted by the ALA provider to ensure that performance guarantees can be delivered for lower priority classes and other ALA users.

Classes C and D support both committed and excess bandwidth. The bandwidth profiles at the UNI and NNI can be configured to be colour aware so that the ALA User's drop precedence marking is respected within these classes. In the case that both classes C and D send excess traffic at the same time, the ALA provider will limit the bandwidth share of Class D. Typical use cases for Class D would be to support (wireless) guest access at the end-user premises, or to limit the bandwidth of a background application such as push video.

8.2 Bandwidth Profiles

ALA uses the bandwidth profile algorithm defined in MEF26 [i.7]. This supports two information rates (committed and excess) and colour awareness. Figure 8.1 shows the ALA bandwidth model.

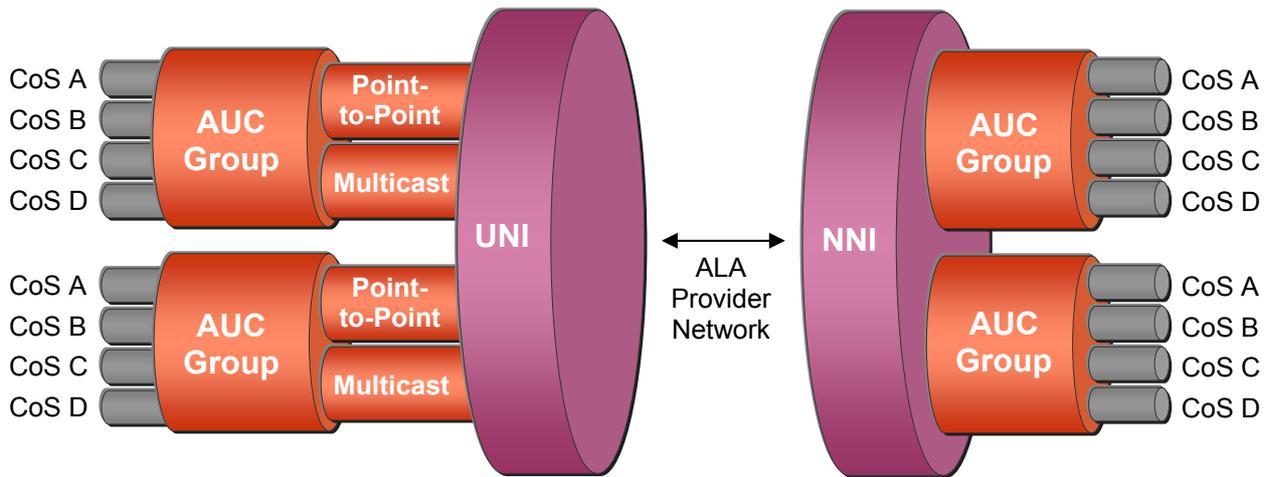


Figure 8.1: ALA Bandwidth Model

Separate bandwidth profiles are defined for ingress and egress frames at an interface. This supports asymmetric upstream and downstream bandwidth allocation. At the UNI and NNI, bandwidth profiles are defined per class of service. There is a separate set of bandwidth profiles for each interface and AUC Group. Each AUC Group is associated with a single interface and contains one or more AUCs. In the case where there is only one AUC in each AUC Group, this model is aligned with the MEF 26 Bandwidth Profile per Class of Service Identifier model.

AUC Groups at the UNI provide a way for the ALA provider to define the total amount of bandwidth offered to an ALA User who uses a point-to-point and a multicast AUC. A single set of bandwidth profiles may be shared by one Point-to-Point AUC and one Multicast AUC. This supports a concept of ‘video bandwidth’ where an ALA User can choose whether to send video using multicast or unicast delivery but must not break the ALA provider’s contracted capacity limit at the UNI. In the absence of a multicast service, the Group membership at the UNI is one, i.e. the bandwidth profile applies to a single point-to-point AUC. Multiple AUC groups may be present at a UNI to support multiple ALA Users.

Aggregate bandwidth profiles at the NNI are useful in cases where NNI or backhaul bandwidth for AUCs is contended. AUC Groups at the NNI allow a single bandwidth profile to be shared by multiple AUCs. This bandwidth profile may be used by the ALA Provider to configure an aggregate policer per class of service for ingress traffic at the NNI and an aggregate shaper per class of service for egress traffic at the NNI. The AUC Group at the NNI will be identified by an S-VLAN tag that is used by a single ALA User at that NNI. Multiple AUC groups may be present at a given NNI to support multiple ALA Users.

8.3 Bandwidth Profile Enforcement

Having defined the bandwidth profiles that apply to service frames at the UNI and NNI, the ALA provider can implement per-Class-of-Service policing and shaping functions to ensure that ingress and egress traffic is compliant with these bandwidth profiles. When an AUC Group at the UNI contains traffic from both a Point-to-Point and Multicast AUC, the policing / shaping function will need to be between the last multicast replication point and the UNI.

The ingress bandwidth profile is used to implement policing on ingress to the network (see MEF 26 section 7.6.2). The MEF bandwidth profile algorithm will mark a frame as Green, Yellow or Red. Frames marked Yellow by the ingress bandwidth profile may be delivered by the AUC, but SLS

performance objectives do not apply for these frames. Frames marked Red by the ingress bandwidth profile are dropped.

The egress bandwidth profile regulates the amount of egress traffic that will be sent across a particular NNI or UNI (see MEF 26 section 7.6.3). The ALA provider can shape traffic in order to ensure that a egress traffic complies with an egress bandwidth profile. Traffic shaping is discussed in MEF 10.2 section 10.3. Given an egress bandwidth profile with parameters (CIR, CBS, EIR, EBS) it is possible to configure either a single rate or double rate shaper. The queue length of the shaper will need to be configured such that the maximum delay introduced by the queue is allowed by the end-to-end performance target. A queue management algorithm such as WRED can be used to ensure lower frame loss for frames marked green.

For a single rate shaper, the shaping rate could be configured as CIR + EIR. CBS and EBS will be bounded by the maximum output burst of the shaper. In order to limit the number of egress frames marked green to the CIR, a single-rate egress policer would need to be used.

MEF 10.2 section 10.3 describes a double-rate shaping algorithm. This algorithm attempts to smooth the traffic to improve its conformance with the CIR and will only send yellow frames up to the maximum rate of the EIR when the egress buffer becomes full.

8.4 Service Level Specification

The AUC Service Level Specification specifies the frame delivery performance objectives between the ALA UNI and NNI. The ALA Service Definition includes SLS attributes for Frame Delay Performance, Inter-Frame Delay Variation Performance, Frame Loss Ratio Performance, and Availability Performance. The SLS for Multicast also includes an attribute for channel change latency. These attributes can be measured by either the ALA User or the ALA Provider. The ALA provider may choose to either monitor conformance of an AUC to these performance objectives pro-actively or only take steps to measure the performance in the case of a dispute with the ALA User.

8.5 Multiple ALA Users at a customer premises

If multiple ALA users are supported by the ALA provider at a single customer premises then the bandwidth available to all of the UNIs at the customer premises needs to be shared between the AUCs that terminate on those UNIs. The UNI bandwidth is first likely to be partitioned into committed and excess bandwidth. Committed bandwidth at all of the UNIs will be split between the committed information rates in each class for each AUC Group. The ALA provider will meet the standard service level specification applying to committed bandwidth for each class of service for each ALA User at a UNI. The excess bandwidth will be shared between AUCs according to the policies of the ALA provider.

9 Ethernet ALA multicast architecture

ALA provides multicast AUCs to support multicast within the ALA provider domain. ALA point to point AUCs can also be used to support multicast services by ALA users however in this case multicast is not provided within the ALA providers network.

9.1 Supporting multicast services using multicast AUCs

Multicast capability within an ALA provider's network offers potential bandwidth savings for ALA users and to take advantage of this, ALA supports a Multicast ALA User Connection (Multicast AUC).

An ALA Multicast AUC allows an ALA user to inject multicast traffic into an ALA provider's network at the NNI and have this stream replicated and delivered to appropriate members of the Multicast AUC. A multicast AUC permits multiple ALA end users UNIs to receive frames from a single AUC at the NNI. Multicast traffic is delivered downstream from the NNI either unconditionally (all multicast traffic injected at the NNI is forwarded to all UNIs that are connected by the Multicast AUC) or conditionally (multicast traffic injected at the NNI is forwarded to those UNIs that have requested the traffic using a multicast control protocol).

Where ALA is implemented using a 1:1 VLAN transport solution then a Multicast AUC must use a dedicated multicast VLAN.

Where ALA is implemented using a N:1 VLAN transport solution then it may use a VLAN that is shared with unicast connections. In this case the same VLAN is used to carry multiple unicast AUCs and a single multicast AUC.

The multicast architecture for Ethernet ALA is shown in figure 9.1

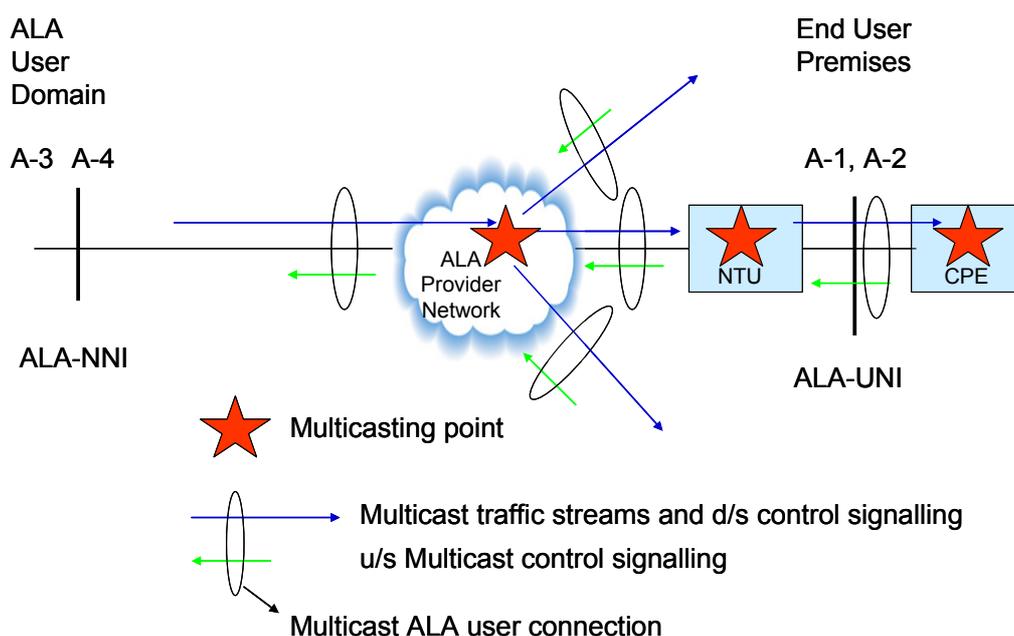


Figure 9.1 Multicast architecture

A multicast ALA user connection acts as a container for a number of multicast streams (or channels). The multicast ALA user connection has a defined bandwidth profile as for a point to point ALA user connection, a defined forwarding behaviour (conditional or unconditional forwarding), and a Service Level Specification. At the ALA NNI the multicast AUC has a bandwidth profile that defines the total bandwidth for all multicast traffic in each direction that the ALA provider will accept.

The bandwidth profile at the UNI defines the total bandwidth downstream and upstream for multicast traffic that will be passed over the UNI. Because the bandwidth profile at the NNI represents the total sum of all multicast traffic that will be accepted into the provider's network it will be greater than the multicast traffic that can be accepted at each UNI. Each bandwidth profile is likely to be asymmetrical because in the downstream direction the multicast traffic contains video/audio media streams and multicast control messages but in the upstream direction the multicast traffic only contains multicast control messages.

A multicast stream sent into the ALA provider's network within an ALA multicast connection may be replicated to any number of ALA end users that can be reached via the ALA provider's network. This replication is carried out by the ALA provider.

The mechanisms used by the ALA provider to replicate traffic are beyond the scope of the ALA architecture.

The Ethernet ALA multicast service supports IGMP snooping that allows end users to request individual multicast streams within an ALA multicast user connection. IGMP snooping is supported at the UNI.

IGMP snooping is also supported at the NNI and provides a capability to ensure that a multicast stream is only sent into the ALA provider network if at least one ALA end user served by that ALA provider network requires it.

IGMP snooping behaviour is defined as part of the ALA service definition [3] and the UNI [4] and NNI [5] definitions.

9.2 Supporting multicast services over point to point AUCs

When ALA is implemented using a 1:1 VLAN transport solution it is possible to support a multicast service over a point to point ALA User Connection. In this case the ALA provider delivers the multicast frames over the point to point ALA User Connection as for any other frames. No replication points are provided in the ALA provider's network. The ALA user processes multicast control protocol messages to/from end users within their own network and replicates multicast frames as appropriate. The ALA service provides a point to point connection between the ALA user's replication point and the end user.

The ALA user is responsible for all VLAN dimensioning for the service, the ALA user should therefore agree and purchase bandwidth from the ALA provider to suit their specific services needs. Each individual AUC can be dimensioned to support the delivery requirements at the ALA UNI (number of concurrent channels, High Definition content, Standard Definition content, etc).

There is no optimisation of bandwidth in the ALA provider's network in this case. The ALA user could choose to purchase a non-contended service in which case the sum of the individual AUC's bandwidth should be equal to the sum of all AUCs bandwidth at the NNI. Alternatively the ALA user can choose to contend the service using their own concurrency rules in which case the

aggregate bandwidth at the NNI will be less than the sum of the AUC bandwidths at the UNI. In this case the ALA user will be responsible for any admission control.

9.3 Multicast AUC End Point mappings

The UNI/NNI AUC end point map is used to identify the AUC for a Multicast AUC in the same way it is used for a point to point AUC.

9.3.1 Multicast AUC end point mappings at the UNI

The multicast AUC end point mapping behaviour depends on which UNI type is being presented.

9.3.1.1 Frame handling at a port based UNI

A multicast AUC can be mapped to a port for a port based UNI. In this case all frames received on the port will be mapped to the multicast AUC traffic. This means the port must be dedicated to the multicast service.

9.3.1.2 Frame handling at an S-tag based UNI

A multicast AUC is mapped to an S-tag as it is for point to point AUCs. A multicast AUC may be assigned to the default VLAN in which case the multicast traffic will be untagged or priority tagged at the UNI.

Since at most one AUC can map to a given S-tag if the multicast AUC is assigned to the default VLAN then it follows that all point to point AUCs (and any other multicast AUCs) at the UNI must be S-VLAN tagged.

9.3.1.3 Frame handling at a Customer edge port based UNI

A multicast AUC is mapped to the configured multicast C-tag. If the multicast AUC maps to the default VLAN for the interface then it is untagged at the interface.

9.3.2 Multicast AUC end point mappings at the NNI

A multicast AUC is identified by S-tag at the NNI, double tagging of multicast AUCs is not supported.

9.3.2.1 Frame handling at the NNI in a 1:1 VLAN architecture

In a 1:1 VLAN architecture the multicast AUC is identified by S-tag. An S-tag that is used for a multicast AUC at the NNI is not used by any double tagged point to point AUCs at that NNI.

9.3.2.2 Frame handling at the NNI in a N:1 VLAN architecture

In a N:1 VLAN architecture the multicast AUC is identified by S-tag. It may share this S-tag with zero or more point to point AUCs.

9.4 Multicast capability at the end user premises

In order to support the Ethernet ALA multicast service the CPE needs to be multicast aware and needs to act as a multicast router at the boundary of the end user network and the ALA provider network.

This means that where the end user has multiple multicast clients requesting a multicast channel from a multicast ALA user connection then the CPE must either act as a IGMP proxy router or provide IGMP proxy reporting functionality such that only sends a request to leave a channel over the UNI when the end user network no longer requires the channel.

Depending on the ALA user's service design IGMP messages may also need to be forked at the ALA user's CPE so that they are sent within both the multicast and any point to point AUCs. This is typically done in order to allow ALA user networking equipment (such as a Broadband Access Server) to shape traffic to a customer line. See TR-101 [i.3] for additional information on multicast support in broadband networks.

The ALA user is responsible for providing any CAC required to prevent an end-user from exceeding their available bandwidth and thus disrupting their own service. This means that where an end user can request multiple channels from the network the ALA user must check that the requested bandwidth does not exceed the bandwidth available to the end user (for that service class) in the ALA provider's network.

Some ALA providers may sell an aggregate bandwidth to be shared between multicast services and unicast services (for example selling video bandwidth to incorporate both multicast services and VoD). In this case the ALA user needs to ensure that the combination of unicast and multicast traffic requested by the end user does not exceed the bandwidth available (for that service class) to the end user in the ALA providers network.

Where the ALA UNI is provided via reference point A-1 then, depending on the properties of the ALA provider network, the NTU may provide multicast capability. If the ALA provider NTU does not provide multicast capability then it will be transparent to IGMP passed over the multicast ALA user connection.

Where the ALA UNI is provided via reference point A-2 then the NTU does not provide multicast functionality.

The definitions of the multicast ALA user connection and the IGMP snooping behaviour at the UNI are provided in the Ethernet ALA service definition [3].

9.5 Multicast capability in the ALA provider network

The exact mechanisms by which the ALA provider supports multicast are beyond the remit of the ALA architecture.

The ALA provider shall support multicast ALA user connections over the NNI. The ALA provider shall ensure that a multicast channel need only be sent into the network at the NNI once for each ALA multicast AUC no matter how many of the ALA end users supported by that multicast AUC require the multicast channel.

A multicast AUC may support conditional forwarding or unconditional forwarding of multicast frames. If the multicast AUC supports conditional forwarding then the ALA NNI shall support the multicast control signalling to ensure that within a multicast ALA user connection a multicast channel can be sent into the ALA provider network or pruned from the ALA provider network based on the requirements of the ALA end-users served by the ALA provider network.

An ALA provider may support multiple intermediate multicast replication points within their network prior to the network element to which the end user is physically connected. At these intermediate replication points the ALA provider may determine that it is unnecessary to support the simultaneous delivery of all of the multicast channels and bandwidth that is available to the multicast AUC at the NNI to the next replication point. For example if the NNI serves 10k end users but the last replication point in the ALA provider's network is typically connected to no more

then 100 users it is unlikely all of the multicast channels and bandwidth available at the NNI can be consumed by the end users connected to the last replication point. The dimensioning of the ALA provider's network is beyond the scope of ALA, however the ALA provider should provide a Service Level Specification that defines the number of different multicast streams and the amount of multicast bandwidth that can be simultaneously delivered per end user, or per group of end users. If the ALA provider is optimising the multicast bandwidth of their network in this way, then it becomes their responsibility to discard multicast requests that would result in an excess of multicast traffic at an intermediate replication point.

Where an NNI supports conditional forwarding, i.e. multicast traffic is delivered over the NNI in response to requests from IGMP, it is possible that more multicast traffic will be requested at a NNI than the traffic profile can support. In this case it is the responsibility of the ALA user to provide a connection admission control function that prevents excess traffic being sent over the NNI. If the ALA provider receives more multicast traffic at an NNI than the multicast AUC is dimensioned to support then they may police out excess frames⁹.

Any end user billing functions (including per-usage billing) need to be provided by the ALA user. The ALA provider is not required to provide the ALA user with detailed customer interface records or events.

Conditional access functions for a multicast service need to be supported by the ALA User. The ALA provider is not required to support access control for a multicast AUC. As an option, the ALA provider may offer configuring multicast access lists per end user, allowing ALA users to further secure their multicast content (complementarily to their application layer encryption)

The definitions of the multicast ALA user connection and the IGMP snooping behaviour are provided in the Ethernet ALA service definition [3].

⁹ Since this policing will apply indiscriminately to all frames within the multicast AUC, and not just to frames from the newly added channel that resulted in the limit being broken, this will have a detrimental effect on all multicast traffic in the AUC and will impact all channels in the AUC.

10 ALA CPE Management

The management of ALA provider equipment is outside the scope of ALA since this is internal to the ALA provider's network. This technology is likely to be access network specific, for example GPON networks may use OMCI between the OLT and ONT.

The ALA user may require the ability to manage their equipment at the customer premises, either for operational reasons (for example to monitor access network performance where they have deployed in a wires only model) or as part of a value added service offering. In this case the ALA user management is transported over ALA using a point to point AUC, and may use any number of technologies, for example the Broadband Forums TR-069 protocol or alternatively an SNMP solution. Figure 10.1 shows a TR-069 solution for both x.DSL and PON deployments.

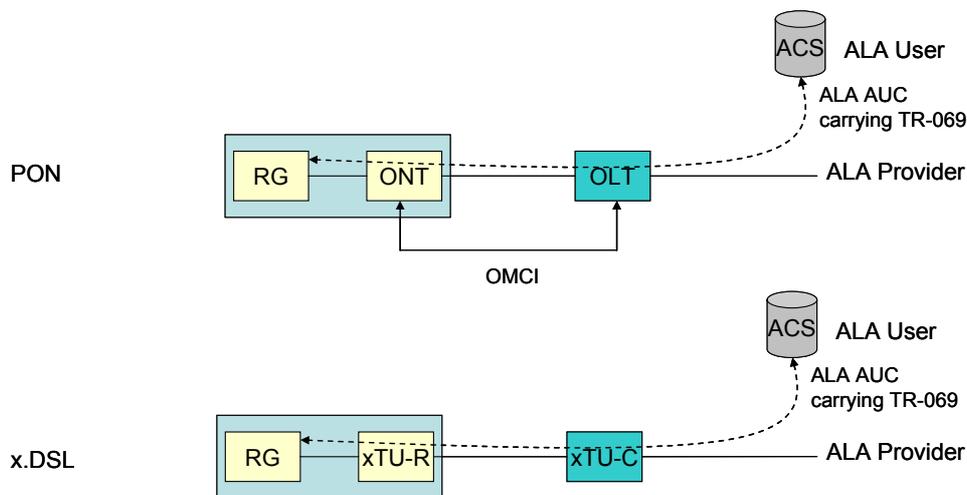


Figure 10.1 ALA CPE management

The ALA user may use a dedicated point to point AUC for carrying management traffic or they may use part of the bandwidth of a point to point AUC that is also carrying end user traffic.

Annex A (informative): ALA Quality of Service implementation options

This Annex provides additional information about the ALA QoS architecture and it provides some examples of QoS implementations that might be used for an ALA deployment.

A.1 QoS solution components

A.1.1 Jitter Control and Queue Allocation

ALA provides support for 4 different Classes of Service. In order to deliver the service attributes of the various Classes, these need to be queued separately at any congestion points. However ALA also supports multiple ALA users per Access line, so the question arises as to whether each ALA User of a given Class needs their own, dedicated queue for that class. This could give rise to a large number of queues in various nodes.

Although having several ALA users on the same access line each of whom used all 4 service classes is perhaps unlikely for commercial, bandwidth and network complexity reasons, there is clearly a case for having more than one Class A user on a line; for example to deliver an additional voice service, and maybe Class B, for a second TV service provider. There are also some scenarios in which multiple Class C users makes sense. This section considers the opportunities for queue sharing (if any) in the case where there are 2 ALA users on a line both of whom use all 4 classes. The same general arguments can however be applied to more than 2 ALA users, and a subset of Classes used by any given ALA user.

A.1.2 Class A

The main purpose of Class A is to support applications which have the most stringent delay and jitter requirements.

If an ALA provider wishes to have deterministic control of the relative Class A performance for multiple ALA users on an interface, he could implement round robin scheduling within Class A with a separate queue for each ALA user and serving the set as a whole to exhaustion before servicing any of the lower priority queues. This approach provides a number of benefits to the ALA provider.

- It allows them to define the traffic management for each ALA users class A service independently of the others.
- The ALA providers product set for Class A services can be refined and expanded without impacting already deployed services.

Note the ALA provider must still account for the performance of all Class A services on a given interface and it is still the case that adding additional Class A services can potentially impact the delay and jitter experienced by existing services on the line.

A single Class A queue for all ALA users is sufficient if the bandwidth sum of all the Class A traffic is much lower than the interface rate. However this would need to remain true as Class A product bandwidths and/or the number of Class A users increased. If a shared queue is used, each Class A user will still need to be individually policed.

A.1.3 Class B

Class B traffic is less jitter sensitive than Class A, and therefore relying on the policers and maximum burst size to constrain the jitter might be acceptable, so queue sharing is more feasible. In order to facilitate line-sharing, both unicast and multicast traffic needs to be policed per ALA user per line; if this is not done then injection of an excessively high bitrate multicast stream e.g. at an access node, could disrupt other ALA users video traffic, and/or completely starve the lower priority data traffic. Therefore there needs to be an aggregate (unicast+multicast) policer per ALA user, after the multicast replication point. This traffic needs to be separated from any other ALA users Class B traffic and therefore may need a separate queue.

A.1.4 Classes C and D

In the upstream direction class C and D cannot share a queue for a given ALA user because it is the inter-queue weighting that provides the controlled sharing during congestion. Since individual ClassC/D users need to be able to specify their own individual weightings (between C and D), these too need to be in separate queues, i.e. a C and a D queue per ALA user (of those Classes) per line. This also allows the defined sharing between different ClassC/D users on the same line to be controlled. In fact both the C/D sharing for a given user, and the defined sharing between ALA users can be accommodated in the single weight applied to each queue.

In the downstream directions class C and D traffic can share the same queue because the ALA user has alternative mechanisms to control the weightings of class C and class D traffic. However there may be some benefits of providing separate queues for class C and D traffic in the downstream direction as a value added service to ALA users.

A.2 An example implementation

This section describes a QoS implementation that could be used to deliver an ALA service for a VDSL architecture as shown in Figure A.1.

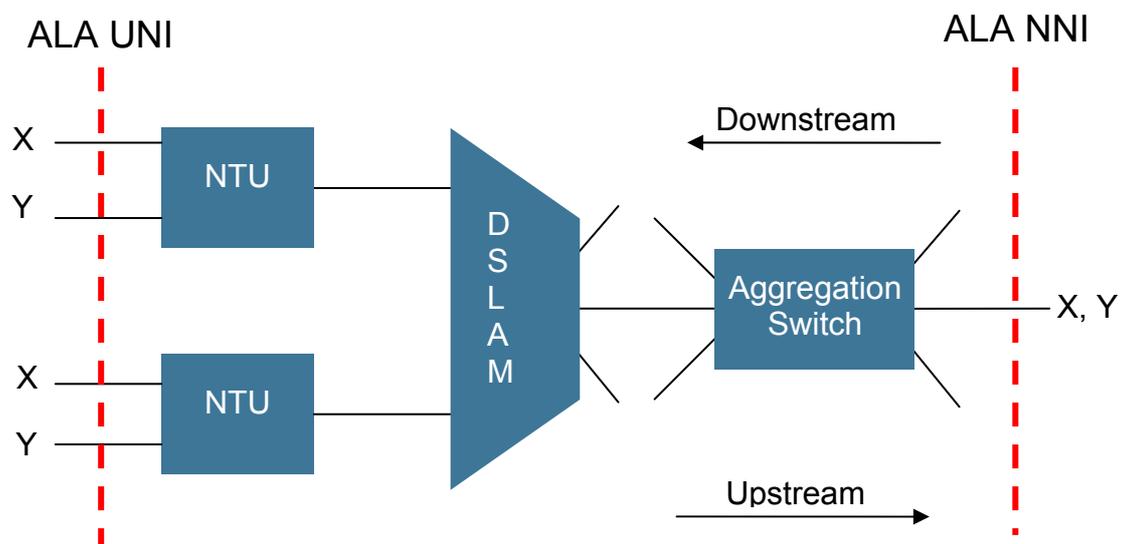


Figure A.1: Reference Network Architecture

In this implementation, the VDSL2 DSLAM in a cabinet is backhauled using a GE link to an aggregation switch at the CO that supports the ALA NNI. There is a VDSL2 NTU at the customer premises which supports two ALA users (X, Y) with an ALA UNI on each of two physical Ethernet ports. The aggregation switch supports multiple cabinets and also hand-off to multiple NNIs.

Upstream traffic for ALA Users X and Y is backhauled over the same physical interface on the Aggregation Switch.

The critical contention ‘point’ in this architecture is the VDSL link between the DSLAM and the NTU. Scheduling across this link is managed by the DSLAM in the downstream direction and the NTU in the upstream direction.

Note: At a wires-only UNI, upstream scheduling at the VDSL2 interface becomes the responsibility of the ALA User CPE.

Figure A.2 shows a scheduling implementation at the NTU VDSL2 interface largely based on the above considerations, i.e. 2, 4-Class ALA users. However as the primary reason for not sharing Class B queues is to control multicast injection, in this example (upstream) Class B traffic can share a queue.

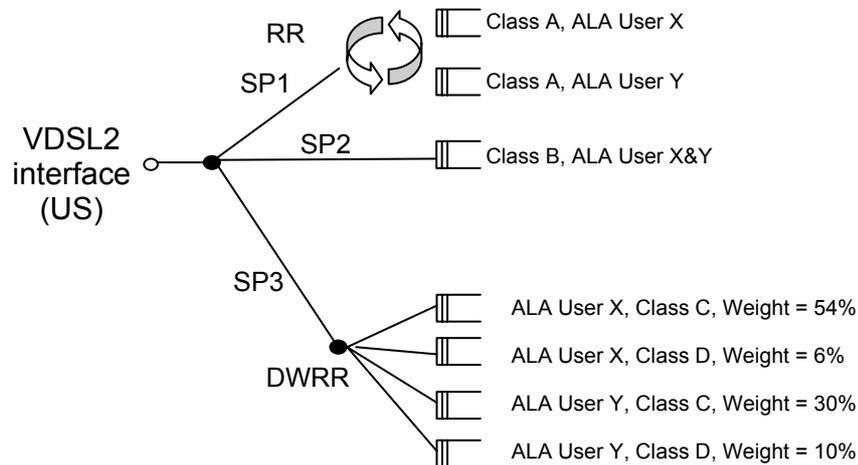


Figure A.2: VDSL interface upstream scheduling

On the NTU, ingress class A and class B traffic is policed according to the ingress bandwidth profile for the AUC group corresponding to each ALA user. This allows traffic from ALA user X and ALA user Y to share the Class B upstream queue, but the Class A traffic for each user is queued separately so the jitter is constrained as described above. Class A traffic is scheduled with strict priority over Class B traffic. Class B traffic is scheduled with strict priority over Class C and D traffic.

Class C and Class D are scheduled using a weighted round robin algorithm. The weighting between Class C and Class D traffic is combined with any weighting desired between the traffic of different ALA users. In this case the bandwidth is split 60:40 between ALA users (as they are taking different products) and the split between Class C and Class D is 90:10 for User X, and 75:25 for User Y. The Class C and Class D queues use WRED to ensure that yellow frames are discarded before green frames.

Figure A.3 shows a possible scheduling implementation at the DSLAM VDSL2 interface.

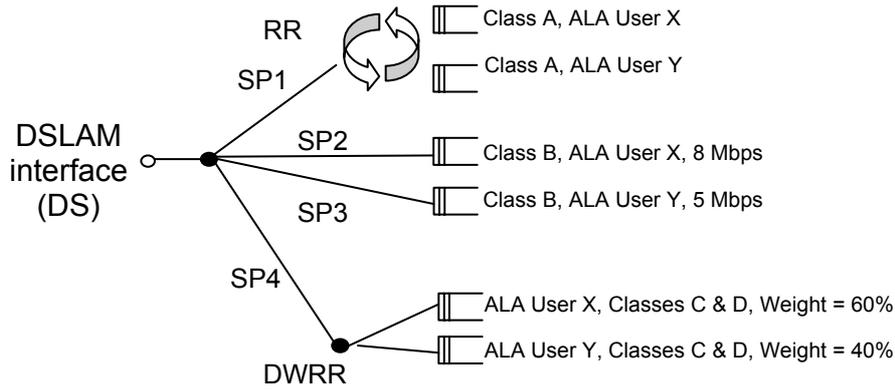


Figure A.3: VDSL interface downstream scheduling – 1

Class A traffic is policed according to the egress bandwidth profile of each 2 user before being scheduled using the 2 top priority queues. Both ALA users use Class B for both video on demand in a Point-to-Point AUC and broadcast video in a multicast AUC. As multicast traffic is replicated in the DSLAM, these AUCs share an AUC Group that is configured with a CIR for class B traffic. This is implemented using a per-ALA user queue for class B traffic that is configured with an 8 Mbps shaping rate for user X, and a 5 Mbps rate for user Y.

Figure A.4 shows a variant of the above in which the DS Class A traffic is in a shared queue. The reason for this is that the amount of jitter introduced by a given packet size is proportional to the PHY line rate. Downstream VDSL rates, are typically significantly higher than the upstream. For rates of 10s of Mbps, there is much less need for packet interleaving to constrain jitter. However this will depend on the product rates and maximum burst size.

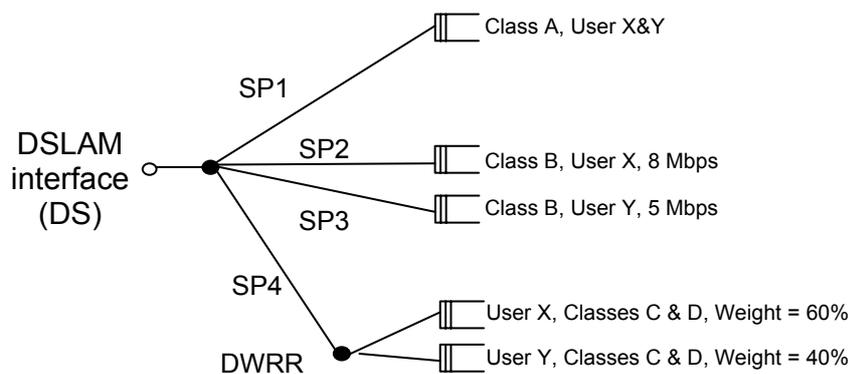


Figure A.4: VDSL interface downstream scheduling - 2

Figure A.5 shows a possible scheduling implementation at the NNI on the Aggregation Switch.

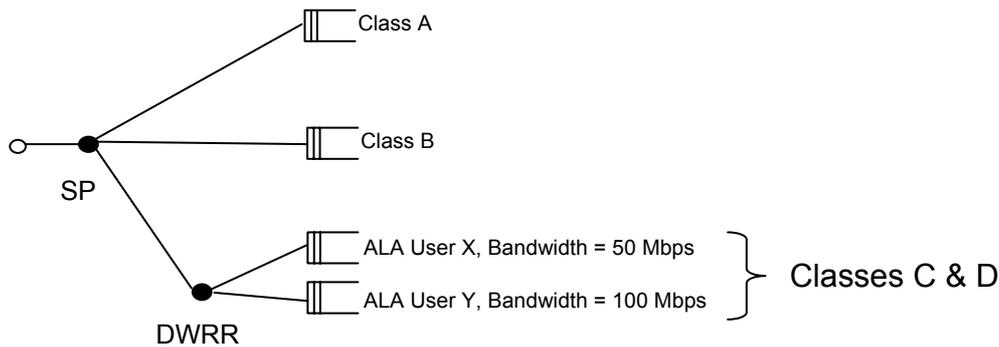


Figure A.5: NNI interface upstream scheduling

The ALA provider shapes the AUCs into per-ALA User S-VLANs at the NNI. An AUC Group at the NNI is defined for each ALA user, and all of an ALA user's AUCs are mapped to the same Group. The aggregate rate for each S-VLAN is the sum of the rates configured for each of the classes.

The Class A and B traffic for each ALA user is policed according to the egress bandwidth profile before it is queued. Strict Priority scheduling is used to schedule all Class A traffic before Class B traffic, which is scheduled before Class C & D traffic. The Aggregation Switch queues Class C and D traffic together, and uses DWRR to schedule different rates for each ALA user depending on the backhaul S-VLAN capacity. The Class C and D queues use WRED to ensure that yellow frames are discarded before green frames.

A.3 ALA Queues and Scalability

The example described in section A.2 results in six queues on a DSLAM VDSL interface and seven at the VDSL modem for 2 ALA users who both use all 4 classes, this number of queues is likely to be the maximum required in any realistic deployment scenario. While 2 ALA users, each using 4 Classes per VDSL line may be unlikely, it is important not to preclude it to avoid 'land-grab'. Supporting eight queues per line would cope with the possibly more likely 'worst-case' scenario of one, 4 Class ALA user, one, 2 Class ALA user (e.g. for a work at home business voice and data connection), and two single Class ALA users (e.g. for a utility and mobile data offload). Note however that this means there needs to be some flexibility with regard to the types of queue for a given number of queues.

A.4 Queuing Implementations for GPON and Point to Point Ethernet

Queuing is used to manage contention at points in the network that can become congested. The above examples are deliberately put in the context of VDSL which can suffer such congestion.

In the case of the NNI it is necessary to define the behaviour so that it is possible to support contended handover.

Where there are no congestion points, then it may well not be necessary to implement the above queue structures and scheduling disciplines in order to meet the SLAs. In particular GPON and point to point Gigabit Ethernet are so fast that much simpler queuing may suffice. A GPON can support multiple ALA users by providing them with separate T-CONTs, but within each T-CONT the queuing can be very simple, and no new capabilities are envisaged to be needed.

History

Document history		
V1.1.1	23 December 2010	Approved version