Broadbanding Australia

NBN Co Fibre Access Service

PRODUCT TECHNICAL SPECIFICATION

VERSION 2.0

DECEMBER 2010

NBNCo Limited

Disclaimer

This document is provided for information purposes only. The recipient must not use this document other than with the consent of NBN Co and must make their own inquiries as to the currency, accuracy and completeness of the information contained herein. NBN Co's position on the subject matter of this document may also be impacted by legislative and regulatory developments in respect of the National Broadband Network.

Environment

NBN Co asks that you consider the environment before printing this document.

Contents

1 Scope and Purpose

1.1 Document Purpose

This document describes the functional and high-level operational aspects of NBN Co's Fibre Access Service (NFAS).

It is intended for a technical audience, who are responsible for integrating NBN Co's NFAS product into their own service delivery architecture.

> This Product Technical Specification represents the culmination of extensive industry consultation, including NBN Co sessions with the Communications Alliance, and a number of technical deep dives with Access Seekers.
>
> The contents of this document represent NBN Co's current position on the subject matter, and should not be relied upon as representing NBN Co's final position, except where stated otherwise.
>
> The views expressed by NBN Co in this document may change.

NBN Co welcomes feedback on this document.

Please provide feedback via email to: feedback@nbnco.com.au with subject "NFAS Product Technical Specification", by the 30th January 2011.

1.2 Scope

This document should be read in conjunction with the following NBN Co publications:

- Product Overview - Fibre Access Services
- Consultation Paper – Access Seeker Accreditation (Fibre Network)

> Note that not all features and product capabilities described within this document will be available in the first release of the NFAS product. An effort has been made to identify capabilities which will be unavailable. However it should not be implied that a feature will be available at first release, merely because there is no indication otherwise.
>
> Refer to the Product Overview document for details of the timing and features of the first and subsequent releases of NFAS.

2 Introduction

NFAS is a wholesale-only Ethernet access product that delivers Layer 2 bit-stream services across a range of access technologies, including:

- FTTP/GPON
- Point-to-Point Ethernet

The NFAS product implements the IEEE802.1ad Provider Bridges VLAN addressing scheme for the operation of point-to-point and multipoint Ethernet Virtual Circuits using Passive Optical Network (PON) and Point-to-Point Ethernet access technologies, suitable for residential, business and infrastructure applications.

NFAS services are constructed from a number of modular product components, each with a different set of attributes and rules. These components are structured to ensure a high degree of configurability and flexibility, to meet the needs of current and future markets. This document details these components, as well as the processes for deploying NFAS services.

In a number of places in this document, NBN Co has provided details of the speeds at which the NBN, components of the NBN or particular wholesale product offerings are capable of operating. However, it is important for acquirers of NFAS to recognise that the speeds actually achieved by an End-User will depend on a number of factors, including the terms of the retail broadband plan, the End-User's chosen hardware, their in-premises connection and backhaul transmission provided by the Access Seeker or third parties.

3 Technical Overview

NFAS services are constructed using a number of configurable, high-level product components:

- User Network Interface (UNI)
- Access Virtual Circuit (AVC)
- Connectivity Virtual Circuit (CVC)
- Network-Network Interface (NNI)

This document provides the technical details of each of these product components, as well as the rules governing their deployment. These product components are depicted in Figure 1.

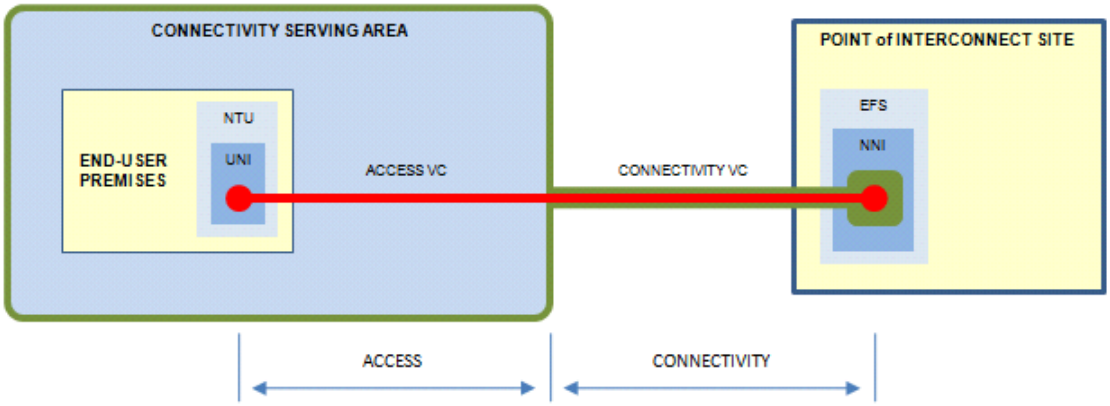

Figure 1 NFAS Product components

The Access VC and UNI product components are specific to an individual End-User Premises, and determine the "Access" portion of the NFAS service. An Access Seeker will dimension these components according to the application requirements of each individual End-User service.

The NNI and Connectivity VC product components together provide an aggregate capacity to a particular Connectivity Serving Area, and are depicted as representing the "Connectivity" portion of the NFAS service. An Access Seeker will dimension these components according to the aggregate End-User service capacity required of that Connectivity Serving Area.

Together, the NNI, CVC, UNI and AVC present a range of attributes that may be configured by Access Seekers, for the delivery of end-to-end services at the Access Seeker's designed service level. Each of these attributes falls into one of the following two categories:

- Configuration Attributes

  Configuration attributes determine how the Access Seeker will interface their network equipment to the NBN Co fibre network. These attributes are defined during the on-boarding phase, and are designed around the Access Seeker's technical product requirements. Examples of NFAS Configuration Attributes are:

  - UNI-D Mode (e.g. tagged., Default-Mapped, etc)
  - UNI-V CODEC Selection (G.711 a-law, G.729 a/b)

An Access Seeker's selection of Configuration Attributes will be validated during the interoperability testing phase, and form the basis of a Configuration Template. Access Seekers are free to define a number of Configuration Templates, noting that any modifications to these profiles must be performed in consultation with NBN Co.

- Service Attributes

  NFAS Service Attributes are those parameters that may be selected at time of service ordering, tailored to the needs of an End-User's service. Examples of NFAS Service Attributes are:

  - Bandwidth Profile
  - Multicast settings
  - Voice features (UNI-V)

  Access Seekers are free to modify Service Attributes on a per-service level, through the standard ordering process. Note that there may be technical limitations on the availability and control of Service Attributes, depending on the Configuration Template chosen.

Details of Configuration and Service Attributes are provided in Sections 9 and 10.

3.1 Connectivity Serving Area

A Connectivity Serving Area (CSA) defines the geographical region that an Access Seeker may address through a single Connectivity VC.

A CSA may cover up to 40,000[1] End-User Premises, enabling an Access Seeker to efficiently deliver services to a large number of End Users.

Access Seekers are free to determine, through the Service Qualification process, which CSA a given End-User Premises belongs to, and what service capabilities may be delivered to that End-User. A feasibility process will also assist the Access Seeker to determine if they have connectivity resources in place to offer services to that location.

Details of each Connectivity Serving Area will be published as part of the NFAS rollout.

3.2 Service Demarcation

The Service and Network Boundary Points for NFAS services are defined at the User Network Interface (UNI) and Network-Network Interface (NNI).

These interfaces define the physical and logical hand-off points for NFAS services. Any infrastructure or service responsibilities beyond these boundaries are the sole responsibility of the Access Seeker.

Note that where a UNI residing on an NTU is cabled to an external physical connector (e.g. RJ-45 connector) located outside of the NTU, the Service and Network Boundary Point will remain at the UNI interface located at the NTU.

[1] The number of End-User Premises served per CSA is subject to final network design

3.3 NFAS Product components

The product components depicted in Figure 1 are described individually:

3.3.1 User Network Interface (UNI)

The UNI is a data (Ethernet) or telephony (analog POTS) physical interface, located at the End-User Premises, and is required to deliver one or more data/telephony services. This interface is housed on an NTU which is capable of supporting multiple, independent UNI ports. Details on the number and types of UNI supported by the different NTUs are provided in Section 12.2.

Refer to Sections 8.1 for details on the various UNI ports supported by NFAS.

3.3.2 Access VC

An Access VC is a logical Ethernet Virtual Circuit dimensioned to the needs of a specific End-User service, operating across the NFAS network between an NNI and single UNI. An Access Seeker will dimension this circuit according to the traffic capacity and performance characteristics required to support an individual End-User's service requirements.

Refer to Section 8.2 for details on the Access VC.

3.3.3 Connectivity VC

A Connectivity VC represents an aggregate "trunk" capacity purchased at the NNI, dimensioned by the Access Seeker to support a number of Access VCs to a particular Connectivity Serving Area. The amount of capacity allocated to a Connectivity VC is aligned to the aggregate needs of its Access VCs, and allows an Access Seeker to manage contention and End-User experience.

The number of Connectivity VCs to a particular Connectivity Serving Area that an Access Seeker will need to purchase is a technical matter, and it will depend on the number and type of End-User services being delivered and the Access Seeker's desired capacity to service those End-Users.

Refer to Section 8.3 for details on the Connectivity VC.

3.3.4 Network-Network Interface (NNI)

The NNI is a physical, aggregated Ethernet interface, directly accessed by the Access Seeker within the Point of Interconnect (PoI), and used to interface NFAS services to an Access Seeker's backhaul or IP core. Where multiple Connectivity Servings Areas are served by a single POI, a single NNI may be used to deliver traffic to one or more Connectivity Serving Areas.

Refer to Section 8.4 for details on the NNI.

3.4 Service Construction

NFAS product components cater for a wide variety of Access Seeker deployment scenarios, providing flexibility and control of capacity, and feature and availability options for premium applications. This model also allows for Access Seekers to re-use and leverage components (such as NNI) for multiple End-Users, allowing the ability to cost-effectively and seamlessly grow a service footprint and applications with minimal impact to existing services.

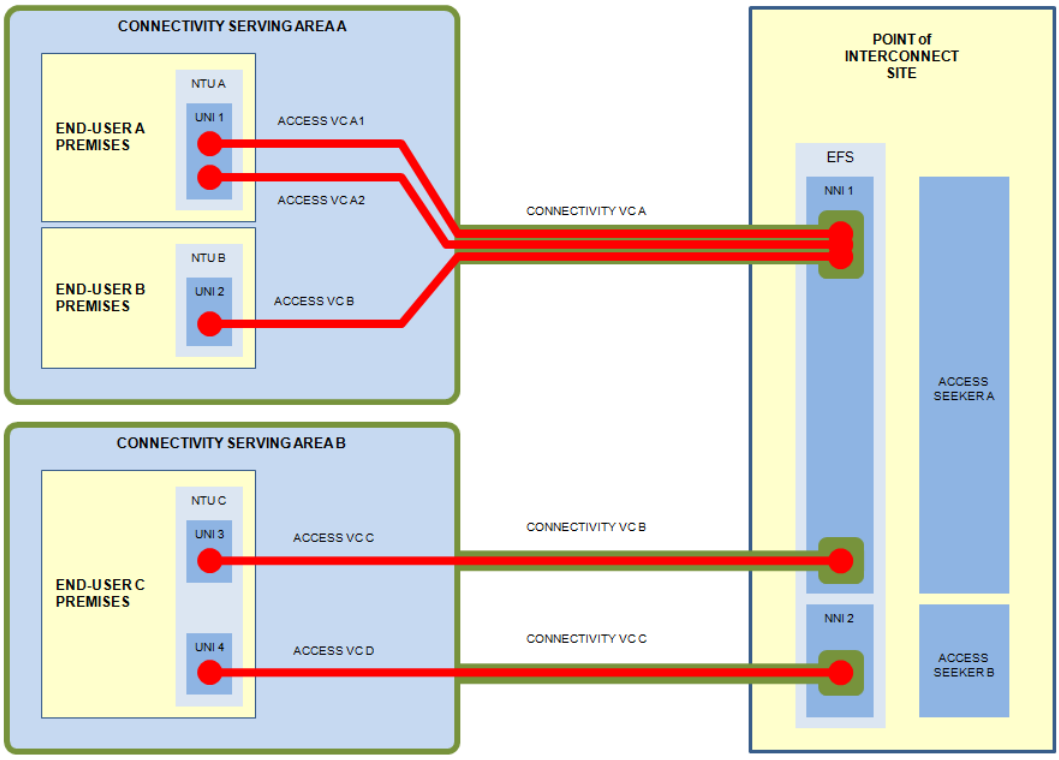

Figure 2 Application of NFAS Technical Constructs

Figure 2 depicts a number of example scenarios whereby the NFAS product components are used to construct end-to-end services. Note the following:

- End-Users A and B are located within the same Connectivity Serving Area, and accessed via a single Connectivity VC (Connectivity VC A).
- End-User A has two Access VCs (Access VC A1 and A2) delivered through a single UNI (UNI 1). Access VC A1 (for example) may be used for an IP-based telephony service, and the other (Access VC A2) used for internet access. These are both provided off the same NNI (NNI 1), by the same Access Seeker (Access Seeker A).
- End-User C is in a different Connectivity Serving Area, and is served using different Connectivity VCs from End-Users A and B.
- End-User C has two services, each provided from a different Access Seeker, delivered on two different UNI (UNI 3 and UNI 4). This may be an IP-based telephony service and internet service, however unlike End-User A, End-User C has chosen to receive these services from two different Access Seekers.

4 Supported Service Types

This section provides a brief overview of the wide variety of data services able to be simultaneously delivered across the NFAS network.

4.1 Unicast Data Services

NFAS supports the flexible delivery of unicast data services using logical circuits that may be used for a variety of data applications, including IP-based telephony services, internet access and video conferencing. NFAS unicast data services support IPv4 as well as IPv6 protocols, allowing Access Seekers a smooth migration path to future IP-based network architectures.

These unicast services provide physical point-to-multipoint (aggregated) connectivity between one or more UNI located at the End-User Premises, and an Access Seeker's centrally-aggregated NNI.

Unicast services are configured with the following general sets of attributes:

- Specification of end-points (i.e. End-User Premises and Point of Interconnect)
- Allocation of bandwidth (traffic capacity)
- Additional features (e.g. SLA options)

Within these broad categories are a range of additional attributes and options, including capacity allocations, Class of Service handling, and reporting.

4.2 Multicast Data Services

NBNCo's Fibre Access Service facilitates the efficient deployment of multi-media content distribution services by Access Seekers through offering layer 2 multicast capabilities.

This capability supports the simultaneous, secure and dynamic delivery of multiple video streams at a variety of bit-rates, from multiple Access Seekers to select End-Users. It also operates in an environment where other simultaneous unicast services (such as telephony and data) are being delivered on the same End-User UNI, though on a different Access VC and Connectivity VC.

NFAS Multicast capabilities seamlessly scale to support a variety of video qualities and deployment options, enabling Access Seekers to offer a range of standard and high-definition video formats, and a variety of channel packages.

4.3 Business Data Services

NFAS provides a range of business-oriented capabilities that are targeted primarily at enterprise applications.

A summary of these NFAS features, as applicable to the business market, is shown in Table 1.

Table 1 Business Feature Summary

Business Feature	Description
Symmetrical Bandwidth*	Symmetrical bandwidth options up to 1Gbps
Optical UNI Interfaces*	UNI options for networking equipment located beyond 100m from the NFAS NTU.
Class of Service	Four classes of service, aligned to IP QoS application behaviour
NTU Power Options*	DC power option for integration into server racks, or DC-powered environments
Capacity Contention	The ability to control the contention ratio for business applications
CE-VLAN Transparency*	The ability to provide “Q-in-Q” for Transparent LAN (TLS) business services
Service OAM*	Increased service visibility and diagnostics to support higher quality SLAs
Direct Fibre Option*	Bandwidth scalability which is expected to allow Access Seekers to adapt to future technological advances and increases in End-User demand for bandwidth
Increased Frame Size	Support for various layer 3+ protocol encapsulations through a 2000 byte MTU
IPv6 Transparency	Seamless migration path to IPv6 addressing
Security	Inherent security through the use of dedicated point-to-point, logical circuits
CPE interface Options	Wide range of UNI-D interfacing options for seamless migration off legacy CPE
Enhanced Reporting*	Greater visibility into service performance and operation
Enhanced SLA Options*	Wide range of assurance options to support Access Seeker’s SLAs
Higher Availability*	Options for increasing access resiliency for critical services

These features may be selectively used by Access Seekers to tailor NFAS services to specific segments within the business market.

* Certain business features are not supported within NFAS first release.

4.4 Telephony Services

NFAS services may be used for the provision of IP-based telephony services to End-Users via two means:

- An integrated Analogue Telephony Adaptor (ATA) port, with integrated SIP capabilities for legacy telephony applications (UNI-V)
- Access to external, Access Seeker-supplied ATA devices using a UNI-D port

An Access Seeker who wishes to use NFAS services for the delivery of IP-based telephony services is expected to provide and manage their own IP-based telephony network capabilities that interface to, and operate across, the NFAS network.

All IP-based protocols and functions that the Access Seeker utilises to implement IP-based telephony services will pass transparently through the NNI, AVC, CVC and UNI-D NFAS product components. Where utilised, the UNI-V will terminate all IP-based telephony protocols and functions.

NFAS supports the provision of voice-grade, IP-based telephony services through the use of specific traffic handling mechanisms that are tailored toward deterministic performance for real-time, conversational applications. The TC_1 traffic class is designed to accommodate the needs of IP-based telephony applications.

Capacity within this traffic class is available to the Access Seeker via the UNI-D or UNI-V interfaces, ensuring a consistent telephony service experience regardless of the interface used.

4.4.1 Legacy Telephony Applications

Using the UNI-V, an Access Seeker may access the in-built Analogue Telephony Adaptor (ATA) port, with integrated SIP capabilities for legacy telephony applications. A range of configuration options enable an Access Seeker to migrate an existing telephony service, with minimal impact to in-building wiring or CPE.

An Access Seeker must interface their core IP-based telephony network with the IP-based telephony functions provided by the internal ATA of the UNI-V port. This will involve a degree of integration testing between the Access Seeker and NBN Co prior to service deployment.

IP-based telephony services deployed using the UNI-V are automatically provisioned with a specific TC_1 capacity allocation.

4.4.2 External ATA Device Support

An Access Seeker may choose to deliver IP-based telephony services to an End-User Premises using a dedicated, external ATA device beyond the NFAS Service Boundary Point. The supply and operation of this device is the responsibility of the Access Seeker.

Such devices are readily available for consumer applications today, and will appear to the NFAS service as a regular data device, connected to a UNI-D port.

The Access Seeker may choose to operate the Access VC in a manner that recognises the relative priority of telephony traffic above other applications sharing the same Access VC, or may optionally dedicate an Access VC wholly to the telephony device.

Under this deployment scenario, the NFAS service is agnostic[2] to the IP-based telephony protocols and data that the Access Seeker utilises for the delivery of IP-based telephony services to an End-User.

When delivering IP-based telephony services using an external ATA, the Access Seeker is able to utilise capacity from any of the NFAS traffic classes (not restricted to TC_1).

[2] Note that specific CoS handling may be configured for voice packets (requires appropriate IEEE802.1p marking)

5 Service Addressing

This section details the options for NFAS service addressing, including IEEE802.1ad S-TAG/C-TAG structure, the allocation of S/C-VID values, and the addressing options available at the UNI-D. It describes the structure of the service frame with regard to fields used for individual service identification.

NFAS supports a common NNI addressing scheme for Connectivity VCs, using an IEEE802.1ad S-TAG to identify individual services.

NFAS supports four different NNI service addressing modes for Access VCs, capable of being selected at a CVC-level. These service addressing modes define how each individual AVC service within a CVC will be addressed by an Access Seeker through the NNI.

5.1 VLAN Tag Structure

When required for CVC/AVC service addressing, each S-TAG and C-TAG is required to contain the following fields[3]:

- S/C-TPID – Tag Protocol Identifier, used to identify the tag type
- S/C-VID – VLAN Identifier, used for service identification
- S/C-PCP – Priority Code Point Identifier, used for priority marking

This is depicted in Figure 3.

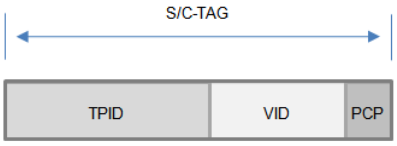

Figure 3 S/C-TAG Structure (4 Bytes)

Where appropriate, these fields will be validated for all service frames ingress to the NFAS network.

5.2 Connectivity VC Addressing

Connectivity VCs are identified at the NNI using an outer IEEE802.1ad S-TAG, contained within each service frame. Each Connectivity VC within an NNI may be addressed and operated independently, allowing adjacent Connectivity VCs to be configured differently.

It is the responsibility of the Access Seeker to ensure that each supplied S-TAG VID field is as per the agreed service configuration. The NFAS service will discard any service frames received at the NNI with an S-VID that does not map to an agreed upon identifier for an active Connectivity VC service.

At egress to the NNI, the NFAS service will insert the S-TAG with the agreed S-VID for identification of the Connectivity VC to the Access Seeker.

[3] Refer IEEE802.1ad for explanation of S/C- TAG fields

Within a Connectivity VC, a number of Access VCs may be present. The mechanism used to address these individual Access VCs depends upon the service being operated through the Connectivity VC.

The following service addressing modes are used at the NNI to access individual Access VC services operating through a Connectivity VC.

5.3 Service Addressing Mode A

Service Addressing Mode A uses a two-level VLAN addressing scheme at the NNI, which is compliant with IEEE802.1ad [8] (Provider Bridges) to identify individual Access and Connectivity VC services. This mode is suitable for traditional unicast data services between the NNI and UNI-D ports.

Figure 4 describes the frame structure for service frames presented at ingress to the NNI using Service Addressing Mode A, highlighting the S-TAG and C-TAG provided by the Access Seeker, required to associate the service frame with an individual Connectivity and Access VC.

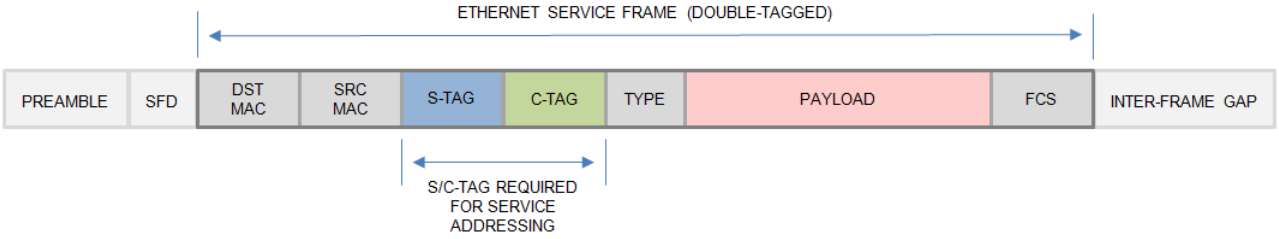

Figure 4 Service Addressing Mode A Frame Format[4]

Services using this addressing mode use the inner IEEE802.1ad C-TAG VID field to address an individual Access VC within a CSA. This C-TAG is visible at the NNI, and may be passed across the UNI boundary, for the identification of Access VCs to any CPE devices at the End-User Premises. The C-VID can be used to address up to 4000 individual Access VCs through a single S-TAG. Note that the same C-VID may appear through different S-TAGs on a given NNI, even where both S-TAGs are directed to the same Connectivity Serving Area. In such cases, the C-VIDs must always address different NTU UNI-D ports.

The S/C-PCP field is used to communicate priority information both across the UNI/NNI boundaries, and within the NFAS network.

Figure 5 describes this addressing mode in operation across the NFAS network, highlighting the scope of the S-TAG and C-TAG.

[4] Refer IEEE802.3 for explanation of service frame fields

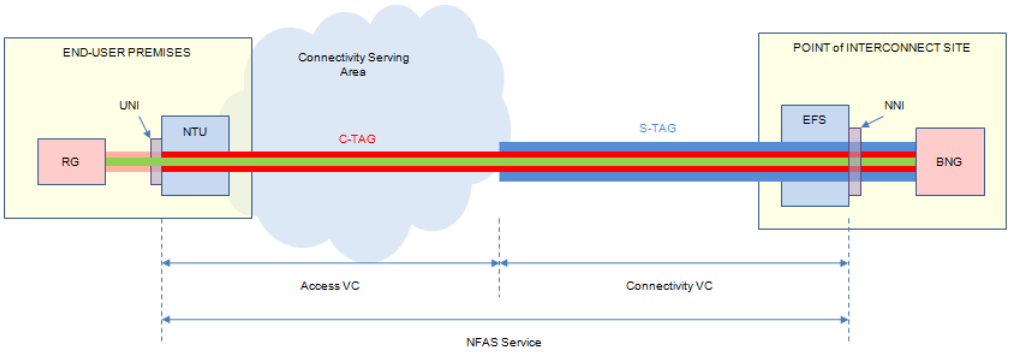

Figure 5 S/C-TAG Structure Applied to an NFAS Service (Service Addressing Mode A)

Service Addressing Mode A requires that traffic flowing in the downstream direction (from the Access Seeker's network into the NNI) must be tagged with the appropriate S/C-VID settings. Traffic flowing in the upstream direction, ingress to the UNI, may utilise one of a range of tagging options (refer Section 8.1.1.3). It is the responsibility of the Access Seeker to ensure that all ingress traffic is compliant with the assigned VID settings for each respective service.

5.4 Service Addressing Mode B

Service Addressing Mode B provides the ability for an Access Seeker to utilise the C-TAG VID for its own purposes, within an End-User's service.

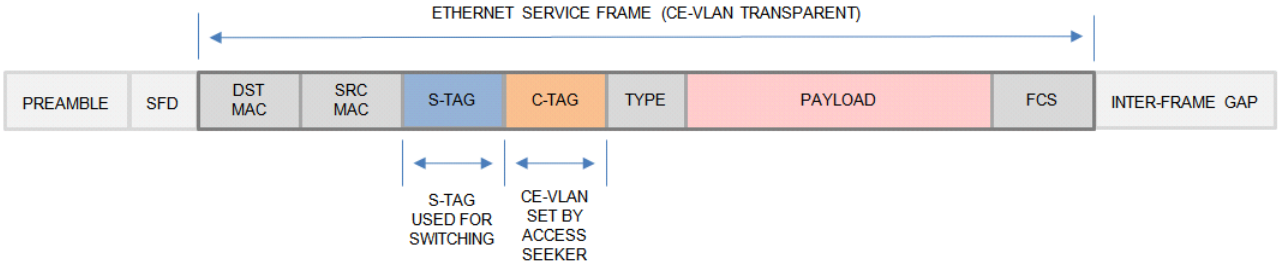

Figure 6 Service Addressing Mode B Frame Format

This capability is shown in Figure 7.

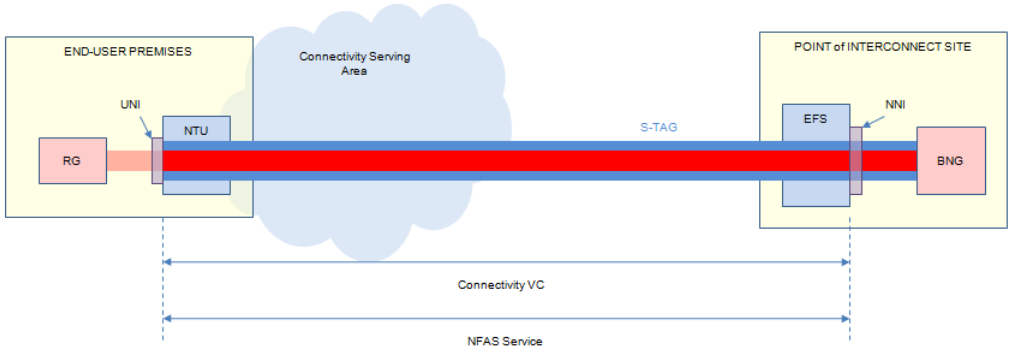

Figure 7 S/C-TAG Structure Applied to an NFAS Service (CE-VLAN Transparent Mode)

Note that under this mode, the S-TAG addresses a Connectivity VC that is delivered to an End-User UNI-D port on a specific NTU.

This allows the inner C-TAG to be used by the Access Seeker for the provision of a CE-VLAN transparent service to the End-User, enabling the following End-User frame formats:

- Tagged
- Priority-Tagged

This mode of operation has the following implications:

- The use of S-TAGs to address individual End-User UNI ports will place additional strain on the Access Seeker's available S-TAG allocation pool. There is a limited number of S-TAG VIDs available per NNI (refer Section 8.4.4.4), and this may become quickly exhausted if this mode is used extensively (requiring the purchase of additional NNI).
- This mode requires that a 1:1 relationship exist between the Connectivity VC and the UNI end-point (note that there is no Access VC).

This service addressing mode is provided as an option for business services that require freedom to communicate VLAN information across a WAN service. This mode is not expected to be used for mass-market applications, and is not offered in first release.

5.5 Service Addressing Mode C

Service Addressing Mode C implements N:1 addressing for IP-based telephony applications using the UNI-V). These services require the frame format as shown in Figure 8 at the NNI:

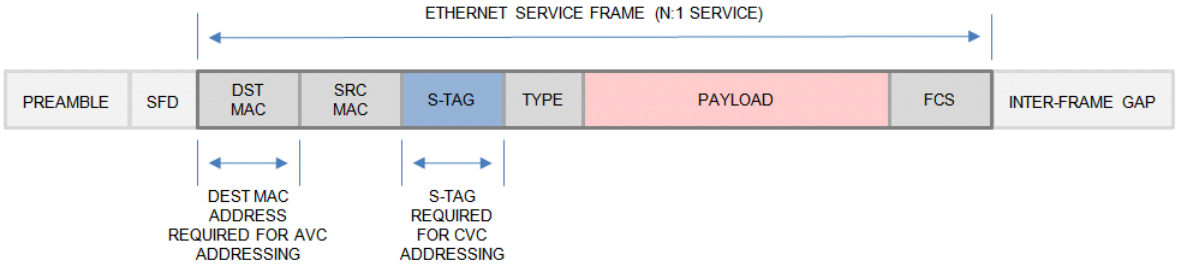

Figure 8 Service Addressing Mode C Service Frame Format

Figure 8 describes the frame structure for service frames presented at ingress to the NNI for this type of service, highlighting the S-TAG provided by the Access Seeker, required to associate the service frame with an individual Connectivity VC, and the Destination MAC field which identifies the individual destination UNI-V.

The example service depicted in Figure 9 demonstrates the use this addressing mode to indicate a Connectivity VC service and Access VC connecting to a UNI-V port, to provide an IP-based telephony service.

Note that under this addressing mode, there are no restrictions imposed by C-TAG VID range limitations on the number of Access VCs that can be addressed through an S-TAG.

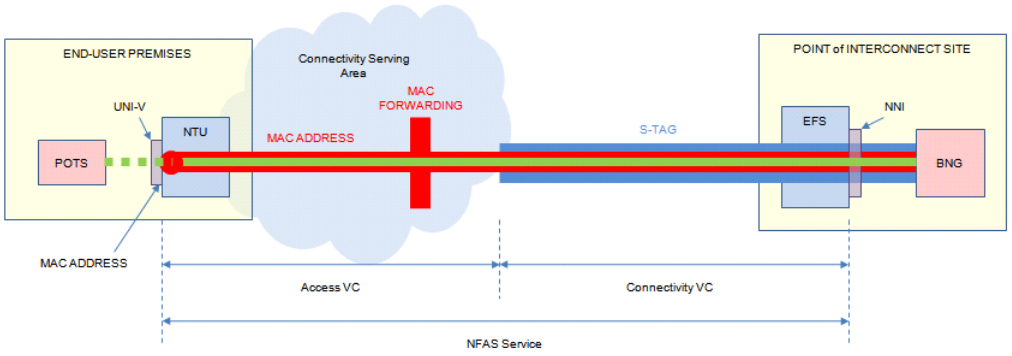

Figure 9 Addressing Mode C Service Demonstrated for UNI-V Interfacing

5.6 Service Addressing Mode D

The example service depicted in Figure 10 demonstrates the use of Service Addressing Mode D to access a Connectivity VC and Access VC connected to a UNI-D port, to provide a multicast service.

For this addressing mode, the NFAS network implements a multicast distribution function that enables connectivity across all CSAs that can be accessed through the NNI.

Note that the NNI frame structure for this addressing mode is similar to that depicted in Figure 8, however the scope of the MAC addressing is increased to multiple CSAs.

Note that for any NFAS services that utilise MAC forwarding, all ingress traffic at the End-User Premises is passed to the CVC. There is no forwarding of service frames directly between UNI-D/V.

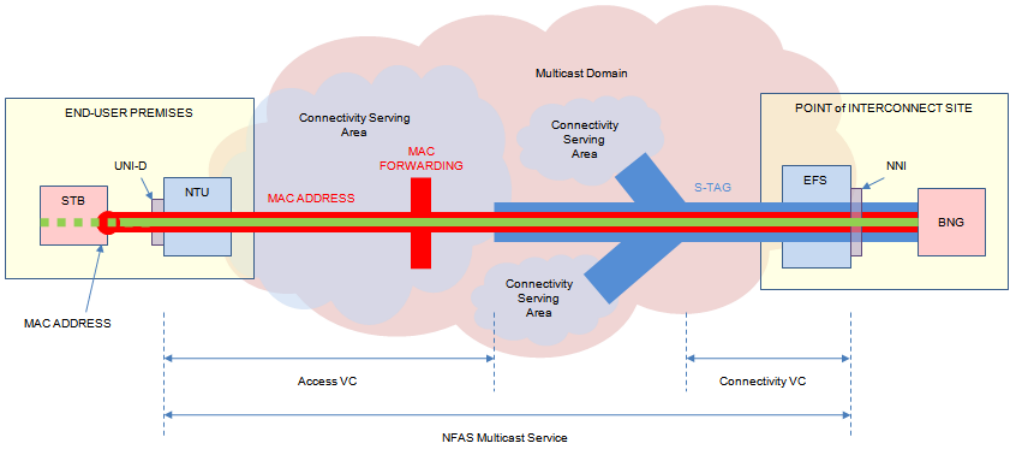

Figure 10 N:1 Service Demonstrated for Multicast Operation

Under this addressing mode, there is no restriction on the number of Access VCs that can be addressed through an S-TAG.

5.7 S/C-VID Allocation

The allocation of S/C-VID values must be co-ordinated between the Access Seeker and NBN Co.

Where required during the ordering process, NBN Co will, by default, allocate each new CVC/AVC an internally-generated S/C-VID. This S/C-VID value will be returned to the Access Seeker, and will be used for accessing the requested service at the NNI.

Access Seekers may optionally elect to nominate the S/C-VID used to address each CVC/AVC service instance through the NNI, for further alignment to their own core network addressing schemes. Note that Access Seekers are encouraged to use NBN Co's S/C-TAG default VID allocations, which will be unique within the Access Seeker's service. This will avoid any potential for S/C-VID mismatch between the Access Seeker and NBN Co.

For service addressing modes at the NNI that rely on MAC addressing for forwarding within the access network, the allocation of a C-VID is not required.

Note that where a C-VID is required at the UNI-D to address an individual Access VC (using a Tagged UNI-D), the value used for each order will be supplied through a Configuration Template, defined by the Access Seeker during the on-boarding phase.

5.7.1 Default S/C-VID Allocation

Figure 11 depicts an example where the VID for both the S-TAG and C-TAG have been allocated internal values (C-VID=A, S-VID=B) by NBN Co.

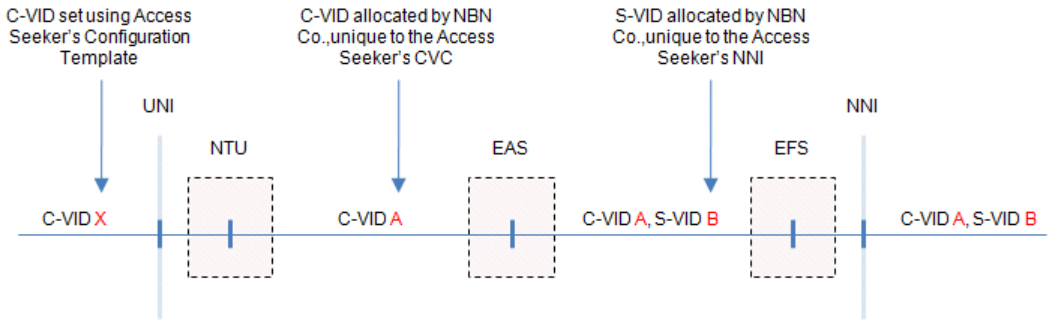

Figure 11 Default VID Mapping (Service Addressing Mode A)

Note that in this example, the C-VID is visible at the UNI, requiring a tagged UNI-D. During the on-boarding process, the Access Seeker is required to nominate the C-VID to be used at the UNI-D to address the Access VC. The C-VID in this case is fixed (C-VID=X), and cannot be changed through the order process.

The default S/C-VID that is visible at the NNI is allocated by NBN Co according to the following:

- S-VID values will be allocated unique to an Access Seeker's NNI.
- C-VID values will be allocated unique to an Access Seeker's Connectivity VC (S-TAG)

Default S/C-VID values will be allocated as the next available, sequential value within the addressing scope of that Access Seeker's services. Where a default value cannot be allocated, due to exhaustion of the address space, the order will be rejected, and the Access Seeker will be required to order a new NNI (for S-VID exhaustion) or Connectivity VC (for C-VID exhaustion). Any S/C-VID exhaustion issues are contained entirely within the Access Seeker's service domain.

Under the default scenario, the Access Seeker must orient their core network addressing scheme to the S/C-VID allocation of the NFAS network.

5.7.2 Access Seeker S/C-VID Allocation

An Access Seeker may request specific S/C-VIDs during the order and on-boarding processes. Note that where the Access Seeker chooses to nominate the S/C-VID settings during the order process, the Access Seeker maintains full responsibility for ensuring that the nominated S/C-VID settings are unique within their addressing scope, and do not overlap with any other active services. NBN Co will reject any service order where the nominated S/C-VID setting clashes with another active service.

For the Access VC, the C-VID may be requested through the order process for presentation at the NNI only. Where required at the UNI, the C-VID for an Access VC will be determined through the Configuration Template (as per the default scenario).

Note that Access Seekers are encouraged to use NBN Co's default per-service S/C-VID allocations, which will be unique within the Access Seeker's CVC and NNI. This will avoid any potential for S/C-VID mismatch between the Access Seeker and NBN Co.

Figure 13 shows an example Access Seeker deployment, using NFAS to deliver End-User services to two Residential Gateways (RG_A and RG_B), from a centralised Broadband Network Gateway (BNG).

Both RG devices use a standard C-VID value (10) to interface to the NFAS service. This value is supplied through the Access Seeker's Configuration Template, and aligned to the standard RG configuration. Each End-User Access VC is mapped to a unique C-VID once inside the NFAS network, allowing each service to be individually addressed. In this example, the C-VID values (C-VID=124 and 125) are supplied by the Access Seeker through the ordering process.

For the CVC, the Access Seeker also chooses S-VID = 1, for addressing through the NNI.

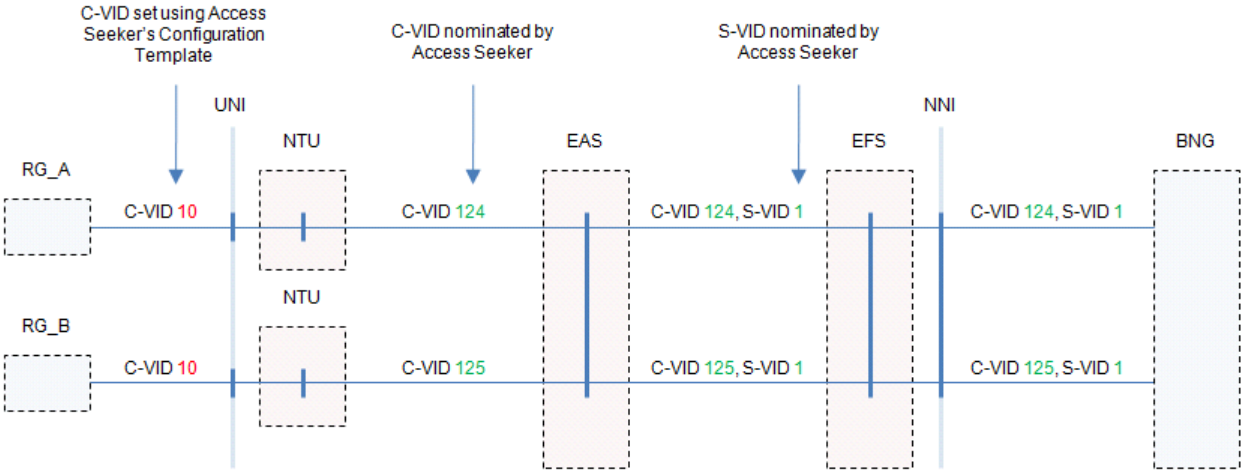

Figure 12 Example of S/C-VID Allocation

5.8 Tag Protocol Identifier (TPID) Formats

Table 2 describes the required TPID values for service frames ingress to the NFAS network. Any received service frames that do not comply with these values will be discarded at ingress.

Table 2 EtherType Requirements

Interface	Mode	S-TPID	C-TPID	Comment
NNI	Addressing Mode A	0x88A8	0x8100	C-TPID value indicated is applicable to inner C-TAG. Note that the inner C-TAG is not examined or validated at the NNI, however must be submitted to the NFAS network as indicated, for switching within the CSA. S-TPID value applicable to outer S-TAG.
	Addressing Mode B		No Restriction	CE-VLAN Transparent services may utilise their own C-TPID settings.
	Addressing Mode C		N/A	Addressing Modes C and D utilise MAC forwarding for the Access VC, and do not require a C-TPID.
	Addressing Mode D		N/A	
UNI-D	Default-Mapped	N/A[5]	N/A[6]	For Default-Mapped and DSCP-Mapped UNI-D, the C-TPID is supplied by the NFAS service at ingress.
	DSCP-Mapped			
	Tagged		0x8100	All ingress service frames received at the UNI-D interface must have the C-TPID shown. Any service frames with C-TPID other than this (defined at time of order) will be discarded at ingress.
	Priority Tagged		0x8100	
	CE-VLAN Transparent		No Restriction	

Any tagged service frames with TPID settings outside of these values will be discarded at ingress.

[5] S-TPID appended by NFAS network and not visible at UNI-D.

[6] Untagged interfaces do not require a C-TAG.

6 Class of Service (CoS)

NBN Co's NFAS product implements four traffic classes that are distinguished in capability and performance, designed to accommodate the widest variety of higher-layer applications. Access Seekers may take advantage of these traffic classes, to provide more tailored performance and effective utilisation of NFAS services.

6.1 NFAS Traffic Classes

The NFAS traffic classes are described in Table 3.

Table 3 NFAS Traffic Classes

Traffic Class	Example Applications	Specification
TC_1	Voice	CIR
TC_2	Interactive streaming and real-time video	CIR
TC_3	Premium data	CIR, PIR
TC_4	Best-effort data	PIR

Access Seekers are free to use these classes to allocate service capacity in a manner that reflects the demands and operation of their end-to-end applications.

These traffic classes are aligned to the application definitions within RFC4594 (refer Appendix C – Class of Service Application), to support an Access Seeker's higher-layer, end-to-end IP Quality of Service policies.

Note that NBN Co may introduce, as part of future products or product upgrades, additional traffic classes or refinements to the traffic class attributes.

6.1.1 TC_1 Description

The TC_1 traffic class is targeted towards real-time, interactive multimedia applications, with the following characteristics:

- Low bit-rate
- Low frame delay, frame delay variation, frame loss
- Highest levels of availability

The attributes of this class are aligned to the characteristics of the DSCP Expedited Forwarding (EF) per-hop behaviour described in RFC4594, section 1.5.3.

TC_1 provides a committed level of premium capacity with no ability to burst above its CIR, suitable for applications that require deterministic performance and are likely to be sensitive to packet loss.

6.1.2 TC_2 Description

The TC_2 traffic class is targeted towards real-time, interactive multimedia conferencing applications, as characterised by the DSCP Assured Forwarding (AF) per-hop behaviour described in RFC4594, section 1.5.2.

TC_2 provides a committed level of capacity with no ability to burst above its CIR.

6.1.3 TC_3 Description

The TC_3 traffic class is targeted towards enhanced data applications. This traffic class complements the TC_4 class, by adding a CIR component.

It is expected that Access Seekers will apply a lower over-booking ratio for TC_3, than for TC_4, as required for data applications that require higher determinism than best-effort delivery.

6.1.4 TC_4 Description

The TC_4 traffic class is targeted towards “best effort” applications, as characterised by the DSCP Default Forwarding (DF) per-hop behaviour, described in RFC4594, section 1.5.1.

Each access seeker will be free to configure its own AVC to CVC overbooking ratios.

6.2 Traffic Class Scheduling

Figure 13 depicts the mechanism used to schedule traffic across each of the four traffic classes.

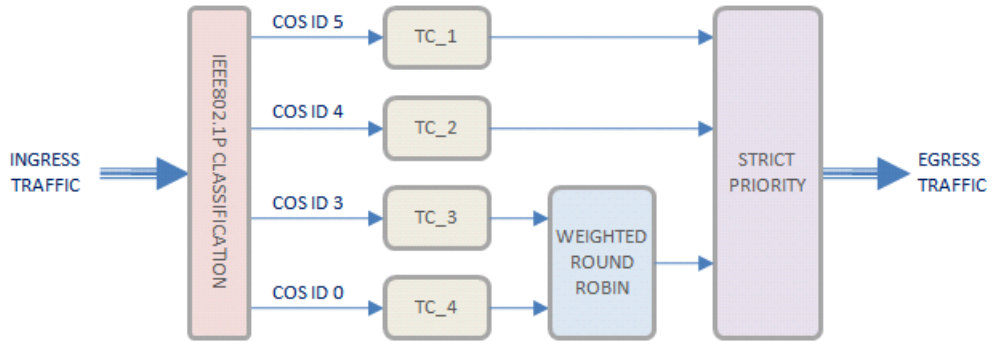

Figure 13 Traffic Class Scheduling

Note that the weightings to be used for the Weighted Round Robin scheduler that services the TC_3 and TC_4 traffic classes are still being considered.

Table 4 Traffic Class Scheduling Description

Traffic Class	Scheduling Mode	Priority	Comment
TC_1	Strict Priority	1	Traffic serviced at the highest priority
TC_2	Strict Priority	2	Traffic serviced at the second highest priority, after TC_1.
TC_3	Weighted Fair Queuing	3	Weighting applied to TC_3 – TBA
TC_4			Weighting applied to TC_4 - TBA

6.3 Bandwidth Profile Parameter Definitions

This section provides clarification of the bandwidth profile parameters used within the NFAS product.

6.3.1 Calculation of Information Rate

All Information Rates are calculated on Access Seeker layer 2 Ethernet service frames, over the series of bytes from the first bit of the Destination MAC Address through the last bit of the Frame Check Sequence. Note that IEEE802.3 physical-layer fields such as the Preamble, Start of Frame Delimiter and Inter-Frame Gap are not included in the Bandwidth Profile.

Note that the effective layer 2 payload rate of any NFAS service will degrade slightly for lowest-sized Ethernet service frames. This is the expected behaviour for Ethernet-Based services whose Bandwidth Profile is Based on the Service Frame definition as per Figure 4. It is the responsibility of the Access Seeker to accommodate any payload rate degradation as a result of layer 2 frame sizes. Note that the effective throughput of an Ethernet service is described in Appendix B – Effective Information Rate.

6.3.2 Committed Information Rate

Committed Information Rate (CIR) defines a level of data throughput for which service frames are delivered according to the performance objectives of the respective traffic class.

6.3.3 Peak Information Rate

Peak Information Rate (PIR) is defined as the maximum data throughput that may be delivered by the service. Note that traffic capacity in excess of the CIR and within the PIR will be carried through the NFAS network without any performance objectives. Traffic that exceeds the PIR will be discarded at ingress to the service.

6.4 Bandwidth Specification Model – Access VC

The Access Seeker is required to select the desired amount of capacity for each Traffic Class required for the Access VC at time of service order. The selectable CIR and PIR capacities are detailed in the Product Overview.

Note that in order to take advantage of the NFAS CoS capabilities, the Access Seeker will be required to have configured the UNI-D as one of:

- Tagged
- Priority-Tagged
- DSCP-Mapped

For Default-Mapped UNI-D, it is not possible to allocate traffic to different traffic classes. Under this configuration, all ingress traffic will be carried within the TC_4 class.

The AVC Bandwidth Profile for NFAS First Release is shown in Table 5.

Table 5 Bandwidth Profile – AVC

Traffic Class	Component	Units	Description
TC_1	CIR	Mbps	CIR requirement for TC_1
TC_2	CIR	Mbps	CIR requirement for TC_2
TC_3	CIR	Mbps	CIR requirement for TC_3
TC_3	PIR	Mbps	PIR requirement for TC_3
TC_4	PIR	Mbps	PIR requirement for TC_4

The allowable settings for each of the Traffic Class capacity allocation parameters are defined in Figure 14.

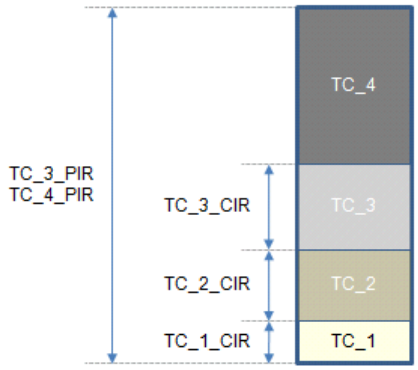

Figure 14 NFAS AVC Bandwidth Specification

Access Seekers are free to simultaneously provision and operate multiple applications, for example:

- Voice-grade IP telephony service
- Streaming, real-time residential and business video applications such as video conferencing, unicast video and multicast IPTV
- Premium data applications that require deterministic performance, such as gaming, business data and Video on Demand
- Data applications that can burst to a peak rate (taking advantage of excess capacity), at rates selectable up to the UNI Interface Rate (suiting TCP-Based applications such as HTTP and FTP)

These capabilities are deemed sufficient to enable migration from existing mass-market broadband services, as well as supporting current and future data services. In the future, additional parameters may be configurable within the AVC Bandwidth Profile, to enable a finer degree of traffic handling.

6.5 Bandwidth Specification Model – Connectivity VC

The Access Seeker is required to nominate the capacity for each required traffic class within the Connectivity VC at time of service order.

The selectable capacities are subject to provisioning rules, and are detailed in the Product Overview. The CVC Bandwidth Profile is shown in Table 6.

Table 6 Bandwidth Profile – CVC

Traffic Class	Component	Units	Description
TC_1	CIR	Mbps	CIR requirement for TC_1. Note TC_1_PIR set to TC_1_CIR by default.
TC_2	CIR	Mbps	CIR requirement for TC_2 Note TC_2_PIR set to TC_2_CIR by default.
TC_3	CIR	Mbps	CIR requirement for TC_3 Note TC_3_PIR set to TC_3_CIR by default.
TC_4	CIR	Mbps	CIR requirement for TC_4 Note TC_4_PIR set to TC_4_CIR by default.

Note that capacity specified within a Connectivity VC bandwidth profile is inclusive of the S/C-TAGs, as per the service frame definition in Figure 4.

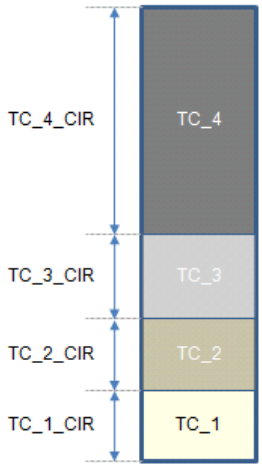

Figure 15 Bandwidth Specification Model (Connectivity VC)

Supported bandwidth profile options for the Connectivity VC are provided in the Product Overview. Note that where an Access Seeker under-dimensions the amount of CVC capacity for a given traffic class, with respect to the aggregate AVC capacity, any NFAS performance objectives for that class are not applicable.

6.6 Traffic Contention and Congestion Management

Access Seekers are free to control their own End-User experience, through contention applied through dimensioning of capacity between the AVC and CVC.

Contention may be applied at the traffic class level, allowing Access Seekers to independently control the economics and operation of each class. This is controlled by careful dimensioning of AVC and CVC capacity, on a Traffic Class basis, to ensure a level of contention appropriate for each respective higher-layer application.

Access Seekers must be aware of the implications of further contending NFAS services, as this will effectively degrade services beyond the performance and throughput guidelines for the NFAS network, as supplied by NBN Co.

6.7 Priority Identification

An Access Seeker may use a number of methods to indicate relative priority of individual service frames depending on the NFAS interface. The available methods differ for the UNI and NNI, as shown in Table 7.

Table 7 NFAS Priority Marking Options

Marking Scheme	UNI	NNI
PCP field (IEEE802.1p)	Y	Y
DSCP (RFC2474)	Y	N
Un-marked	Y (default)	N

6.8 Priority Code Point Encoding

An Access Seeker must align to the IEEE802.1P and DSCP settings indicated in Table 8 to map traffic into NFAS traffic classes at the UNI and NNI.

These ingress assignments are valid for purchased NFAS traffic classes only. Ingress traffic which has a PCP/DSCP assignment that cannot be mapped to a purchased NFAS traffic class will be discarded at ingress. Access Seekers will be required to identify and validate all required UNI-D and NNI p-bit assignments during the on-boarding phase.

Table 8 NFAS Class of Service Encoding

Traffic Class	PCP/DSCP Assignment (Ingress)		
	CoS (UNI/NNI)	DSCP[7] (UNI)	
		DSCP	DSCP (Decimal)
TC_1	5	CS5, EF	40 – 47
TC_2	4	CS4, AF 41 – 43	32 – 39
TC_3	3	CS3, AF 31 – 33 CS2, AF 21 – 23	24 – 31, 16 – 23
TC_4	0	CS1, AF 11 – 13 CS0, Default	8 – 15, 0 – 7

Note that DSCP decimal values 48 – 63 are not mapped to an NFAS Traffic Class, and will be discarded at ingress.

[7] DSCP-mapping available at UNI-D only

NBN Co may in future introduce further traffic classes, and possibly modify PCP/DSCP assignments per traffic class. Access Seekers are advised to maintain a flexible approach to PCP/DSCP mapping in their service architecture.

The mappings for DSCP classes at the UNI are in accordance with RFC4594. Refer Appendix C – Class of Service Application for further details on translating IP QoS classes to NFAS using DSCP.

6.9 Priority Code Point Decoding

Egress CoS decoding is indicated in Table 9.

Note that traffic carried within TC_3 may be re-marked by the NFAS network. This will occur for TC_3 traffic that is in excess of the TC_3 CIR, however within the TC_3 PIR.

Table 9 NFAS Class of Service Decoding

Traffic Class	PCP/DSCP Assignment (Egress)
	CoS (UNI/NNI)
TC_1	5
TC_2	4
TC_3 (In-Profile)	3
TC_3 (Out-of-Profile)	2
TC_4	0

6.10 DSCP Mapping

DSCP mapping is an Access VC option that enables traffic priority to be derived at the UNI Based on the layer 3 TOS field of the Access Seeker's ingress IP service traffic. This capability option is not offered at the NNI.

The DSCP priority encoding mechanism is fully described in RFC4594, and related standards. This feature is intended to cater for scenarios whereby the Access Seeker's CPE supports IP-level QoS, but does not have the capability to mark the relative priority of service frames at layer 2 (using the 802.1P PCP field).

The translation between an Access Seeker's DSCP setting, and NFAS traffic classes is described within Table 8.

Note that this mapping enables an Access Seeker to take advantage of a range of IP DSCP PHB settings, without having to provide layer 2 priority markings.

The Access Seeker is encouraged to perform DSCP-to-PCP translation using CPE equipment. Whilst this capability will allow existing layer 3-Based CPE devices to take advantage of NFAS CoS capabilities, it will not provide access to the full breadth of product features at the UNI.

For example, a UNI that supports DSCP-mapping cannot support more than one Access VC. Standard, non-configurable DSCP settings will be used to map service frames into one of the NFAS traffic classes. All traffic on that UNI must be delivered through a single Access VC, as there is no ability to identify different Access VCs.

6.11 Default (Best Effort) Traffic Handling

For UNI_D configured as Default-Mapped, all ingress traffic will be mapped into TC_4.

For UNI_D configured as Tagged or Priority Tagged, any ingress traffic that does not map to a provisioned AVC Traffic Class will be discarded at ingress.

For UNI_D configured as DSCP-Mapped, any ingress traffic that does not map to a provisioned AVC Traffic Class will be discarded at ingress.

For all NNI configurations, any ingress traffic that does not map to a provisioned CVC Traffic Class will be discarded at ingress.

7 Multicast

The NFAS network provides a layer 2 multicast feature for the support of an Access Seeker's higher-layer IP-based multicast architectures, as used for IPTV applications.

This capability is intended to enable Access Seekers to efficiently deliver a wide variety of streaming video content, to a large number of end-users. It has the following characteristics:

- The Access Seeker is able to inject their entire channel line-up once at the Point of Interconnect
- The NFAS network dynamically distributes the content streams to the appropriate End-Users
- Each End-User subscribed to the Access Seeker's multicast service is able to view the desired content, whilst also enjoying receiving parallel services (such as voice and data)
- The NFAS multicast service can be dimensioned to handle a variety of high-definition, standard-definition and internet-quality video CODECs.

The NFAS network implements layer 2 multicast capabilities as shown in Figure 16. This diagram depicts a multicast service whereby the Multicast CVC terminates on a different NNI than the Unicast CVC. Note that it is possible for the Multicast and Unicast CVC to terminate on the same NNI, if the Access Seeker prefers.

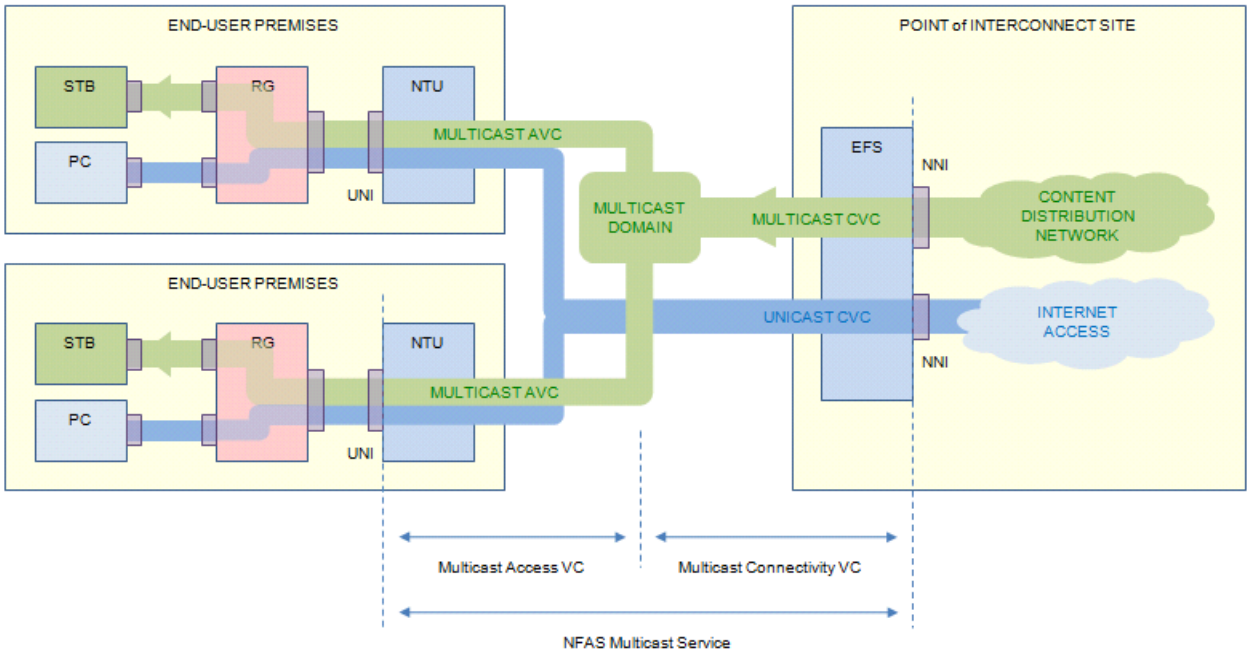

Figure 16 NFAS Multicast Service - High Level

7.1 Multicast Architecture

The NFAS multicast capability is Based on IEEE802.3 Ethernet multicast addressing and operation, using IGMPv3 snooping for interaction with the Access Seeker's IP-layer multicast services.

The NFAS network will "snoop" an End-User's upstream unicast data stream for IP-layer IGMP multicast packets. These IGMP multicast packets are interpreted as channel-change events in an End-User's IPTV service, and are used by NFAS to determine which of the Access Seeker's individual video streams to transmit to the End-User, in the downstream multicast AVC.

Currently, only IPv4 multicast services are supported by the NFAS multicast capability. Future enhancements to NFAS will extend support to IPv6.

7.2 Multicast Service Requirements

Each Access Seeker is required to undergo thorough interoperability testing with NBN Co prior to activation of any multicast services within the NFAS network, through the on-boarding process. This will involve the definition of a Configuration Template for use within the access network (AVC and UNI), which accommodates the multicast capability as well as any other UNI/AVC components being delivered as part of the same service.

In addition, the Access Seeker must provision a multicast CVC, which must be referenced by each End-User service order.

7.3 NFAS Multicast Operation

The NFAS multicast capability is implemented using a dedicated, multicast Access VC, operating in the downstream direction only. This multicast Access VC requires the presence of a bi-directional, unicast AVC for the communication of channel-change and control information from the End-User back into the Access Seeker's network.

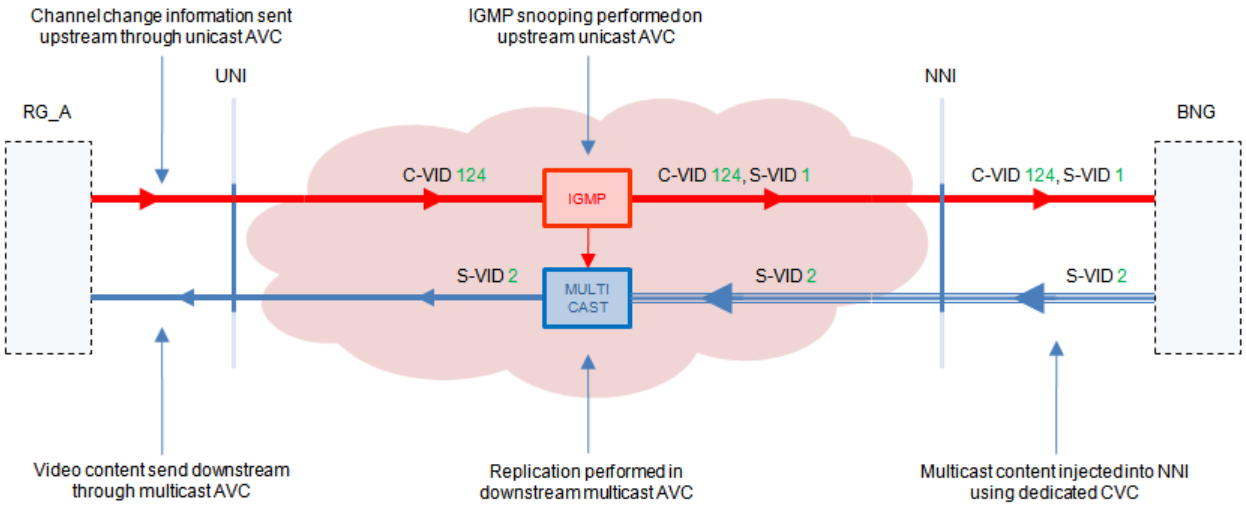

Figure 17 Multicast Operation

Figure 17 depicts the operation of a single multicast service. The upper (red) data flow represents an existing, unicast AVC that supports bi-directional data services. This AVC carries the channel change information from the End-User back into the Access Seeker's core network. The NFAS network intercepts this information, to detect any changes to the multicast data flow, as requested by the End-User.

The lower (blue) data flow represents the downstream, uni-directional multicast traffic flow. This data is injected at the NNI on a single CVC. It is then carried to each Connectivity Serving Area, where it is replicated to End-Users, in accordance with the IGMP information as intercepted in the upstream AVC.

7.4 UNI-D Interfacing

The NFAS multicast AVC is currently supported on UNI-D operating in DSCP-mapped or Default-Mapped mode only.

Whilst the multicast AVC is carried within the NFAS network separately from the unicast AVC, it is presented at the UNI-D as a single, merged data flow.

In the downstream direction, multicast and unicast service frames are expected to be identified by the Access Seeker's Residential Gateway or Set-Top Box at the IP level.

Likewise, in the upstream direction, the NFAS network will identify service frames that are associated with the multicast service, within the unicast AVC, Based on IP-level packet information (IGMP).

The Access Seeker may utilise DSCP-mapping at the UNI-D in order to prioritise upstream IPTV control packets over other applications, however this is not mandatory. This mode also allows the efficient scheduling of multicast service frames at egress to the UNI-D.

7.5 Connectivity VC Interfacing

NFAS multicast services are technically required to be serviced with a dedicated multicast Connectivity VC.

At the NNI, multicast services are addressed using Service Addressing Mode D (Refer Section 5.6).

7.6 Multicast IGMP Reporting

Upstream IGMP information may be passed through the NFAS network, and passed to the Access Seeker through the NNI, through a number of options:

- Passive Reporting
- Proxy Reporting

In addition, an Access Seeker may receive these reports through the following CVCs:

- Unicast CVC
- Multicast CVC
- Both (Unicast and Multicast CVC)

The reporting mechanism must be determined by the Access Seeker during the on-board phase, and cannot be changed through the standard service ordering process.

7.7 Multicast Performance

The Access Seeker must take into account a number of performance attributes when designing their IPTV service architecture to utilise NFAS multicast capabilities.

7.7.1 Capacity Allocation

The amount of capacity purchased in the multicast AVC must take into account the number of simultaneous video streams to a particular End-User. Any requests to simultaneously view channels above the subscribed maximum will be rejected. Further guidelines for capacity allocation will be provided in future.

The amount of capacity purchased in the multicast CVC must take into account the total number of unique video streams to be delivered to all End-Users.

8 Product components

This section describes the technical and operational requirements of each of the product components described in Section 3.3.

8.1 User Network Interface (UNI)

The supported UNI types are as follows:

- Data UNI (Ethernet port) – referred to as “UNI-D”
- Voice/Telephony UNI (Analog POTS port) – referred to as “UNI-V”

Each UNI is logically connected to an NNI via an Access and Connectivity VC, and may support one or more Access VCs.

8.1.1 UNI-D

Each UNI-D is regarded as a fully independent interface, operating in total isolation from any other port residing on the same NTU.

8.1.1.1 UNI-D Interface Attributes

The following interfaces are supported for UNI-D ports:

- 10/100/1000BASET/TX (Electrical)
- 1000BASE-TX (Electrical)
- 1000BASE-SX (Optical)
- 1000BASE-LX (Optical)

Note that there are restrictions as to which physical UNI-D interfaces are supported for different NTUs. Refer Section 12.2 for details of interfaces supported on different NTUs.

8.1.1.2 UNI-D Scalability Factors

Each UNI-D has two capacity metrics that define its ability to carry End-User services.

8.1.1.2.1 Line Rate

The Line Rate defines the rate at which the physical interface will transfer data. The UNI-D supports the following Ethernet Line Rates:

- 10Mbps
- 100Mbps
- 1000Mbps

The Line Rate sets the maximum bound on the information-carrying capacity of the link. Note that the effective throughput of any Ethernet interface is described in Appendix B – Effective Information Rate. Access Seekers are advised to become familiar with the inherent limitations of Ethernet in relation to the impact of framing overhead and asynchronous operation on bandwidth efficiency, and accommodate this within any NFAS capacity allocation.

The Access Seeker may set the Line Rate for a UNI-D in one of two ways:

- Auto-Negotiation - The Access Seeker selects the UNI-D Line Rate to be set by auto-negotiation between the UNI-D and the Access Seeker's CPE device. This may result in a Line Rate that is less than desired. Note that this is offered as an option for electrical UNI-D interfaces only.
- Fixed – The Access Seeker specifies the desired Line Rate for the UNI-D. This is the default and recommended mode.

The Access Seeker is advised to carefully consider the interface type, and/or auto-negotiation parameters at time of ordering.

8.1.1.2.2 Information Rate

The Information Rate defines the amount of logical capacity assigned to the UNI. This is calculated as the sum of all Access VC Bandwidth Profiles active on the UNI-D.

The UNI-D is capable of supporting an aggregate Information Rate up to the active Line Rate. For example[8], a UNI-D that has an auto-negotiated Line Rate of 100Mbps is capable of simultaneously supporting two Access VCs each with a CIR of 50Mbps.

A feasibility check is required upon addition of any Access VC to a UNI-D, to determine whether sufficient Information Rate capacity exists on the interface to support the incremental Access VC bandwidth profile. This feasibility check takes into account the fixed Line Rate of the UNI-D. A feasibility check will fail if the amount of available Information Rate capacity on the UNI-D is less than the Access VC capacity requirements. Note that where a UNI-D is configured for Auto-Negotiation, it is the responsibility of the Access Seeker to ensure that the requested Information Rate can be accommodated by the UNI-D.

Note that once provisioned, Access VC capacity will not be automatically re-adjusted as a result of changing Line Rates as a result of Auto-Negotiation. Should a UNI-D undergo a decrease in Line Rate whilst active (i.e. an End-User swaps CPE, and the UNI-D auto-negotiates to a lower Line Rate than previous), the End-User may experience increased discard on any provisioned Access VCs.

8.1.1.2.3 Access VC Support

A UNI-D can support a maximum of 8 Access VCs. Each of these Access VCs, and the associated UNI-D, must be under the control of the same Access Seeker.

A feasibility check will be performed upon addition of any Access VC to a UNI-D, to determine whether the number of allowable Access VCs has been exceeded.

[8] Note that this is a simplistic example that does not take into account Ethernet protocol overhead

8.1.1.3 UNI-D Interfacing

There are four options for addressing services at the UNI-D, shown in Table 10.

Table 10 Access VC Addressing Modes

UNI-D Mode	Maximum Number of Access VCs Supported at UNI-D	Comments
Default-Mapped	1	Default-Mapped service frames, as per IEEE802.3
Priority Tagged	1	Service frames tagged with a NULL VID, as per IEEE802.1Q
Tagged	8	Service frames tagged with a non-NULL VLAN, as per IEEE802.1Q
DSCP-Mapped	1	Untagged service frames, as per IEEE802.3, where priority information is encoded into the DSCP field, as per RFC2474
CE-VLAN Transparent	1	Connectivity VC interfaces direct to the UNI_D

The addressing mode must be specified at time of ordering a UNI-D, and determines how the Access Seeker interfaces to the Access VC and UNI-D. These modes have no impact of the operation or allocation of AVC C-TAGs at the NNI.

Note that a UNI-D must be operated in only one of the modes described in Table 10. The chosen Access VC addressing mode therefore impacts the capability of the UNI-D, and its ability to support multiple services.

8.1.1.3.1 Default-Mapped UNI-D

For a Default-Mapped UNI-D, all service frames ingress to the UNI-D will be encapsulated with a C-TAG by the NFAS service. Downstream service frames will be stripped of the C-TAG at egress to the UNI-D.

Note that in a Default-Mapped configuration, a UNI-D can support only a single Access VC.

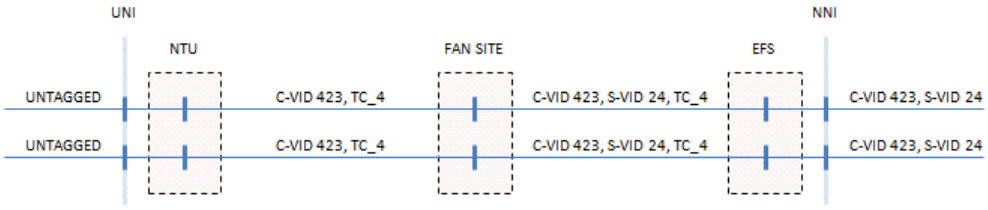

Figure 18 Example Default-Mapped UNI-D – Different Connectivity VCs

Figure 18 depicts a single NTU, with two Access VCs delivered on different Default-Mapped UNI-D. These services are operated by different Access Seekers, and carried using different Connectivity VCs and NNI. This situation (for example) may describe a Smart Metering service operating at the same residence as a Home Alarm service. Both services will be provided by completely separate Access Seekers, and operate in complete isolation.

For this example, note the following:

- For Default-Mapped UNI-D, all service frames at the UNI-D are mapped into a single Access VC. The C-VID assigned to the Access VC is not passed across the UNI-D.
- Service frames are mapped into the default transport class (TC_4).
- Both Access Seekers have elected to re-map the C-VIDs at egress to the NNI. Because these services are delivered on different Connectivity VCs, their C-VIDs (both set to C-VID=423) can overlap.
- Because the Connectivity VCs both terminate on different NNI, their S-VIDs can overlap (both set to S-VID=24).

8.1.1.3.2 Priority-Tagged UNI-D

For a Priority-Tagged UNI-D, ingress service frames must contain a C-TAG with VID set to 0. The C-PCP field is used to determine any further CoS handling by the NFAS service. All ingress service frames will have the C-VID set to the agreed C-VID at ingress to the UNI.

At egress to the NNI, the NFAS-supplied C-VID will remain.

At egress to the UNI-D, the C-VID will be reset to 0, and the C-PCP will be preserved as per the value provided by the Access Seeker at the NNI.

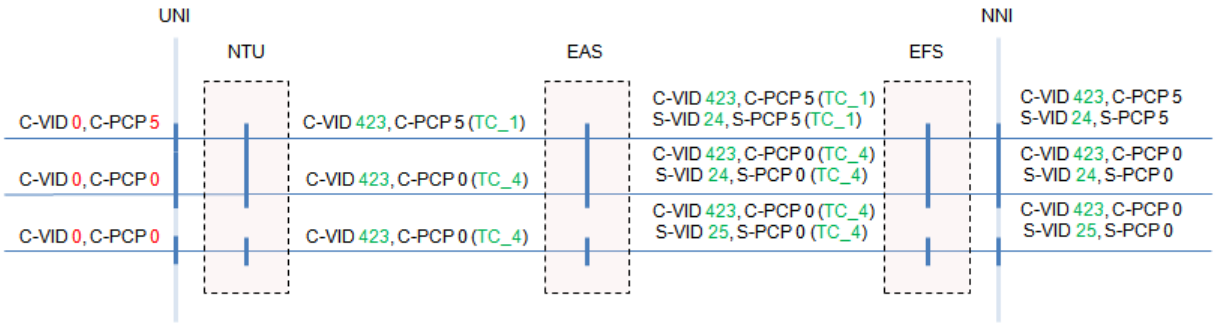

Figure 19 Example Priority-Tagged UNI-D

Figure 19 depicts two completely separate services being offered across different Priority-Tagged UNI-D at an NTU.

The top-most service marks traffic for TC_1 and TC_4 using the C-PCP fields of the C-TAG. Note that the C-VID field is re-written by the NTU at ingress to the internally allocated setting of C-VID=423. The bottom-most service contains traffic marked for TC_4 only.

For this example, note the following:

- The C-VID and S-VID fields overlap at the NNI between the two services, enabled by the ability to re-assign C-VID settings across S-TAGs, and for the S-VID to be unique within an NNI.
- In the upstream direction, the S-PCP is derived from the C-PCP.
- At egress from the NNI, the C/S-PCP values are preserved, irrespective of any re-mapping performed on the C/S-VID fields.

8.1.1.3.3 DSCP-Mapped UNI-D

Traffic allocated to a DSCP-Mapped UNI-D will have a C-TAG inserted at ingress to the UNI-D (as per a Default-Mapped UNI-D). By default the C-PCP field of the inserted C-TAG will be derived through a translation from the service frame DSCP field (refer Section 6.8), and will determine any further CoS handling by the NFAS service.

The C-TAG VID will be set to the internally allocated C-VID at ingress to the UNI-D. If the service frame contains a C-TAG at ingress, it will be over-written with these values. At egress to the UNI-D, the C-TAG is stripped.At ingress to the UNI-D, once the DSCP-mapping to C-VID has taken place, the traffic will be handled as per Priority-Tagged UNI-D, described in Section 8.1.1.3.2. At egress to the NNI, the internally supplied C-TAG VID/PCP is passed through to the Access Seeker.

Note that this option does not exist for downstream traffic handling (from the NNI to UNI-D). It is expected that the Access Seeker will encode all downstream priority information at layer 2.

8.1.1.3.4 Tagged UNI-D

An NFAS service may use tagged traffic across the UNI-D in order to identify one of multiple Access VCs. A single C-VID is used to identify each Access VC.

Where the Access Seeker requests the UNI-D to be Tagged, the Access Seeker may also request a single, specific C-TAG VID setting that the Access Seeker's CPE device will use to identify traffic for each Access VC (at the UNI-D). If none is requested, NBN Co will provide its own, internal C-TAG VID setting, allocated during the ordering/activation process.

Figure 20 depicts Access VCs being delivered across Tagged UNI-D at a single NTU. Two services are delivered on the top-most UNI-D and aggregated onto a common Connectivity VC/NNI. A single service is delivered in the bottom-most UNI-D (on a different Connectivity VC/NNI).

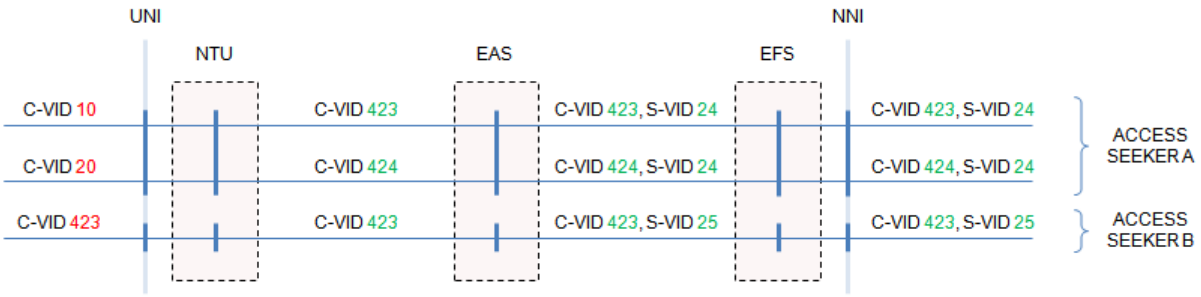

Figure 20 Example Tagged UNI-D

Note that for the upper-most service, the C-VID values of 10 and 20 are used to identify different Access VCs at the UNI-D, both of which aggregate to the same Connectivity VC. These C-VID values are re-mapped by the UNI-D at ingress, and then re-mapped again at the NNI at egress. For these services, the Access Seeker has elected to re-map the S-VID to a value that is unique within their own aggregation network.

The bottom-most service is carried with an S-VID value as allocated by NBN Co at time of service activation.

8.1.1.3.5 CE-VLAN Transparent UNI-D

Figure 21 depicts the operation of a UNI-D in CE-VLAN Transparent mode.

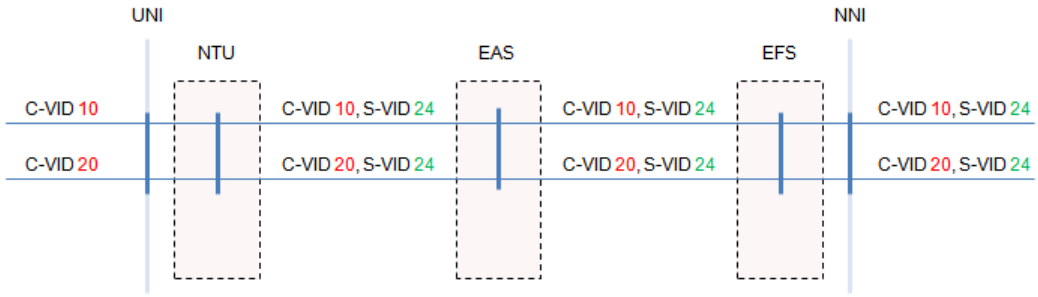

Figure 21 Example CE-VLAN Transparent UNI-D

In this mode, the S-VLAN is used for switching down to the UNI-D level. Within this, the Access Seeker is able to configure the full range of CE-VLAN VIDs.

In the example shown, the Access Seeker has configured S-VID 24 at the NNI to deliver a CE-VLAN Transparent service to the UNI. Across this S-TAG, the End-User is operating CE-VLAN values 10 and 20.

8.1.1.4 UNI-D Functional Attributes

8.1.1.4.1 Frame Forwarding

The UNI-D implements forwarding of service frames as per IEEE802.1ad[8], section 8.6.

Table 11 UNI-D Frame Forwarding Details

MAC Address	Application	Default Behaviour	Optional Configurable Behaviour
01-80-C2-00-00-00	Bridge Group Address	Discard	None
01-80-C2-00-00-01	IEEE Std 802.3 PAUSE	Discard	None
01-80-C2-00-00-02	LACP/LAMP	Discard	None
	Link OAM	Discard	None
01-80-C2-00-00-03	IEEE Std. 802.1X PAE address	Discard	None
01-80-C2-00-00-04 - 01-80-C2-00-00-0F	Reserved	Discard	None
01-80-C2-00-00-10	All LANs Bridge Management Group Address	Discard	None
01-80-C2-00-00-20	GMRP	Discard	None
01-80-C2-00-00-21	GVRP	Discard	None

01-80-C2-00-00-22 - 01-80-C2-00-00-2F	Reserved GARP Application addresses	Discard	None
01-80-C2-00-00-30 - 01-80-C2-00-00-3F	CFM	Tunnel	Peer[9]

Note the following definitions:

- Discard – Service Frame will be discarded at ingress to the NFAS network
- Peer – The Service Frame will be terminated within the NFAS network
- Tunnel – The Service Frame is passed to the A/CVC and carried through the NFAS network

8.1.1.4.2 Auto Negotiation

All electrical UNI-D ports provided at the NTU each individually support auto-negotiation as per IEEE802.3ab. Alternately, an Access Seeker may request, through the service order, a specific interface rate of 10Mbps, 100Mbps or 1000Mbps.

All interfaces are configured as full duplex. Auto-negotiation is not supported on optical UNI-D ports.

8.1.1.4.3 MAC Address Limitations

Each UNI-D is capable of supporting up to four simultaneous MAC source addresses. This imposes a limit on the number of layer 2 devices that an Access Seeker can connect directly to a UNI-D. Any attempt to connect a number of devices directly to a UNI-D that exceeds this limit will result in traffic from the newly-attached devices being discarded.

The NFAS network will learn the first four MAC source addresses detected at ingress to the UNI-D, based upon ingress service frames. A MAC address ageing function ensures that any obsolete MAC addresses are removed from the active list, after a period of 300 seconds.

Note that this limitation applies for the UNI-D, irrespective of the service type, and does not imply MAC address-based forwarding for unicast services based on 1:1 VLANs.

Access Seekers are encouraged to use a layer 3 device to interconnect to the UNI-D, to avoid any issues arising from MAC address restrictions.

8.1.1.4.4 Operations, Administration and Maintenance (OAM)

Options for peering with the Connectivity Fault Management (CFM) functions of the NFAS network are reserved for future release.

8.1.1.4.5 Resiliency

By default, the UNI-D is an unprotected physical interface. If an unprotected UNI-D suffers a failure, all services being delivered across that UNI will be disrupted.

Resiliency options are available using Point-to-Point Ethernet deployment. These details are still being investigated.

[9] Note feature not available first release.

8.1.2 UNI-V

Each UNI-V provides a standard analogue telephony port for the provision of POTS services via IP Telephony. Refer Section 12.2 for the number of UNI-V interfaces supported per NTU.

The Access Seeker is responsible for interfacing to, and controlling all required features of the integrated UNI-V to support the IP-based telephony service. Each UNI-V is associated with a SIP User Agent for independent configuration.

Note: In the first release, a single UNI-V and single SIP UA per NTU is supported.

8.1.2.1 General Configuration Overview

8.1.2.1.1 Ethernet Connectivity

Each UNI-V is automatically provisioned with a separate, dedicated Access VC. It is not possible to direct UNI-V service frames to/from an existing Access VC, connected to a different UNI port. This dedicated Access VC is used to carry all UNI-V traffic including:

- SIP Signalling messages
- RTP Media
- Management and Operations traffic, such as configuration file transfers.

Each Access VC will be automatically provisioned as 150kbps TC_1. Refer to section 6.1.1 for a description of TC_1 performance.

8.1.2.1.2 VLAN Configuration

All UNI-V will be placed into an N:1 VLAN per CSA. At the NNI, the N:1 VLAN will be identified with a single (S-TAG) VLAN tag. Refer to Section 5.5 for a description of the Service Addressing Mode C.

Each CSA with UNI-V services will require at least one Connectivity VC. The Connectivity VC will need to be dimensioned for the aggregate concurrent telephony channel capacity required plus signalling requirements by the Access Seeker. The diagram below illustrates an example VLAN Configuration across multiple CSAs:

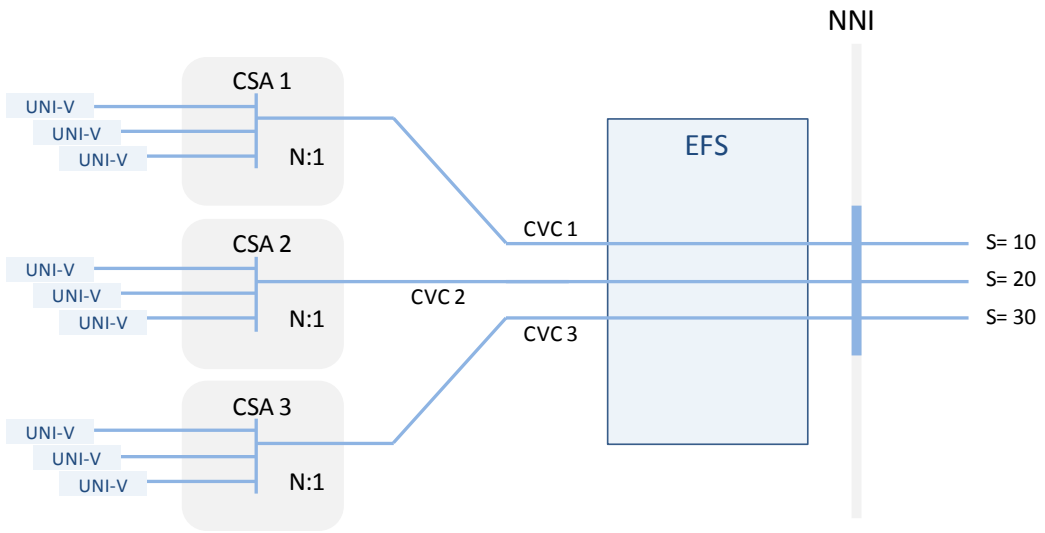

Figure 22 – UNI-V Example VLAN Configuration across multiple CSAs

8.1.2.1.3 Traffic Management and Identification

All traffic types from the UNI-V are treated as TC_1 within the NBN Co network and at NNI egress will be marked with an 802.1P PCP value of 5. Access Seekers must ensure NNI ingress traffic to UNI-V services is also marked with PCP value 5.

To enable prioritisation of traffic within the Access Seeker infrastructure, the NTU will mark IP packets generated by the UNI-V in the upstream direction with the following DSCP markings:

Table 12 – UNI-V DSCP Markings

Traffic Type	DSCP Marking (Decimal)
SIP Signalling	40
RTP Media	41
Management and Operations	42

8.1.2.1.4 IPv6 Support

The NFAS UNI-V supports IPv4-Based SIP services in the first release. NBN Co intends to support IPv6-Based SIP services in future.

8.1.2.1.5 Layer 3 Connectivity

It is the responsibility of the Access Seeker to manage allocation of IP addresses and associated network parameters to the SIP User Agent associated with each UNI-V. DHCP will be used as the mechanism to manage address distribution and configuration file transfers.

Access Seekers must provide DHCP server infrastructure and assign the following parameters:

- IP Address (IPv4)
- Subnet Mask
- Default Router Address (IPv4)
- DNS server (required if a hostname is used for proxy server SIP URI)
- Configuration Server (Option 66)

Within the NFAS infrastructure, DHCP Option 82 fields will be populated with the Service ID of the Access VC attached to a given UNI-V.

8.1.2.2 Configuration of the SIP Profile

Access Seekers will be responsible for creating and maintaining an XML Based configuration template for each active UNI-V in use on the network. The UNI-V ATA will download its configuration file upon boot-up of the NTU.

As part of the on-boarding process, NBN Co will work with Access Seekers to develop the XML configuration file. The file will contain all configurable parameters and in conjunction with NBN Co, Access Seekers may vary parameters to suit their soft switch or telephony feature set as required.

8.1.2.3 Supported Telephony Features

Most IP-based telephony features are implemented in the Access Seeker's soft switch infrastructure. However, certain telephony features involve NTU support through different electrical characteristics on the UNI-V Interfaces.

The following features are supported by the NTU (assumes support from the Access Seeker Soft switch):

- Call Waiting
- Calling Number Display
- Call Blocking (Calling Number Blocking)
- Distinctive Ring
- DTMF Digit Recognition

8.1.2.4 Local Number Portability

Local number portability will be the responsibility of the Access Seeker.

8.1.2.5 Dial Plan Configuration

Each SIP profile on the NTU will have an associated telephony dial plan which enables Access Seekers to configure appropriate digit recognition rules. Unlike traditional circuit switched networks, digits are typically transmitted *enbloc* in packet voice solutions. The NTU will "wait" for the entire digit string to be input before sending the string to the Access Seeker's soft-switch infrastructure.

Since numbers on the PSTN vary in length, specific patterns must be recognised to minimise post-dialling delays. Patterns are region specific and it is the Access Seeker's responsibility to ensure an appropriate dial plan is initially provisioned and, importantly, maintained for telephony End-Users.

Access Seekers may also use the dial-plan to provide "special" access numbers such as short dialling between branch offices or for internal services such as voicemail.

8.1.2.5.1 Guiding Standard

The dial plan will be specified in accordance with RFC3435 [16] - Media Gateway Control Protocol (MGCP).

8.1.2.5.2 Dial Plan Syntax Overview

Refer to RFC3435 – Media Gateway Control Protocol.

8.1.2.5.3 Recommendations

It is recommended that dial plans be created in accordance with the Australian PSTN numbering plan available from the ACMA Website:

http://www.acma.gov.au/WEB/STANDARD/pc=IND_REG_TEL_NUMPOL

8.1.2.5.4 Digit Tone Detection

The ATA integrated within the NTU supports both in-band and RFC2833 [14] for DTMF tone detection and transmission across an IP network. NBN Co recommends Access Seekers implement in-band DTMF detection.

8.1.2.6 Physical Port Characteristics

Each UNI-V will exhibit the characteristics described in Table 13. Note that these figures are provided for guidance only, and are subject to change.

Table 13 UNI-V Physical Port Characteristics

Parameter	Specification
Guiding Standards	Communications Alliance AS/ACIF S002:2005 Analogue interworking and non-interference requirements for Customer Equipment for connection to the PSTN
Physical Interface	Miniature 6-position socket as specified in ANSI/TIA 968 A 2002 [10].
Maximum Loop Length	150 metres
Loop Voltage	42 to 56VDC
Signalling	User Selectable – refer to configuration template table.
Ring Voltages	40 to 72V RMS
Ringer Equivalence	3 REN per UNI-V
Loop Current	20mA
Impedance	Refer to diagram below

8.1.2.6.1 Interface Impedance

Each UNI-V will exhibit a complex impedance, compliant with the Australian PSTN as shown below. The configuration is in accordance with Communications Alliance Technical Specification S002.

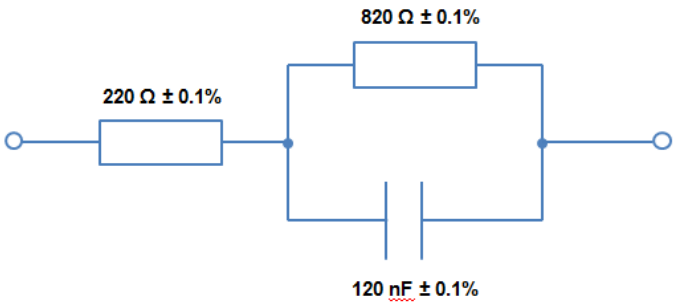

Figure 23 Impedance Configuration for UNI-V Interfaces

8.1.2.7 Service Tones

The following service tones are provided by each UNI-V.

Table 14 Service Tones

Tone	Description	Cadence or Tone
Standard Dial Tone	The standard Australian dial tone indicating the line is ready to accept digit input.	A combination of three frequencies at 400, 425 and 450Hz, lasting for up to 30 seconds prior to pressing a digit on the keypad.
Alternate Dial Tone	A “stutter” dial tone indicating the line is ready to accept digit input, however a special condition such as “message waiting” is enabled on the line.	To be advised.
Confirmation Tone	The tone which is used to confirm the acceptance of a feature being enabled	A single continuous tone at 425Hz
Ringing Tone	The tone used to indicate the called party is in the ringing state	A 425Hz tone repeated twice with a sequence of 200ms ON, 200ms OFF followed by a 2 second delay.
Busy	The tone which is used to indicate the called party is busy	A repeated 400Hz tone turned on and off every 375ms.
Network Busy	The tone to indicate the network is congested or unavailable.	A repeated 400Hz tone turned on and off every 375ms, but with alternating 10dB attenuation every second tone.

8.2 Access VC

8.2.1 Overview

The Access VC implements the C-VLAN component of an IEEE802.1ad Provider Bridge, as described in Section 5.

An Access Seeker may deliver multiple End-User applications (such as voice and video) using a single Access VC (using Class of Service to manage the capacity between applications), or using two, dedicated Access VCs (one per application).

8.2.2 Access VC Scalability

The maximum number of Access VCs that can be supported on a single UNI port depends on the UNI type and operation. Refer Section 8.1.1.2.3.

Access VCs are isolated from each other via the use of distinct S-TAG/C-TAG VIDs, and can be individually dimensioned according to the service needs of each End-User. An Access VC can be scaled in capacity (through its bandwidth profile), within the bounds of the product constructs and the physical limits of the underlying access network technology.

8.2.3 DHCP Option 82 Support

An Access VC may be optionally configured to provide support for DHCP Option 82, as per Section 3.9.1, TR-101 [22].

DHCP Option 82 allows for two fields to be set:

- Circuit-ID
- Remote-ID

NBN Co will insert DHCP Option 82 fields into upstream DHCP DISCOVER messages ingress to the Access VC at the UNI-D. The fields will be set as follows:

Circuit-ID – The Circuit-ID will not be populated. If a CPE attached to an Access VC populates the Circuit-ID field, the NBN Co infrastructure will strip this field.

Remote-ID – The Remote-ID will be set to the NBN Co Service ID and will use the following format. The first three bytes will signify the AVC product prefix and the next 12 bytes will be a unique string identifying the Access VC. Concatenated together, these values will form the Access VC Service ID. If CPE attached to the Access VC populates the Remote-ID field, the NBN Co infrastructure will replace it with the format below.

Figure 24 – DHCP Option 82 Remote-ID Field Format

8.2.4 PPPoE Intermediate Agent Support

An Access VC may be optionally configured to provide support for PPPoE Intermediate Agent, as per Section 3.9.2, TR-101 [22].

PPPoE Intermediate Agent allows for two fields to be set:

- Circuit-ID
- Remote-ID

NBN Co will insert PPPoE Intermediate Agent Option 82 fields into upstream PPP PADI messages ingress to the Access VC at the UNI-D. The fields will be set as follows:

Circuit-ID – The Circuit-ID will not be populated. If a CPE attached to an Access VC populates the Circuit-ID field, the NBN Co infrastructure will strip this field.

Remote-ID – The Remote-ID will be set to the NBN Co Service ID and will use the following format. The first three bytes will signify the AVC product prefix and the next 12 bytes will be a unique string identifying the Access VC. Concatenated together, these values will form the Access VC Service ID. If CPE attached to the Access VC populates the Remote-ID field, the NBN Co infrastructure will replace it with the format below.

Figure 25 – PPPoE Intermediate Agent Remote-ID Field Format.

8.3 Connectivity VC

This section details the technical interface and operational requirements of the Connectivity VC.

8.3.1 Overview

The Connectivity VC implements the S-VLAN component of an IEEE802.1ad Provider Bridge. This is an Ethernet Virtual Circuit that provides connectivity between an NNI and Connectivity Serving Area. It is dimensioned with a specific, configured amount of bandwidth capacity to deliver a higher-layer service (or number of services) to a range of Access VCs within a particular Connectivity Serving Area.

Under certain circumstances, the Connectivity VC may be used to deliver services directly to the UNI-D, bypassing the need for an Access VC (referred to as “CE-VLAN Transparent” operation).

The Connectivity VC may also be configured as an N:1 VLAN, for services delivered using the UNI-V interface.

The NNI, and all Connectivity VCs delivered through it, are specific to a single Access Seeker. It is possible that an Access Seeker may have multiple Connectivity VCs delivered using a number of NNI at a given location.

An Access Seeker may request to cancel a CVC. A CVC cancellation can only proceed once all member AVCs have been cancelled.

8.3.2 Connectivity VC Operation

Refer to Section 5.1 for the modes in which a Connectivity VC may be operated.

8.3.3 Connectivity VC Scalability

A single Connectivity VC can support up to 4000 Access VCs, and is delivered to a single Connectivity Serving Area. Each of the 4000 Access VCs is addressed using a single, unique C-TAG VID, locally significant to the Connectivity VC. The number of Connectivity VCs that an Access Seeker may purchase to a given Connectivity Serving Area is limited only by the NNI resources that the Access Seeker has purchased at the PoI.

Note that under CE-VLAN Transparent mode of operation, a Connectivity VC is used to deliver a single service directly to a UNI-D.

Connectivity VCs are isolated from each other on an NNI via the use of distinct S-TAG VIDs, and can be each individually dimensioned according to the service needs of each Connectivity Serving Area or UNI. Connectivity VCs using different service modes are able to co-exist on the same NNI.

Note that where an Access Seeker requires access to more than 4000 Access VCs on a given Connectivity VC, it is necessary to utilise additional Connectivity VCs.

8.3.4 Connectivity VC Interfacing

The Connectivity VC is directly accessed by the Access Seeker at the NNI. The VLAN tagging options for interfacing to the Connectivity VC at the NNI are described in Section 5.

For CE-VLAN Transparent operation, the S-TAG is stripped at egress to the UNI-D, and applied by the NFAS service at ingress to the UNI-D. Under this mode, the S-TAG is not visible by the Access Seeker at the UNI-D.

The Connectivity VC S-VID will be validated at ingress to the NNI. Any traffic that does not comply with this tagging structure, or contains S-TAG VID settings that are not as per agreed values, will be discarded at ingress to the NNI.

8.3.5 Connectivity VC DSCP Mapping

The Connectivity VC does not support DSCP mapping (including CE-VLAN Transparent mode).

8.3.6 Connectivity VC 802.1P Discard

Under congestion, any discard of service frames from a Connectivity VC will be in accordance with Section 6.

All discard events at the ingress to the Connectivity VC are recorded as service metrics.

8.3.7 Connectivity VC Capacity Management

Access Seekers are advised to monitor the growth of CVC capacity within an NNI, and allow for any ordering lead-times when determining capacity head-room. Service metrics are provided to enable the effective management of Connectivity VC capacity, and to enable the proactive identification of potential areas of service degradation in a manner that allows appropriate lead-times for action.

Ongoing Connectivity VC management involves a number of factors. As Access VCs are incrementally added to a Connectivity VC, the following must be considered:

- The aggregate utilization of capacity within a Connectivity VC will increase, leading to potential discard. Once a certain level of discard per traffic class has been reached, an Access Seeker may wish to consider augmenting the Connectivity VC with either more capacity, or with another Connectivity VC. This is to ensure a consistent experience among End-Users, and ample lead-time to accommodate seamless growth in End-User services.
- Any contention ratios and economics that the Access Seeker is trying to achieve through leveraging Connectivity VC capacity across multiple Access VCs will change, impacting the end-to-end user experience
- The available C-TAG VID pool within an S-TAG will decrease. A Connectivity VC can accommodate a finite number of Access VCs, before the C-TAG VID space for that Connectivity VC is exhausted. This will impact the ability to service incremental Access VCs,

as the C-TAG pool is finite (refer Section 8.3.3). The Access Seeker is encouraged to monitor the number of Access VCs active within a Connectivity VC, taking note of the rate of growth and the lead times for additional, parallel Connectivity VCs. This is necessary to ensure that incremental End-User service growth is not held up by the lead-times associated with Connectivity VC deployment.

8.4 Network-Network Interface (NNI)

The Network-Network Interface (NNI) is a physical interface that is used to deliver one or more Connectivity VCs to the Access Seeker. The NNI physical interface will be cabled to an Optical MDF within an NBN Co Point of Interconnect facility, and forms a Service Boundary Point. An Access Seeker is expected to interface the NNI directly to their backhaul, or local networking equipment for connection to their core network.

Each NNI is configured as a member of an NNI group using IEEE802.1ad Link Aggregation, which associates a number of Ethernet links together into a logical bundle. An Access Seeker may create NNI groups for the purpose of redundancy, capacity augmentation, or both. An NNI group must be configured as one of single-chassis, chassis-diverse, or site-diverse.

Each NNI group is capable of delivering CVCs that provide access to all CSAs available to that Point of Interconnect. This allows an Access Seeker to access all End-Users within a CSA, through a single NNI group.

An Access Seeker will be required to undergo inter-operability testing, and NNI group accreditation as part of the on-boarding process, before any NNI can be ordered. Refer NBN Co Consultation Paper – Access Seeker Accreditation (Fibre Network).

An Access Seeker may request through the order process for an NNI interface to be provisioned in a “down” state. Billing will commence when the NNI is provisioned and placed in the “down” state.

An Access Seeker may request to cancel at the NNI or NNI group level. If the group is specified, then there is no need to specify the individual NNI interfaces. A cancellation can only proceed once all member CVCs have been cancelled.

All NNI requests are subject to feasibility, and may experience a long lead-time (approximately 1-2 months).

8.4.1 NNI Interface Attributes

The following physical interfaces are may be selected by the Access Seeker upon ordering an NNI:

- 1000BASE-LX (10km)
- 1000BASE-ZX (40km)
- 10GBASE-LR (10km)
- 10GBASE-ER (40km)

8.4.2 NNI Resiliency

A number of resiliency options exist for Access Seekers to configure a higher availability across the NNI.

8.4.2.1 Single Chassis

Single Chassis configuration allows an Access Seeker to logically bundle a number of physical NNI off the same NBN Co EFS chassis.

Under this option, the interfaces will be grouped into a standard IEEE802.1ad Link Aggregation relationship, allowing a load-shared operation. For example, if a physical link within the NNI group goes down, its load will be transferred to the remaining links within the group. NBN Co is not responsible for any further traffic congestion that may result.

Within a Single Chassis NNI group, there can be an even/odd number of physical interfaces, allowing a group to consist of 1, 2, 3, 4, 5, 6, 7 or 8 physical links. The effective capacity of an NNI group in single chassis mode will be equal to the aggregate interface capacity. For example, if there are 8 physical GE interfaces within an NNI group, then the maximum CVC capacity may be 8Gbps.

The Single-Chassis option exists at all points of interconnect. NBN Co will make every effort to ensure that member links of an NNI group terminating on the same chassis exhibit a degree of diversity across line cards and internal control cards, however this cannot be guaranteed.

8.4.2.2 Chassis-Diverse

An NNI group may be delivered across diverse chassis within the same NBN Co Point of Interconnect site.

For chassis-diversity, the interfaces will be grouped into an active/standby (1+1) relationship, defined at the chassis level. For example, if an NNI group has a total of 4 physical NNI, then two will be provisioned on one chassis, and two on another. One chassis will be active, the other standby. If a failure occurs within any one of the active links, then a full switchover to the standby chassis will occur.

Within a chassis-diverse NNI group, there must be an even total number of physical interfaces, allowing NNI groups of 2, 4, 6 or 8 physical NNI. This requires that individual NNI are ordered and activated in pairs.

The Chassis-Diverse option exists at all points of interconnect.

8.4.2.3 Site-Diverse

An NNI group may be delivered across different NBN Co Point of Interconnect sites.

For site-diversity, the interfaces will be grouped into an active/standby relationship (as per chassis-diversity), defined at the site level. For example, if an NNI group has a total of 4 physical NNI, then two will be provisioned in one site, and two in another. One site will be active, the other standby. If a failure occurs within any one of the active links, then a full switchover to the standby site will occur.

Within a site -diverse NNI group, there must be an even total number of physical interfaces, allowing NNI groups of 2, 4, 6 or 8 physical NNI. This requires that individual NNI are ordered and activated in pairs.

The Site-Diverse option only exists for points of interconnect that are able to access more than one CSA.

8.4.3 NNI Reporting

NBN Co will report on the following metrics for each NNI group, on a weekly/monthly/quarterly basis:

- Aggregate, historic throughput
- Discard history
- CVC accumulation trend (i.e. historic growth of S-TAGs)
- Event history and metrics (e.g. uptime)

This is in addition to per-CVC metrics at the NNI, and is provided to assist Access Seekers in proactively managing NNI resources, and take into account lead-times when ordering additional NNI capacity.

8.4.4 NNI Scalability Factors

8.4.4.1 NNI Group Scalability

Each NNI group may comprise of up to 8 physical NNI, which may be required to be ordered individually, or in pairs, depending on the resiliency option chosen.

All interfaces within an NNI group must be the same interface speed. Note that the optical characteristics may be different (e.g. all must be GE, but some can have long-range optics, some with short-range optics). Note that the ability to deliver additional NNI into an NNI group is at the discretion of NBN Co.

A feasibility check will be required upon addition of any Connectivity VC to a NNI Group, to determine whether the number of allowable Connectivity VCs has been exceeded, and whether the aggregate Information Rate has exceeded the group's effective Line Rate. NBN Co will provide an NNI group ID where a new group is created as part of the order process.

8.4.4.2 Line Rate

For an NNI group configured as Single-Chassis, the aggregate Line Rate is determined by the sum of the constituent Interface Rates.

For an NNI group configured as Chassis-Diverse, or Site-Diverse, the aggregate Line Rate is determined by the sum of the constituent Interface Rates divided by two.

For example, an NNI group with 4x1Gbps links configured as Single-Chassis will have an effective Line Rate of 4Gbps. This means that the Access Seeker can provision up to 4Gbps of Connectivity VC capacity onto the NNI group.

The same NNI group operated as Chassis-Diverse, or Site-Diverse will have an effective Line Rate of 2Gbps.

8.4.4.3 Information Rate

The NNI Information Rate is calculated as the sum of all Connectivity VC bandwidth profiles active on the NNI group.

The NNI is capable of supporting an aggregate Information Rate up to the active Line Rate (taking into account adjustments for NNI groups as per Section 8.4.4.2.

A feasibility check is required upon addition of any Connectivity VC to an NNI, to determine whether sufficient Information Rate capacity exists on the interface to support the incremental Connectivity VC bandwidth profile. This feasibility check takes into account the programmed Line Rate of the NNI. A feasibility check will fail if the amount of available Information Rate capacity on the NNI is less than the Connectivity VC aggregate CIR.

Note that since Auto-Negotiation is not supported on NNI interfaces, the NNI Line Rate cannot change once the interface is active, unless a modification is made to the NNI group.

8.4.4.4 Connectivity VC Support

An NNI group can support up to 4000 Connectivity VCs, including any mix of 1:1 and N:1 Connectivity VC types. Any limitation on CVC support is imposed by the S-TAG address space, and not NNI interface capacity or capabilities.

A feasibility check will be required upon addition of any Connectivity VC to a NNI, to determine whether the number of allowable Connectivity VCs has been exceeded.

8.4.5 NNI Functional Attributes

8.4.5.1 Frame Forwarding

The NNI implements forwarding of service frames as Table 15, providing all Connectivity VC VLAN tag conditions are met.

Table 15 NNI Frame Forwarding Details

MAC Address	Application	Default Behaviour	Optional Configurable Behaviour
01-80-C2-00-00-00	Bridge Group Address	Discard	None
01-80-C2-00-00-01	IEEE Std 802.3 PAUSE	Discard	None
01-80-C2-00-00-02	LACP/LAMP	Peer	None
	Link OAM	Discard	None
01-80-C2-00-00-03	IEEE Std. 802.1X PAE address	Discard	None
01-80-C2-00-00-04 - 01-80-C2-00-00-0F	Reserved	Discard	None
01-80-C2-00-00-10	All LANs Bridge Management Group Address	Discard	None
01-80-C2-00-00-20	GMRP	Discard	None
01-80-C2-00-00-21	GVRP	Discard	None
01-80-C2-00-00-22 - 01-80-C2-00-00-2F	Reserved GARP Application addresses	Discard	None
01-80-C2-00-00-3X	CFM	Tunnel	Peer

Note the following definitions:

- Discard – Service Frame will be discarded at ingress to the NFAS network
- Peer – The Service Frame will be terminated within the NFAS network
- Tunnel – The Service Frame will be passed to the A/CVC and carried through the NFAS network

8.4.5.1.1 Operations, Administration and Maintenance (OAM)

Options for peering with the Connectivity Fault Management (CFM) functions of the NFAS network are reserved for future release.

8.4.5.2 LACP/LAMP

The NNI supports optional peering of Link Aggregation control frames with the NFAS service.

There is no capability to support LACP/LAMP transparency across the NFAS service. All frames identified as LACP/LAMP will be discarded at ingress at the UNI. Any frames identified as LACP/LAMP at ingress to the NNI will be treated according to the NNI's Group policy for Link Aggregation (where NNI protection is active) or discarded.

8.4.5.3 Class of Service

The NFAS traffic class model will operate transparently across an NNI Group, under all diversity configurations.

8.4.5.4 Frame Distribution Function

Service frames will be distributed across an NNI group Based on a fixed algorithm TBA.

9 Interfacing to the NFAS Network

This section details the mechanisms and processes that an Access Seeker must be aware of when planning to integrate NFAS services into their own network.

9.1 Configuration Template

A Configuration Template simplifies the ordering process by capturing the static attributes of the UNI and AVC components required for interfacing CPE, and operating services across the NFAS network. For an Access Seeker, these settings will be closely aligned with the needs of a particular retail product, and are not expected to change across different End-User services.

The Configuration Template is defined and validated during the on-boarding phase, and cannot be modified through the service ordering process. Any changes to the Configuration Template must be managed in conjunction with NBN Co, and may require further accreditation testing. Access Seekers are free to define a number of Configuration Templates, to cater for any variations in their service delivery.

Configuration Templates are required for the access (UNI, AVC) portion of the NFAS network only. Connectivity components (CVC, NNI) do not require a Configuration Template.

During the service order process, the Access Seeker must reference a Configuration Template, as well as provide settings for each of the variable, service-level attributes for each component.

Table 16 describes two example Configuration Templates. Configuration Template 1 describes a single, basic data service, delivered on a UNI-D, using no Class of Service. Configuration Template 2 describes a triple-play (voice, data and video) service, delivered using a UNI-V and UNI-D.

Table 16 Configuration Template Examples

Configuration Template	UNI	UNI Configuration Attributes	AVC	AVC Configuration Attributes
1	UNI-D_1	Default-Mapped	AVC_1	Unicast DHCP Option 82 Enabled TC_4 Enabled
2	UNI-V_1	SIP Details	AVC_1	TC_1_CIR = 150kbps
	UNI-D_2	Priority Tagged	AVC_2	Unicast DHCP Option 82 Enabled TC_1, TC_4 Enabled
			AVC_3	Multicast

These templates also facilitate the addition and deletion of End-User product components, through migrating specific access services between templates. For example, if an End-User currently with a basic data service wanted to purchase a triple-play package (requiring the addition of multicast and voice circuits), then the Access Seeker would simply transition them from Template #1 (using the above example) to Template #2.

The following sections detail the configuration attributes that must be defined by the Access Seeker during the on-boarding process, used to form a Configuration Template. These attributes are not available through the automated ordering process, and must be defined in conjunction with NBN Co.

Each specific Configuration Template will be validated during the accreditation phase, and used to order access services.

9.2 Service Modification

Any structural modifications to active access services must be within the scope of an Access Seeker's accredited Configuration Templates.

For example, if an Access Seeker requests the addition of a UNI-V and corresponding AVC to an existing End-User service, then the resulting configuration must have been approved during the on-boarding process. A feasibility check will be required to ensure that the addition of resources can be accommodated. This modification request will then be able to proceed as a migration from one Configuration Template to another, but may involve a service disruption.

9.3 Access VC Configuration Attributes

The following Configuration Attribute options are available for the Access VC:

- AVC Type (Unicast (1:1), Unicast (N:1), Multicast (N:1))
- Access Loop Identification (DHCP Option 82, PPPoE IA)
- Installation Options (Standard, Additional Validation)
- C-TAG Mapping at UNI
- Activation of individual Traffic Classes (TC_1, TC_2, TC_3, TC_4)
- Reporting Options (Standard, Service-Level, Performance)
- Activation of Service OAM
- Specification of IP version (IPv4, IPv6)

9.4 UNI Configuration Attributes

The following Configuration Attributes are available for the UNI:

- Installation options (Standard, Business, Infrastructure, Other)
- UNI Type (UNI-D, UNI-V)

For the UNI-D, the following additional configuration attributes are available:

- VLAN mode (Default-Mapped, DSCP-Mapped, Priority-Tagged, Tagged, CE-VLAN Transparent)

For the UNI-V, the following additional configuration attributes are available:

- IP Configuration (IPv4 address, 802.1p traffic class)
- SIP Configuration (authentication mode, signalling port)
- Voice Codec (G.711 a-law, G.729 a/b)
- Dial Plan

The Access Seeker will be required to determine these settings during the on-boarding phase, which will be programmed into the Access Seeker's Configuration Template.

10 Service Management

10.1 Business to Business (B2B) Interface

10.1.1 Systems Management Architecture

NBN Co's system architecture is built around the key principles of Service Provider enablement and automation. By placing the management controls of each End-User service in the hands of the corresponding Access Seeker, operational overheads and manual order processing are minimised.

NBN Co system interfaces will span across a number of different domains as shown in Figure 26.

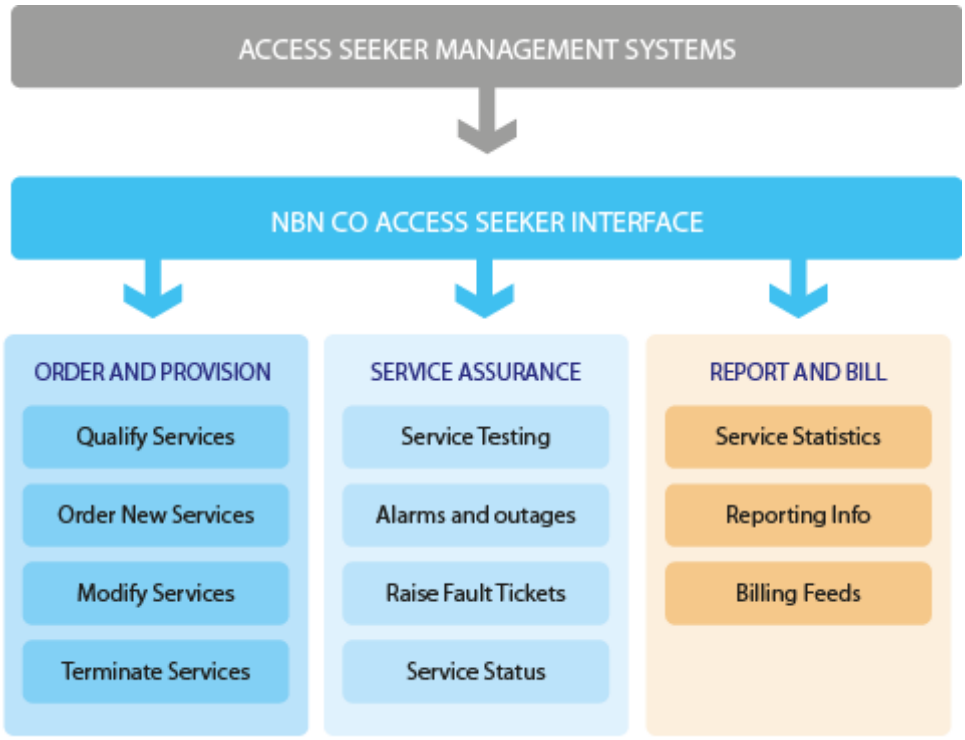

Figure 26 NBN Co Service Management Architecture

10.1.2 Online Gateway for Service Management

The B2B systems interface enables direct, machine-to-machine transactions to occur between an Access Seeker and NBN Co for the purpose of service provisioning, assurance, reporting and billing. In brief, the B2B interface:

- Prevents "rekeying" or "double handling" of End-User orders between Access Seeker systems and the NBN Co ordering interface
- Enables near real-time order status/progress data to be communicated to Access Seekers
- Enables a high volume of transactions to be exchanged
- Provides a high level of integration and interoperability across vendors and Access Seeker systems through being implemented using an open standards approach

- Provides a technology independent interface which can be leveraged through infrastructure upgrades

10.1.3 High Level Architecture

Access to NBN Co's B2B system is based on Web Services. Web Services provide an open, transaction Based interface capable of scaling to meet the high transaction volumes predicted for the system.

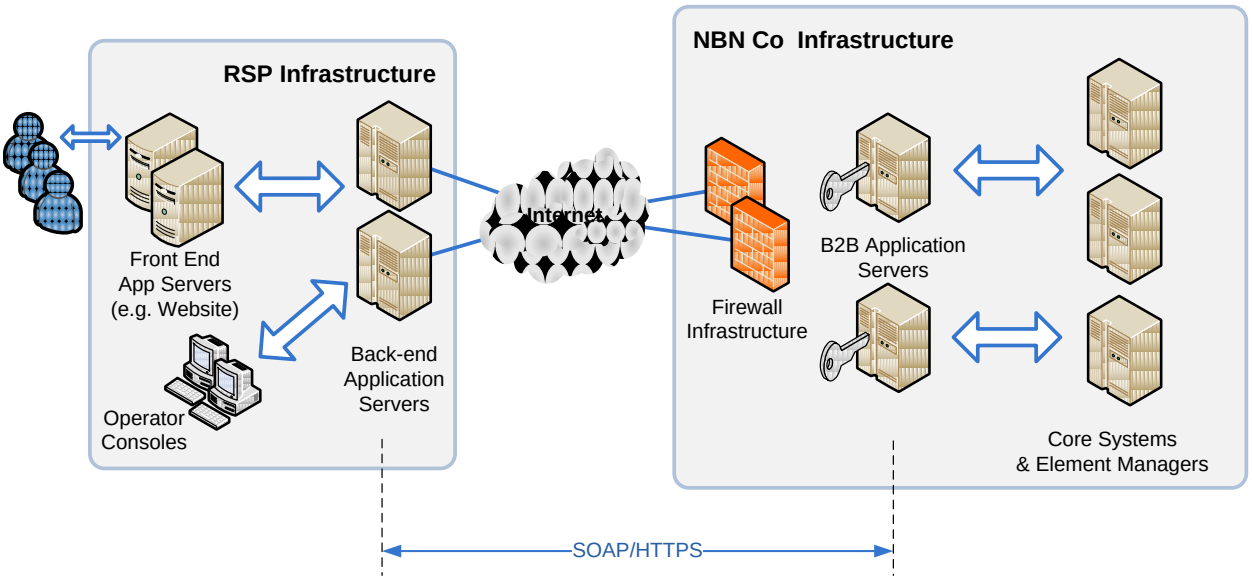

Figure 27 High Level B2B System Architecture

10.2 Service Ordering

A complete, end-to-end NFAS service consists of a serial arrangement of a number of access components (UNI, Access VC) and a number of connectivity components (connectivity VC and NNI). In most cases, it is expected that the connectivity components (and perhaps some access components) will be shared with other End-User services.

A number of the attributes of each service component will be statically defined through a Configuration Template, with the dynamic, per-service configurable attributes defined through each End-User service order. This section details all dynamic, per-service configurable attributes.

All NFAS service orders must reference a Configuration Template, defined during the on-boarding phase. This template details the interface-level settings for each service. In addition, there are a range of attributes that must be specified per-service, such as speed and assurance options.

Each access service order must contain the following:

- Physical location of requested service (End-User premises, GNAF, NTU ID, etc)
- Specification of Configuration Template
- Specification of all service attributes required to fulfil the Configuration Template.

10.3 Access Component Management

Access components consist of the UNI and Access VC, and are expected to be ordered at a high frequency, with a high degree of automation. These orders are expected to be triggered by End-User interest in an Access Seeker's retail product offerings. The Access Seeker is responsible for ordering the appropriate NFAS access components through a Configuration Template, aligned to a retail package the End-User has selected from the Access Seeker.

A high-level access component order process is shown in Figure 28. This diagram shows how the Service Qualification and Order Processes work in conjunction to provide an initial check for service and product availability, and an order workflow that aims to keep the Access Seeker informed throughout the delivery phase.

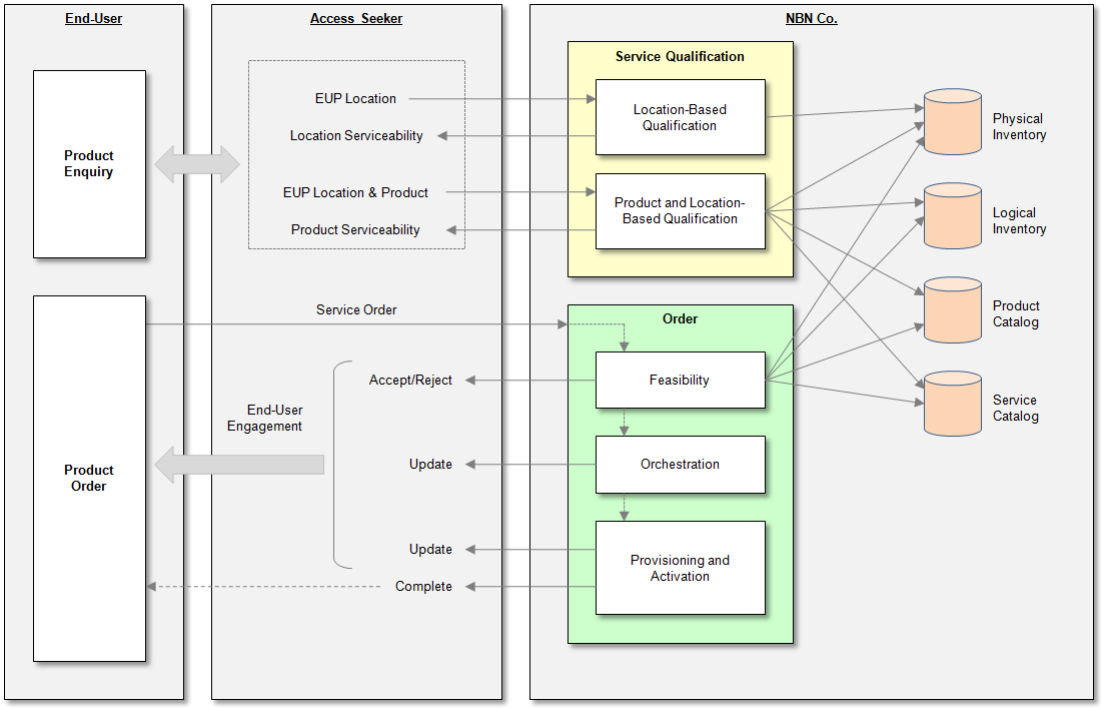

Figure 28 High-Level Access Component Order Process

These processes are described at a high level in the following sections. Note that complete process details are beyond the scope of this document.

10.3.1 Service Qualification

A two-level Service Qualification process is provided, to allow Access Seekers to accurately and rapidly determine the level of service that can be delivered to an End-User before progressing to the order phase. This ensures that services can be ordered confidently, and delivered in a deterministic manner. The following Service Qualification capabilities are provided:

- Location-Based Service Qualification
- Location and Product Based Service Qualification

These qualification processes are intended to act as a front-end to the service ordering process. Further details provided through the order process will be used to identify any additional feasibility or commercial elements required to generate an accurate, expected lead-time and quote.

10.3.1.1 Location-Based Service Qualification

A Location-Based Service Qualification will consist of an automatic inventory lookup Based on preliminary End-User premises information, to determine the following:

- Whether an NBN Co connection is currently available to the End-User Premises. If it is not currently served, then an expected date will be provided when service will be available at that location Based on NBN Co's published rollout plans.
- An indication of network capabilities will be returned, Based on the access technology available at that location

This level of detail will not indicate which products may be delivered at the location. It is intended only to help Access Seekers check End-User Premises coverage information during the NBN Co rollout period.

As an example, a Location-Based Service Qualification for an End-User Premises may indicate that the premises is located in a rural area, and is currently served by NBN Co's footprint. It may also indicate that the location is served by satellite and is capable of a 12Mbps service.

Alternatively, the Location-Based Service Qualification may fail because an NBN Co service is not yet available at the nominated premises. An indication of when this premise will be available for connection will be returned, along with the expected level of capability (e.g. maximum CIR of 12Mbps).

In order to perform a Location-Based Service Qualification, the Access Seeker is required to provide the following information:

- End-User Premises location (specified as a GNAF ID, GPS co-ordinate, Street Address or NTU ID)

If successful, a Location-Based Service Qualification will yield the following data, returned to the Access Seeker:

- Indication of whether the requested location is available for NBN Co connection (including the Connectivity Serving Area).
- Indication of the access technology used to service this location (highlighting any capability or capacity limitations)

If unsuccessful, one or more of the following results will be returned:

- NBN Co connection not available to the premises (with expected date of coverage)
- Feasibility required (for specific products only)

The Location-Based Service Qualification process will not validate whether an Access Seeker has existing or sufficient Connectivity VC or NNI resources in place to accommodate the requested service.

10.3.1.2 Service Qualification by Location and Product

A Product and Location-Based Service Qualification will consist of an automatic product catalog and network inventory lookup Based on preliminary End-User premises information and desired product capabilities, to determine the following (in addition to the Location-Based Service Qualification):

- Whether there is sufficient spare network capacity and resources within the network to deliver the service. In the case of capacity limitations, a response will highlight what subset of the requested capabilities can be delivered.
- Any additional information regarding the degree of complexity in delivering the requested product to the End-User Premises (highlighting the need for special product features). This is intended to help refine the lead-time estimate for delivery.

As an example, a Product and Location-Based Service Qualification for a mass-market 50Mbps service to an End-User Premises located in a rural area may produce a response that the requested service cannot be delivered, because though the premises is available for connection, it is served by Satellite and therefore only capable of receiving a 12Mbps service.

Alternately, a Product and Location-Based Service Qualification for a 100Mbps business service with CE-VLAN transparency to an End-User Premises located in a metro area may produce a response that the requested service can be delivered, however the lead-time will need to take into account the installation of a specific NTU, and a necessary Connectivity VC.

In addition, the Product and Location-Based Service Qualification process assists an Access Seeker in determining the following:

- Whether the Access Seeker is setup to provide services to the area in question (I.e. does the Access Seeker currently have an NNI and Connectivity VC capable of being used for delivering services to the requested location?).
- Whether the Access Seeker has any existing, shared network resources that could be leveraged for the requested service (I.e. an existing UNI port which a new Access VC could be delivered on)
- Whether there are any additional product components that are required to deliver the service (i.e. a Connectivity VC)

In order to perform a Product and Location-Based Service Qualification, the Access Seeker must provide the following information:

- Product Type (specified using the Access Seeker's defined Configuration Template)
- End-User Premises location (specified as a GNAF ID, GPS co-ordinate, Street Address or NTU ID)
- Requested date of service

If successful, a Product and Location-Based Service Qualification will yield the following data, returned to the Access Seeker:

- NTU ID (if not already provided)
- Indication of whether an unallocated UNI port is present on the NTU
- UNI ID and details of any free UNI ports (interface type, etc)
- Indication of support for requested product
- Indication of whether the requested Access VC bandwidth profile can be accommodated on the access network
- Details of the Connectivity Serving Area
- Assessment of installation category (may be used to provide a general indication of lead-time[10])

If unsuccessful, one or more of the following results will be returned:

- End-User Premises outside of existing service footprint (with expected date of coverage)
- No spare UNI ports on NTU (NTU ID provided for Access Seeker's own in-service UNI lookup)
- Feasibility required, by default or by option (for specific products only)
- Insufficient spare capacity on Access Network (with indication of available capacity)
- Access Network does not support requested service (e.g. requesting 100Mbps to an End-User premises located in an area served by satellite)

[10] Estimate only. No commitments will be made to lead-times shorter than those in the Access Seeker's NFAS contract.

- Requested product not supported (either by Access Seeker, or network)

The Product and Location-Based Service Qualification process will not validate whether an Access Seeker has sufficient Connectivity VC or NNI resources in place to accommodate the requested service. The returned CSA ID is expected to be sufficient to allow an Access Seeker to look up their own inventory database, to determine whether they have the following:

- A Connectivity VC to the CSA, of the correct functional type
- Sufficient capacity within the Connectivity VC to serve the Access VC
- Sufficient C-TAG address space within the Connectivity VC to accommodate the Access VC
- An NNI serving the Connectivity VC

It is the responsibility of the Access Seeker to ensure that these functional items are in place, and capable of accommodating the incremental access service.

Any lead-time estimates provided by the Product and Location-Based Service Qualification will not be commercially binding, and will not over-ride the Access Seeker's contracted lead-times.

10.3.2 Feasibility

The Service Qualification enables a high degree of automated, flow-through ordering Based on automated, real-time inventory checking. In some cases, however, it may be necessary to escalate an order for further feasibility work.

10.3.2.1 Standard Feasibility

A Standard feasibility check will validate the presence of all required service attributes, and network resources as identified by the Configuration Template and order details. This is a necessary step for any service order.

The following Standard Feasibility verification checks are performed for access (UNI/AVC) orders:

- The nominated Configuration Template is active, and accredited for use by the Access Seeker
- Any service attributes provided with the access order are valid within the Configuration Template, and legal NFAS settings (for example, certain combinations of NFAS traffic class speeds are not allowed)
- The Connectivity VCs as specified in the access order are active, and compatible with the nominated Configuration Template
- The Connectivity VCs as specified in the access order are able to accommodate the Access VCs as defined by the Configuration Template. Note that checks are only performed for addressing purposes (for example, that the number of Access VCs allowed per Connectivity VC has not been exceeded), and not for capacity, however requests to deliver an Access VC CIR/PIR higher than the CVC CIR/PIR will be rejected.
- The resources required to fulfil the access order are available. Note that in the case of a service modification, a delta will be calculated between what resources are already in place, and what additional resources are required.

The following Standard Feasibility verification checks are performed for CVC orders:

- The NNI as specified in the CVC order is active, and in operation by the requesting Access Seeker
- The NNI as specified in the CVC order is able to accommodate the requested Connectivity VC. Note that checks are only performed for addressing purposes (for example, that the number of Connectivity VCs allowed per NNI has not been exceeded), and not for capacity.

- Any service attributes provided with the CVC order are valid within allowable service attributes for a CVC. Note that there is no Configuration Template required for CVC or NNI orders. This assumes that the Access Seeker has passed inter-operability testing for the requested CVC service attributes, and any NNI features that may be required for support.

10.3.2.2 Desktop Feasibility

A Desktop feasibility check may be required where the service qualification fails (i.e. there are insufficient network resources to fulfil an order), or if the requested features necessitate a further analysis.

This may be the case for business services, in particular those that require point-to-point fibre access tails, or particular features such as CE-VLAN transparency.

These scenarios may result in a manual intervention to further analyse the data stored within the inventory database, and possibly correlate with any work-force related activities that may be underway as part of capacity augmentation, feature upgrade, network rollout, etc.

10.3.2.3 Detailed Feasibility

A Detailed feasibility check will be required when it is known that the requested features require a manual intervention. This is required for higher-end business services, where point-to-point fibre availability (for access diversity) may need to be deployed or assessed, or any facilities-related requests.

10.3.3 Orderable Service Attributes

Each of the product components contains a number of sub-attributes which must be supplied during the ordering process. These attributes must be provided in conjunction with a Configuration Template, which outlines the number and relationships between access (UNI and AVC) components.

Legend for “Type” column:

M – Mandatory. The Access Seeker must provide requested details.

S – Selection. The Access Seeker must select one of the offered options.

O – Optional. The Access Seeker may specify a value, if not then a default will be applied.

Where an attribute is market within the Type column as “Selection” (S), then the Setting column will indicate whether, for a given package, the attribute is:

“D” – Default. This attribute will be selected by default

“O” – Optional. This attribute will be offered as an optional selection.

“-“ – Not Applicable. This attribute is mandatory, or not currently offered.

10.3.3.1 Network Termination Unit (NTU) Service Attributes

Note there are no orderable service attributes for an NTU. An Access Seeker cannot directly order an NTU – the NTU device will be selected and provided by NBN Co as part of the deployment process for a UNI. Note that certain orderable attributes of the UNI and AVC components may trigger deployment of a different NTU type.

10.3.3.2 UNI Service Attributes

Cells shaded red indicates features that are not available for the first release of the NFAS product.

Table 17 Service Attributes – UNI (Generic)

Component	Attributes	Attribute Description	Selectable Options	Type	Setting
Installation Options	Location[11]	Ability to specify the location of the UNI. This assumes that a range of NTU options exist (indoor, outdoor, etc) - TBC	Indoor	S	D
			Outdoor		O
			Other		-
	Installation Lead Time	Lead time options for installation	Standard	S	D
			Lead Time Option 1		-
	UNI Initial State	The initial administrative state of the UNI once the service is provisioned.	Enabled	S	D
			Disabled		-
Assurance Options	Response Time Options	Options for enhanced response times	Standard	S	D
			Response Option 1		-
			Response Option 2		-
			Response Option 3		-
	Restoration Time Options	Options for enhanced restoration times	Standard	S	D
			Restoration Option 1		-
			Restoration Option 2		-
			Restoration Option 3		-
	Availability Options	Options for enhanced availability. Note that non-default availability	Standard	S	D
			Availability Option 1		-

[11] Selection of NTU location is TBA

		figures will imply different access delivery.	Availability Option 2		-
			Availability Option 3		-
Additional Options	End-User Premises cabling	Different levels of End-User Premises cabling able to be ordered with NTU installation	Default cabling	S	D
			Cabling Option A		-
			Cabling Option B		-
			Cabling Option C		-
	Power Type	Ability to request DC powered NTU for business applications	AC	S	D
			DC		-
	Backup Power Unit	Whether the NTU is supplied with a backup power unit.	Backup Power Unit Supplied with Battery	S	D
			Backup Power Unit Supplied (No Battery)		O
			No Backup Power Unit		O
		Backup power monitoring	Enabled	S	D
			Disabled		O

For a UNI-D order, the following additional service attributes apply.

Cells shaded red indicate features that are not available for the first release of the NFAS product.

Table 18 Service Attributes – UNI-D

Component	Attributes	Attribute Description	Selectable Options	Type	Setting
NTU Port	NTU Port ID	Ability to specify the NTU port that a UNI-D will be delivered upon. There may be circumstances where an Access Seeker will want a new UNI-D delivered on a specific NTU port. If not specified, NBN Co will automatically allocate the first free port as part of the ordering process. If an Access Seeker requests a specific port and it is not available, then the order is rejected.	NTU UNI-D PORT ID 0 (default) 1 – UNI-D Port 1 2 – UNI-D Port 2 3 – UNI-D Port 3 4 – UNI-D Port 4 Note that the default option will assign the first free UNI-D port.	M	M
PHY Attributes	Speed & Duplex	Configure the physical parameters of the Ethernet port. Speed and duplex configuration is available for copper-based UNI-D only.	Auto-negotiation	S	O
			10Mbps Full Duplex		O
			100Mbps Full Duplex		D
			1000Mbps Full Duplex		O
PHY Attributes	Interface Type	Physical interface options for the UNI-D. Note that certain Interface types may trigger a non-standard NTU deployment. The optical options will only be available for Business/Infrastructure applications.	10/100/1000BASE-T	S	D
			1000BASE-T		-
			1000BASE-SX		-
			1000BASE-LX		-

For a UNI-V order, the following additional service attributes apply.

Cells shaded red indicate features that are not available for the first release of the NFAS product.

Table 19 Service Attributes – UNI-V

Component	Attributes	Attribute Description	Selectable Options	Type	Setting
NTU Port	NTU Port ID	Ability to specify the NTU port that a UNI-V will be delivered upon. There may be circumstances where an Access Seeker will want a new UNI-V delivered on a specific NTU port. If not specified, NBN Co will automatically allocate the first free port as part of the ordering process. If an Access Seeker requests a specific port and it is not available, then the order is rejected.	NTU UNI-V PORT ID 0 (default) 1 – UNI-V Port 1 2 – UNI-V Port 2 Note that the default option will assign the first free UNI-V port.	M	-
SIP Attributes	SIP Configuration	The following SIP parameters must be supplied.	Outbound SIP proxy IPv4 address	M	D
			SIP Username	M	M
			SIP Password	M	M

10.3.3.3 Access VC Service Attributes

The following Access VC service attributes must be defined by the Access Seeker for every ordered access service. It is the responsibility of the Access Seeker to ensure that the requested Service Attributes correspond with the nominated Configuration Template.

Cells shaded red indicate features that are not available for the first release of the NFAS product.

Table 20 Service Attributes –Access VC

Component	Attributes	Attribute Description	Selectable Options	Type	Setting
End-Point Identification	UNI ID	Identification of the UNI that the Access VC is terminated on. The UNI ID for an AVC end-point will be determined through the Configuration Template. The Configuration Template will relate each AVC as a child of a UNI. A service order which references this template instantiates a UNI, which will be referenced using this field.	UNI ID	M	-
	Connectivity VC ID	Identification of the Connectivity VC that the Access VC is to be delivered on. This will implicitly indicate the CSA that the Access VC is located	CVC ID	M	-

		within.			
Installation Options	Installation Lead Time	Lead time options for installation.	Standard	S	D
			Installation Option 1		-
Assurance Options	Response Time Options	Options for enhanced response times. These are likely to be inherited from the UNI attributes.	Standard	S	D
			Response Option 1		-
			Response Option 2		-
			Response Option 3		-
	Restoration Time Options	Options for enhanced restoration times. These are likely to be inherited from the UNI attributes.	Standard	S	D
			Restoration Option 1		-
			Restoration Option 2		-
			Restoration Option 3		-
C-TAG Mapping	C-TAG (CVC)	An Access Seeker may choose a locally-significant C-TAG within the CVC.	Requested C-TAG at CVC (NULL for NBN Co-Supplied VID) Default = NULL C-TAG: (1 – 4000)	M	D
Bandwidth Profile	Bandwidth Profile (Upstream)	Upstream traffic class allocations applicable to tagged/priority-tagged/DSCP-mapped UNI-D. Optional traffic class allocations not available for Default-Mapped UNI-D.	AVC_TC_1_CIR	O	D
			AVC_ TC_2_CIR	O	-
			AVC_ TC_3_CIR	O	-
			AVC_ TC_3_PIR	O	-
			AVC_TC_4_PIR	M	D
	Bandwidth Profile (Downstream)	Downstream traffic class allocations applicable to tagged/priority-tagged/DSCP-mapped UNI-D. Optional traffic class allocations not available for Default-Mapped UNI-D.	AVC_TC_1_CIR	O	D
			AVC_ TC_2_CIR	O	-
			AVC_ TC_3_CIR	O	-
			AVC_ TC_3_PIR	O	-

			AVC_TC_4_PIR	M	D
Multicast	Multicast Channel Groups	A list of permitted channels for the ONT conditional access function must be provided.	Multicast Channel Groups (Refer Section 7)	M	-
	Multicast Domain	Where AVC_Type = Multicast, the Access VC must be mapped to a nominated multicast domain at the PoI. Specify the multicast domain's service ID.	Multicast Domain S-ID.	M	-

10.4 Connectivity Component Ordering

Connectivity components consist of the NNI and Connectivity VC, and are expected to be ordered at a low frequency, with a low degree of automation. These orders are expected to be triggered by an Access Seeker's infrastructure needs, in meeting the aggregate needs of its retail product offerings. The Access Seeker is responsible for proactively monitoring and ordering the appropriate NFAS connectivity components, taking into account any lead-times or the management of the End-User experience.

10.4.1 Service Qualification

Service qualification for connectivity components requires a manual service qualification request. There is no automated process defined.

10.4.2 Feasibility

It is anticipated that most connectivity orders will result in manual feasibility work.

10.4.2.1 Standard Feasibility

Where possible, a Standard feasibility check will validate the presence of all required service attributes, and network resources, to fulfil the order. This is a necessary step for any service order.

10.4.2.2 Desktop Feasibility

A Desktop feasibility check may be required where the service qualification fails (i.e. there are insufficient network resources to fulfil an order), or if the requested features necessitate a further analysis.

Connectivity services (such as the Connectivity VC or NNI) may require a desktop feasibility, where certain features are requested. Such features may consist of:

- Creation of a new NNI Bundle
- Provision of an NNI bundle
- Allocation of a new Traffic Class on an existing Connectivity VC

These type of features may result in a manual intervention to further analyse the data stored within the inventory database, and possibly correlate with any work-force related activities that may be underway as part of capacity augmentation, feature upgrade, network rollout, etc.

10.4.2.3 Detailed Feasibility

A Detailed feasibility check will be required when it is known that the requested features require a manual intervention. This is required for higher-end business services, where point-to-point fibre availability (for access diversity) may need to be deployed or assessed, or any facilities-related requests.

10.4.3 Order Process

Connectivity components (CVC, NNI) are ordered using a manual process. NBN Co will provide an order form which may be submitted electronically for processing.

10.4.4 Service Attributes

10.4.4.1 Connectivity VC

Cells shaded red indicate features that are not available for the first release of the NFAS product.

Table 21 Service Attributes – Connectivity VC

Component	Attributes	Attribute Description	Selectable Options	Type	Setting
End-Point Identification	NNI ID	Identification of the NNI bundle that the Connectivity VC is to be terminated on.	NNI Bundle ID (Existing)	M	D
	B-END CSA/UNI_D ID	Identification of the B-END that the Connectivity VC is terminated on. This will typically be a Connectivity Serving Area (CSA), however for CE-VLAN transparent services, the Connectivity VC will terminate on a UNI-D.	CSA ID	S	D
			UNI_D Order Reference		-
S-TAG Mapping	S-TAG (NNI)	An Access Seeker may choose a locally-significant S-TAG at the NNI. Optional parameter, if blank NBN Co will assign.	Requested S-TAG (0 for NBN Co-Supplied S-TAG) Default = 0 S-TAG: (1 – 4000)	M	D
Bandwidth Profile	Bandwidth Profile (Upstream)	The CVC_CIR capacity allocation for each individual traffic class in the upstream direction.	CVC_TC_1_CIR	O	O
			CVC_TC_2_CIR	O	-
			CVC_ TC_3_CIR	O	-
			CVC_ TC_4_CIR	M	M

	Bandwidth Profile (Downstream)	The CVC_CIR capacity allocation for each individual traffic class in the downstream direction. Note that the CVC is currently restricted to symmetric operation only. Therefore the upstream/downstream settings must be identical.	CVC_TC_1_CIR	O	O
			CVC_TC_2_CIR	O	-
			CVC_ TC_3_CIR	O	-
			CVC_ TC_4_CIR	M	M
Service Options	Enhanced Reporting	Ability to receive additional service and performance reports on the circuit	Standard	S	D
			Service		-
			Performance		-
		Trend Analysis to aid in Connectivity VC capacity management	Disabled	S	D
			Enabled		-
	Service OAM	Whether Service OAM is provided on the Connectivity VC.	Disabled	S	D
			Enabled		-
Installation Options	Connectivity VC Type	Type of Connectivity VC.	1:1	S	D
			CE-VLAN Transparent		-
			Multicast		-
	Lead Time	Lead time options for installation and modification	Standard	S	D
			Expedited		-

10.4.4.2 NNI

Cells shaded red indicate features that are not available for the first release of the NFAS product.

Table 22 Service Attributes – NNI

Component	Attributes	Attribute Description	Selectable Options	Type	Setting
Service details	Physical Location	Physical location of NNI	Specification of PoI Site	M	M

	NNI Bundle Membership	NNI Bundle Identifier. Used to nominate an existing (or new) Bundle that this NNI will be a member of.	Member of New Bundle (Nominate Bundle ID)	S	D
			Add to Existing Bundle (Provide Bundle ID)		O
Installation Options	Installation Type	Allowance for non-standard installation (e.g. including facilities)	Standard	S	D
			Other		-
	Installation Lead Time	Lead time options for installation	Standard	S	D
			Installation Option 1		-
	NNI Initial State	The initial administrative state of the NNI once the service is provisioned.	Enabled	S	D
			Disabled		-
Assurance Options	Response Time Options	Options for enhanced response times	Standard	S	D
			Response Option 1		-
	Restoration Time Options	Options for enhanced restoration times	Standard	S	D
			Restoration Option 1		-
	Availability Options	Options for enhanced availability.	Single Chassis (default)	S	D
			Chassis Diversity		-
			Site Diversity		-
NNI Type	Interface Type	Physical interface type.	1000BASE-LX	S	D
			1000BASE-ZX		O
			10GBASE-LX		O
			10GBASE-ZR		O
			Other		-

10.4.5 Service Termination

Termination of any shared connectivity components will require that all access components using those resources are fully terminated prior.

10.4.6 Service Relocation

It is not possible to relocate an existing Connectivity VC and/or NNI to a different port or physical location.

10.5 IP-Based Telephony Service Management

An IP-based telephony service as deployed using a UNI-V consists of a special arrangement of specific access components.

10.5.1.1 UNI-V Status

The following status information will be available to Access Seekers via the B2B Management System.

Table 23 UNI-V Status

Parameter	Description
Rx Invite Requests	This attribute counts received invite messages, including retransmissions
Rx Invite Retransmissions	This attribute counts received invite retransmission messages.
Rx Non-invite Requests	This attribute counts received non-invite messages, including retransmissions
Rx Non-invite Retransmissions	This attribute counts received non-invite retransmission messages
Rx Response	This attribute counts total responses received
Rx Response Retransmissions	This attribute counts total response retransmissions received
Tx Invite Requests	This attribute counts transmitted invite messages, including retransmissions
Tx Invite Retransmissions	This attribute counts transmitted invite retransmission messages.
Tx Non-invite Requests	This attribute counts transmitted non-invite messages, including retransmissions
Tx Non-invite Retransmissions	This attribute counts transmitted non-invite retransmission messages
Tx Response	This attribute counts total responses transmitted
Tx Response Retransmissions	This attribute counts total response retransmissions transmitted.
Failed to Connect Counter	This attribute counts the number of times the SIP UA failed to reach/connect its TCP/UDP peer during SIP call initiations
Failed to Validate Counter	This attribute counts the number of times the SIP UA failed to validate its peer during SIP call initiations

Timeout Counter	This attribute counts the number of times the SIP UA timed out during SIP call initiations.
Failure Received Counter	This attribute counts the number of times the SIP UA received a failure error code during SIP call initiations.
Failure to Authenticate Counter	This attribute counts the number of times the SIP UA failed to authenticate itself during SIP call initiations.
Analogue Port Off Hook Timer	This attribute is a high water mark that records the longest period of a single off-hook detected on the analogue port. Time is measured in milliseconds
Analogue Port Releases	This attribute counts the number of analogue port releases without dialling detected (abandoned calls).
Call Setup Failures	This attribute counts call setup failures
SIP Status	This attribute shows the current status of the SIP agent. Values are as follows: 0 = Ok/initial. 1 = Connected. 2 = Failed – ICMP error. 3 = Failed – Malformed response. 4 = Failed – Inadequate info response. 5 = Failed – Timeout.
Received Octets	This attribute counts the total number of octets received over the Access VC.
Sent Octets	This attribute counts the total number of octets sent over the Access VC.

10.5.2 Service Qualification

Refer Section 10.3.1.

10.5.3 Feasibility

Refer Section 10.3.2.

10.6 Service Assurance

This section describes the Service Assurance tools available for the management of Access Seeker services.

10.6.1 Service Assurance Strategy

NBN Co's Service Assurance Strategy aims to provide an integrated, efficient toolset for the rapid detection, diagnosis and resolution of network issues. This allows Access Seekers a level of operational visibility and control over their services that promotes the integration of NFAS services into their own network, and the rapid detection and diagnosis of service-impacting events.

10.6.1.1 Access Seeker Service Assurance

NFAS Service Assurance tools are used by Access Seekers for both proactive and reactive resolution of network issues. There are three levels at which an Access Seeker may employ the NFAS Service Assurance capabilities:

- B2B interface
    - On-demand status and diagnostic tools
- Periodic reports and metrics compilations
- Service-Level Ethernet OAM (not available first release)

All Access Seekers will have access to a range of Service Assurance capabilities (with commercially-available enhancements), as provided through the B2B interface. The goal of this is to provide the following, in an on-demand and historic fashion:

- Indication of the health and operational status of NFAS services for reactive fault detection
- Service-level metrics for long-term capacity planning
- Diagnostic capabilities for reactive fault resolution

10.6.2 Trouble Tickets

Trouble tickets can be lodged electronically using a range of interfaces (toll-free phone, email, web portal, B2B interface). It is assumed the Access Seeker has performed service testing and diagnosis via the assurance tools prior to logging a fault. A required field in the trouble ticket record will be the test reference numbers to provide NBN Co with a view of the test results at the time of the fault (particularly important for intermittent faults).

An Access Seeker has the ability to perform the following tasks regarding Trouble Tickets:

- Raise Trouble Ticket
- Query Trouble Ticket
- Modify Trouble Ticket
- Close Trouble Ticket
- Report Trouble Ticket

11 Network Attributes

This section details network-level attributes and characteristics that are relevant to the delivery of end-to-end services by Access Seekers.

11.1 Network Coverage

Refer to the Product Overview for details and timing of service and feature coverage footprints.

11.2 Maximum Frame Size

The NFAS network supports a maximum layer 2 Ethernet frame size of 2000 bytes at the NNI, inclusive of the S-TAG and C-TAG as depicted in Figure 19. This maximum frame size limitation may be referred to as the layer 2 Maximum Transfer Unit (MTU) of the NFAS network.

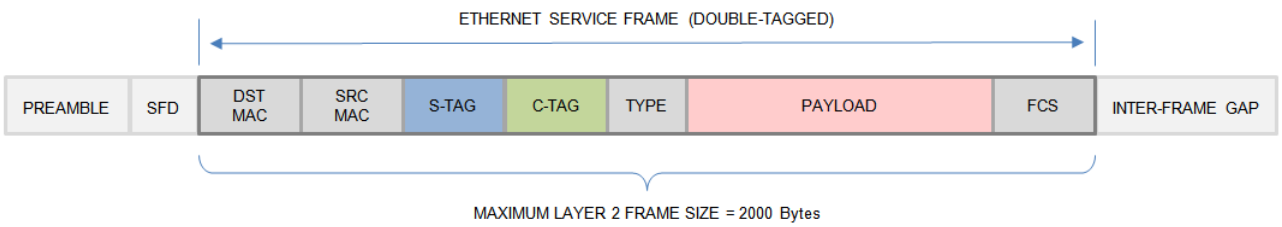

Figure 29 NNI Service Frame Definition

At the UNI-D, service frames are limited to the following:

- Default-Mapped and DSCP-Mapped : 1992 bytes
- Tagged and Priority-Tagged : 1996 bytes

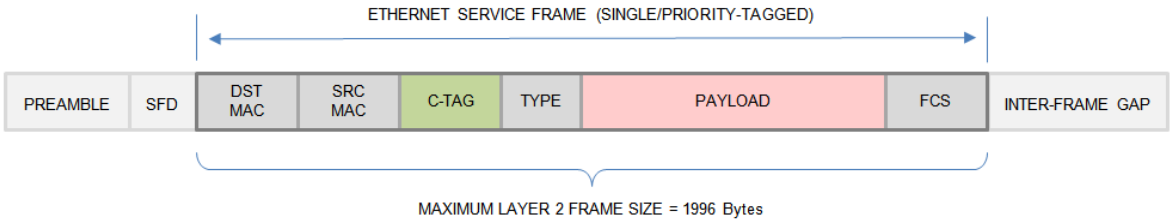

Figure 30 UNI-D Maximum Frame Size Definition (Single/Priority-Tagged Mode)

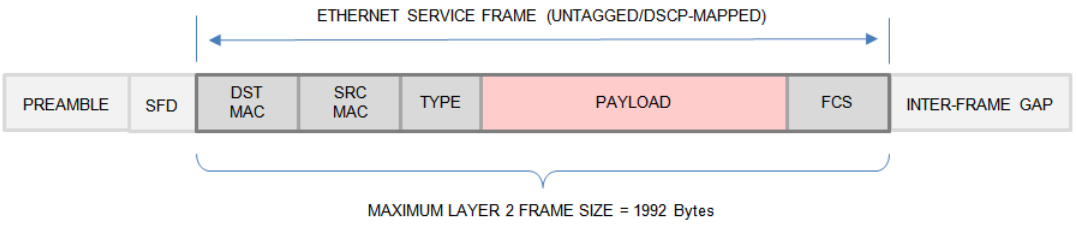

Figure 31 UNI-D Maximum Frame Size Definition (Default-Mapped/DSCP-Mapped Mode)

Any ingress service frame received at the UNI-D that exceeds this length will be discarded.

Any ingress service frame received at the UNI-D that is less than 64 bytes (excluding any VLAN tag applied by the Access Seeker) will also be discarded.

11.3 Security

This section describes the various security-related features and attributes of the NFAS offering.

11.3.1 Traffic Isolation

The NFAS network supports the delivery of Ethernet services across a shared infrastructure. Isolation of individual services across this shared infrastructure is critical for providing secure, deterministic and reliable services.

11.3.1.1 Isolation of traffic between Access Seekers

The use of IEEE802.1ad Provider Bridging ensures that all traffic within the NFAS network is isolated between all users at layer 2. This provides isolation of layer 3 IP domains between Access Seekers.

All 1:1 ingress service frames are mapped to internal S/C-TAG values within the NFAS network, ensuring all layer 2 traffic is traceable to a specific Access Seeker and UNI/NNI port. These internal mappings are performed at ingress, meaning that an Access Seeker cannot inadvertently or maliciously tag their traffic to overlap with an alternate Access Seeker. Any circuit mis-configurations are automatically detected and alarmed to NBN Co operational staff.

Additionally, the requirement for each UNI port to be under the control of a single Access Seeker provides a physical demarcation between Access Seekers offering services on the same NTU.

11.3.1.2 Isolation of traffic between End-Users

The same mechanisms employed in Section 11.3.1.1 also provide isolation for the protection of End-User traffic.

Additionally, the NFAS network does not permit direct communication between UNI ports or Access VCs. All direct communication between End-Users must be performed beyond the NNI, within the Access Seeker's service domain.

11.3.2 User Authentication

NFAS unicast services provide layer 2 virtual circuits between the UNI and NNI. These layer 2 circuits place the onus for End-User authentication within the responsibility of the Access Seeker. Although functions described within Sections 8.2.3 and 8.2.4 are provided to assist Access Seekers in identifying End-Users, the NFAS service does not take any steps to authenticate, or restrict the operation of, End-User services as operated by the Access Seeker.

11.3.3 Network Neutrality

The NFAS service does not interfere with, or differentiate between sites, hosts or domains of the internet services as provided by Access Seekers. The unicast services that NFAS provides are agnostic to the IP-layer data and protocols that operate within the Access Seeker's network, for the purpose of Network Neutrality.

Any interaction between the NFAS network and an Access Seeker's IP services are limited to the following:

- DSCP-mapping of IP QoS classifications into NFAS traffic classes at the UNI (selectable by the Access Seeker)
- Snooping IGMP/MLD traffic for the purpose of offering layer 2 multicast (for multicast services only)
- DHCP Option 82 and PPPoE Intermediate Agent insertion for End-User identification (selectable by the Access Seeker)
- Control and operation of the UNI-V SIP protocol functions (required for use of the UNI-V)

12 Deployment Guidelines

12.1 Delivery Options

NFAS supports the following access technologies for delivery of last-mile connectivity to the End-User Premises:

- GPON
- Point-to-Point Ethernet

The default deployment option is GPON. The deployment of Point-to-Point Ethernet technology as a last-mile alternative to GPON is at the discretion of NBN Co.

12.2 Network Termination Unit (NTU)

NFAS services are delivered to an End-User Premises using a physical Network Termination Unit (NTU). The purpose of this device is as follows:

- Deliver multiple, independent UNI to facilitate simultaneous services from multiple Access Seekers
- Provide optional service continuity in the event of a power failure
- Define a Network and Service Boundary Point at the End-User Premises, providing full management and visibility to operational staff
- Cater for a wide range of deployment and operational scenarios

There following NTU variants may be deployed by NBN Co, depending on the feature set requested:

- Residential NTU (NTU-R)
- Business NTU (NTU-B)

12.2.1 Residential NTU (NTU-R)

The NTU-R is intended for residential deployments only, primarily for single-dwelling premises[12].

[12] NTU-R is also applicable for Multi-Dwelling Units where fibre access is deployed to each tenancy

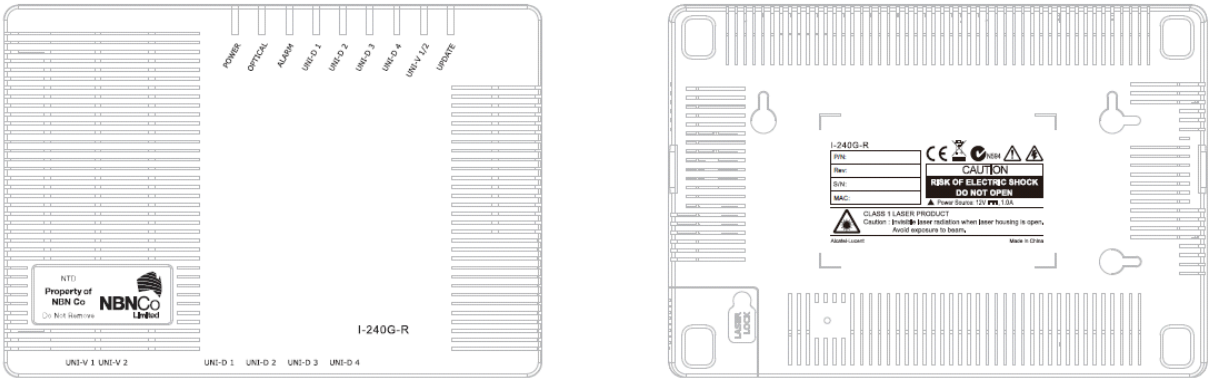

Figure 32 Indoor NTU-R

The indoor NTU-R variant is depicted in Figure 32.

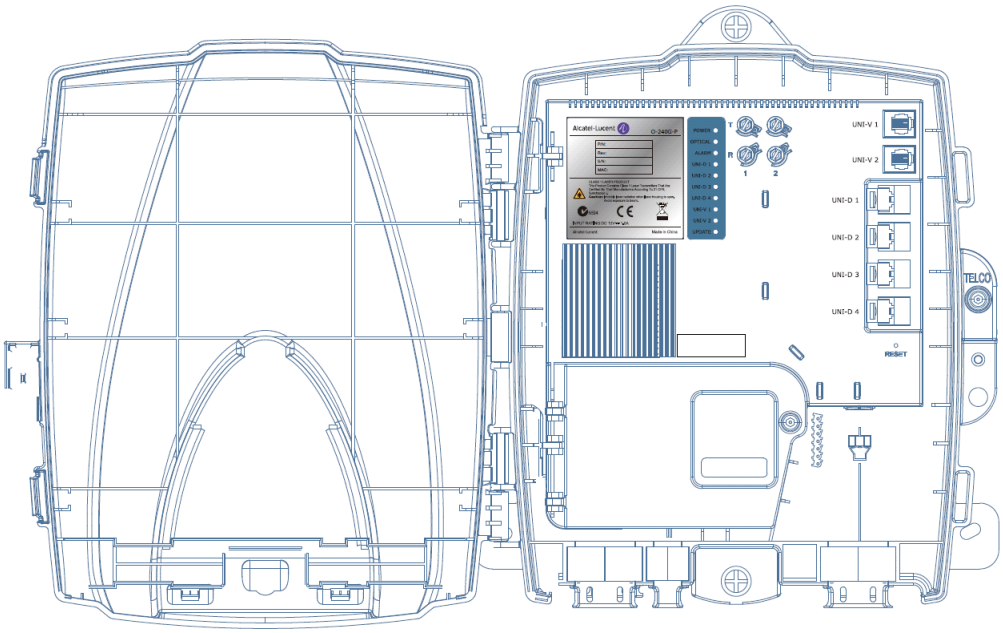

Figure 33 Outdoor NTU-R

The outdoor NTU-R variant is depicted in Figure 33.

The indoor and outdoor NTU-R variants are functionally identical. Both are capable of delivering individual services up to an NTU aggregate of up to 100Mbps CIR, and up to 1000Mbps PIR. For example, an NTU-R may support the delivery of 4x25Mbps CIR services, each delivered on a different UNI-D. Likewise, multiple 1Gbps PIR services may be supported by the NTU-R.

The NTU-R is expected to be available in a range of variants, suitable for outdoor and indoor deployments. The default deployment option is internal.

12.2.1.1 Physical Interfaces

The NTU-R supports a range of physical UNI ports, to facilitate simultaneous service delivery from multiple Access Seekers, including multiple services from each Access Seeker.

The NTU-R supports the following UNI ports:

- Four electrical 10/100/1000BASE-T Ethernet UNI-D ports.
- One or two SIP-based analogue telephony ports.

Figure 34 shows the arrangement of UNI-D and UNI-V ports on the indoor NTU-R.

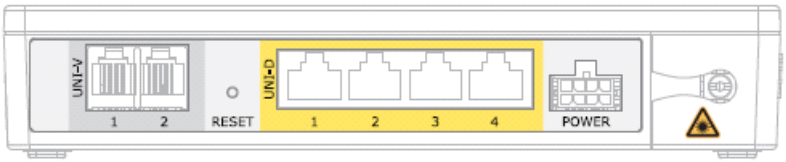

Figure 34 Interface Port Depiction on Indoor NTU-R

12.2.1.2 Power Supply

The NTU-R requires a local power source at the End-User premises. The NTU-R is supplied with a power supply that must be connected to a dedicated standard 240V, 10A Australian General Purpose Outlet (GPO). The NTU should be installed within 10 metres of the Power Supply. Typical power consumption of the NTU is approximately 15W, depending on the type and number of active UNI.

The battery backup PSU will be external to the NTU-R and should be installed indoors adjacent to the 240V AC power outlet. This battery backup capability will apply only to services delivered using a UNI-V port.

The battery backup unit will contain the following LED indicators showing battery status:

- Mains Power OK
- Battery Status (OK, needs replacing, failed).
- Running on battery backup (lit when mains power fails).
- A series of beeps are also provided to indicate the above conditions.

NBN Co's remote monitoring service will enable Access Seekers delivering services through a UNI-V to identify when a NTU is not being supplied with power.

12.2.2 Business NTU (NTU-B) (Future Release)

The NTU-B is intended for business deployments only.

The NTU-B is capable of delivering individual services up to 1Gbps. In addition, this device is capable of delivering multiple 1Gbps services across its UNI-D ports. For example, an NTU-B may support the delivery of 2x1000Mbps services, each delivered on a different UNI-D.

This device is capable of being operated in either a PON-based or Point-to-Point Ethernet-Based environment.

Physical details of the NTU-B are still being considered.

12.2.2.1 Physical Interfaces

The NTU-B supports four electrical or optical Ethernet UNI-D ports. An Access Seeker must specify the type and rate of the physical interface at time of ordering. The physical interface for each UNI-D may be independently configured.

The NTU-B does not support integrated UNI-V interfaces.

12.2.2.2 Power Supply

The NTU-B requires a local power source at the End-User premises. The NTU-B may be supplied with an AC or DC power supply.

Appendix A – Relevant Documents

Table 24 Standards

Standards Body	Reference	Specification
Metro Ethernet Forum	1	Technical Specification MEF 6.1, “Ethernet Services Definitions - Phase 2”, April 2008
	2	Technical Specification MEF 10.2, “Ethernet Services Attributes Phase 2”, October 2009
	3	Technical Specification MEF 15, “Requirements for Management of Metro Ethernet Phase 1 Network Elements”, November 2005
	4	Technical Specification MEF 17, “Service OAM Requirements & Framework – Phase 1”, April 2007
	5	Technical Specification MEF23, “Carrier Ethernet Class of Service – Phase 2”, October 2010
	6	Technical Specification MEF 26, “External Network Network Interface (ENNI) – Phase 1”, January 2010
IEEE	7	IEEE802.3-2008
	8	IEEE802.1ad
	9	IEEE802.1ag
	10	IEEE802.3ah
IETF	11	RFC2474 – “Definition of the Differentiated Services Field (DS Field) in the IPv4 and IPv6 Headers”
	12	RFC2598 – “An Expedited Forwarding PHB”
	13	RFC2597 – “Assured Forwarding PHB Group”
	28	RFC4594 – “Configuration Guidelines for DiffServ Service Classes”
	14	RFC2833 – “RTP Payload for DTMF Digits, Telephony Tones and Telephony Signals”
	15	RFC3261 – “ SIP: Session Initiation Protocol”
	16	RFC3435 – “Media Gateway Control Protocol (MGCP) Version 1.0”

ITU-T	17	Y.1731 – “OAM functions and mechanisms for Ethernet based networks”
	18	G.984 – “Gigabit Passive Optical Network”
	19	G.992 – “Asymmetric Digital End-user Line”
	20	G.993 – “Very High Bit-Rate Digital End-user Line, Version 2”
Broadband Forum	21	TR-069 – “CPE WAN Management Protocol v1.1”, December 2007
	22	TR-101 – “Migration to Ethernet-based DSL Aggregation”, April 2006
	23	TR-144 - Broadband Multi-Service Architecture & Framework Requirements”, August 2007
	24	TR-156 – “Using GPON Access in the context of TR-101”, December 2008
Comms Alliance Ltd.	25	G632:2007 Quality of Service parameters for networks using the Internet Protocol Industry Guideline
	26	National Broadband Network Wholesale Service Definition Framework – Ethernet, December 2009
NBN Co	27	Product Overview - Fibre Access Services, December 2010
	28	Consultation Paper – Access Seeker Accreditation (Fibre Network), December 2010
	29	Consultation Paper – Connecting to the National Broadband Network (Fibre Network), November 2010

Appendix B – Effective Information Rate

The relationship between expected service information rate and Service Frame size is in accordance with RFC2544 Appendix B: Maximum Frame Rates Reference. Note figures quoted in RFC2544 Appendix A are for an Ethernet UNI operating at 10Mbps. These figures may be scaled for UNI rates of 100Mbps and 1Gbps.

The calculation of Maximum Effective Layer 2 Throughput takes into account the service frame boundaries as defined in Section 6.3.1.

Calculations that accommodate a maximum layer 2 Ethernet frame size of 2000 Bytes, on a 100Mbps interface are provided in Table 25. Note these are theoretical figures that indicate the throughput limitations for varying frame sizes under ideal conditions, due to the impact of layer 2 overhead only.

Table 25 Effective Service Throughput as a Function of Frame Size

Frame Size	Line Rate	Maximum Effective Layer 2 Throughput (100Mbps Service)	Frames per Second
64 bytes	100 Mbps	76.19 Mbps	148,810
128 bytes	100 Mbps	86.49 Mbps	84,459
986 bytes	100 Mbps	98.01 Mbps	12,425
1518 bytes	100 Mbps	98.70 Mb/s	8,127
2000 bytes	100Mbps	99.01Mbps	6,188

Access Seekers must be aware of these service throughput limitations, as well as other limitations inherent in the Ethernet protocol, when dimensioning Access VC and Connectivity VC capacity.

Appendix C – Class of Service Application

The following table is reproduced from RFC4594 (figure 3, “DSCP to Service Class Mapping”). It indicates a broad range of application categories, and their respective DSCP associations. The right hand column titled “NFAS Traffic Class” indicates the Traffic Class that each DSCP category maps into. This table presents a broad guideline as to the usage of each NFAS Traffic Class for IP-level QoS applications.

Table 26 RFC4594 DSCP to Service Class Mapping

Service Class Name	DSCP Name	DSCP Value	Application Examples	NFAS Traffic Class
Network Control	CS6	110000	Network routing	N/A (Discarded)
Telephony	EF	101110	IP Telephony bearer	TC_1
Signalling	CS5	101000	IP Telephony signalling	
Multimedia Conferencing	AF41, AF42, AF43	100010, 100100, 100110	H.323/V2 video conferencing (adaptive)	TC_2
Real-Time Interactive	CS4	100000	Video conferencing and interactive gaming	
Multimedia Streaming	AF31, AF32, AF33	011010, 011100, 011110	Streaming video and audio on demand	TC_3
Broadcast Video	CS3	011000	Broadcast TV & live events	
Low-Latency Data	AF21, AF22, AF23	010010, 010100, 010110	Client/server transactions Web-based ordering	
OAM	CS2	010000	OAM&P	
High-Throughput Data	AF11, AF12, AF13	001010, 001100, 001110	Store and forward applications	TC_4
Standard	DF (CS0)	000000	Undifferentiated applications	
Low-Priority Data	CS1	001000	Any flow that has no BW assurance	