GNA – Global Network Architecture
Web: http://gna-re.net/

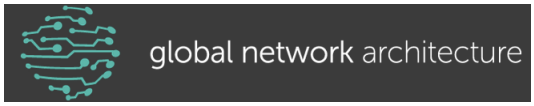

Document name:	Service Description: Multi-Layer Transport Service for the Global Network Architecture
Author(s):	Andrew Lee, David Wilde, Gerben van Malenstein, Jeronimo Bezerra, Joe Metzger and Migiel de Vos
Contributor(s):	GNA Technical Group
Date:	January 2017
Version:	v1.0

Service Description: Multi-Layer Transport Service for the Global Network Architecture

Table of Contents

Executive Summary

This document describes the attributes and characteristics of various data network transport services that are relevant in the context of the GNA. Not all transport services need to be supported in order for network infrastructure to be GNA Compliant. Not all transport services listed need to be offered, nor is this an exhaustive list of potential services.

Introduction

In the context of the GNA, a multi-layer transport service is the service of moving data between Endpoints, independent of location, application, or network layer. This document describes some classes and sub-classes of Transport Services and attributes to define different transport use cases. This document will be updated with new Transport Services and attributes as needed to reflect the research and education network needs.

General Transport Service Attributes

Disclaimer: *This document uses the terms Endpoint, Edge Domain, Transport Domain, Segment and Transport Service as well as the definition of Point-to-Point and Multi-Point provided by the GNA Glossary document. Readers are encouraged to get familiar with these concepts before reading this document. The Appendix provides more details of both Point-to-Point and Multi-Point services.*

All Transport Services must have the following attributes defined. These attributes will be inherited by the sub-class, described in following sections.

Attribute	Options
List of Endpoints	- Two or more Endpoints to be part of the service
Topology	- Point-to-Point - Multi-Point
Scheduled services	- Start and finish time - Duration
Endpoint throughput guarantees	- Best effort - Guaranteed bandwidth

Specific Sub-class Service Attributes

A given Transport Service can belong to one of the following sub-classes:

- Layer 1 or Wavelength transport
- Layer 2 or Ethernet transport
- Layer 3 or IPv4/IPv6 transport
    - Private Layer 3 (IPv4/IPv6 VPN)
    - Academic Layer 3 access
- Any Layer or Layer Independent Transport
    - Transport based in a combination of multiple layers

The subsequent sections describe the attributes specific to each sub-class.

Layer 1 or Wavelength Transport Service

Layer 1 or Wavelength Transport Service is based on the optical layer. Edge Domains can interface with Endpoints on the OTN, DWDM, CWDM or fiber level. Transport Domains could use optical or switching technologies to transport traffic between Edge Domains. Layer 1 Transport Services have the following attributes:

Attribute	Options
Endpoint PHY & throughput	- Wavelength - OTN - Ethernet/Fibre Channel/RDMA/SDH

Layer 2 or Ethernet Transport Service

Layer 2 or Ethernet Transport Service is the Transport Service based on Ethernet headers, such as destination MAC address, VLAN, and PCP fields. Layer 2 Transport Services have the following attributes:

Attribute	Options
Protocol/Ethertype support	- Ethernet /Any Ethertype
Endpoint PHY & throughput	- Ethernet
Endpoint MTU	- MTU supported: at least 9018 Bytes
Throughput control	- Rate-limiting - Shaping
Filtering	- MAC address and Ethertype ACL
Service multiplexing (if service is based on port or attribute)	- Port-based Endpoint - VLAN-based Endpoint
Protocol encapsulation (assumes port-based Endpoint)	- Q-in-Q (802.1q tunneling) - Layer-2-Protocol-Tunneling (L2PT) - MACSec tunneling
MAC address learning (assumes Multi-Point topology)	- Maximum number of MAC addresses allowed
Class of Service	- CoS passed transparently - CoS acted upon
VLAN translation	- Support of VLAN translation

Layer 3 or IPv4/IPv6 Transport

Layer 3 or IPv4/IPv6 Transport Service is the Transport Service based on Layer 3 headers, such as destination IPv4 and IPv6 addresses. Layer 3 Transport Services have the following sub-classes:

Academic Layer 3 Access – IPv4/IPv6 Access to Academic Sites

Academic Layer 3 access is a general Transport Service that connects all Research and Education Networks (REN), using both IPv4 and IPv6 protocols, being the simplest way any REN could reach other RENs in the world. Academic Layer 3 access has the following attributes:

Attribute	Options
QoS	- ToS/DSCP passed transparently - ToS/DSCP acted upon
Endpoint IP MTU	- IP MTU supported: 9000 Bytes
Routed protocol support	- IPv4/IPv6 protocols supported
Routing protocol support	- Static - MP-BGP
Filtering	- ACL (per interface ACL or prefix-list)
Maximum Number of Routes	- Maximum number of routes allowed

Layer 3 or IP VPN Service

Layer 3 or IP VPN service is the Transport Service that creates a private IPv4/IPv6 routing environment with a selected number of Endpoints and number of prefixes. Layer 3 or IP VPN Transport Services have the following attributes:

Attribute	Options
QoS	- ToS/DSCP passed transparently - ToS/DSCP acted upon
Endpoint IP MTU	- IP MTU supported: 9000 Bytes
Routed protocol support	- IPv4/IPv6 protocols supported
Routing protocol support	- Static, RIPv2, OSPF/OSPFv3 or MP-BGP
Maximum Number of Routes	- Maximum number of routes allowed
Filtering	- ACL (per interface ACL or prefix-list)

Any Layer or Layer Independent Transport

Any Layer or Layer Independent Transport is the Transport Service that uses the flexibility provided by the Software-Defined Networking (SDN) approach, available at some RENs and Software-Defined Exchange Points (SDX). With the Any Layer Transport, Edge and Transport Domains will transport traffic matching protocol fields from different protocol layers, for instance, destination MAC address plus destination TCP port. The Any Layer Transport Service has the following attributes:

Attribute	Options
Protocol support	- Ethernet
Protocols/ports supported	- Combination of attributes from Layers 2 to 4 to characterize the Service
Endpoint IP MTU	- Maximum MTU supported: 9000 Bytes
Filtering	- ACL/Prefix List

Example Use-cases

In this section, examples per use-case will be provided.

Use-case #1: Layer 1 Point-to-Point

To enable Layer 1 Point-to-Point service, Endpoints and Edge Domains have to define one approach:

- Alien Wave: user connects directly to the Edge Domain's MUX/DEMUX
- Capacity: user connects to an assigned Edge Domain's transponder

If decided for *Alien Wave*, a lambda/channel will be assigned to user, as well interface characteristics (optical fiber type and connector type). For example, *Channel 42, optical fiber's connector LC/PC duplex, Single Mode in both sides.*

If decided for *Network Capacity*, the Edge Domain will assign a terminal transponder accordingly to users needs.

Example: *Fibre Channel or Ethernet, 100 Gbps, Multi-Mode or Single Mode interface and optical fiber connectors*.

Use-case #2a: Layer 2 Point-to-Point

For Layer 2 Point-to-Point transport, the Ethernet protocol is available. When requesting this service, user has to provide the following attributes:

- If the Service is Port-based or VLAN-based;
    - If the Service is VLAN-based, VLAN IDs for both ends
- Identify the Endpoints;
- Bandwidth requirements;
- Number of MAC addresses expected;
- If Class of Service (CoS) will be used;
- If specific Layer 2 filters (Ethertype, for instance) are required.

If bandwidth reservation is required by user and supported by the Edge and Transport Domains, a CoS value should be selected to apply any QoS policy (shaping, rate-limit, or bandwidth guaranteed).

Example A: *User requests VLAN ID 450 from site A to site B, with minimum bandwidth of 100 Mbps. Only 10 MAC addresses are expected in this VLAN*.

Example B: *User requests a port-based service, where the Edge Domain will encapsulate all traffic. Traffic should be limited to 450 Mbps.*

Use-case #2b: Layer 2 Multi-Point

For Layer 2 Multi-Point transport, only the Ethernet protocol is available. When requesting this service, user has to provide the following attributes:

- VLAN IDs to be used at each site;
- Who the remote Endpoints are;
- Bandwidth requirements;
- Number of MAC addresses expected;
- If Class of Service (CoS) will be used.

Example: *For Super Computing, a demonstration requires three Endpoints with the same VLAN ID: 1799. Minimum bandwidth of 100Mbps is being requested. Only 20 MAC addresses are expected in this Service*.

Use-case #3a: Academic Layer 3 Access

In the academic community, providing IPv4/IPv6 access to academic networks is the most basic service. Through this service, Endpoints can exchange data through a dedicated IPv4/IPv6 network, with more bandwidth and minimum effort.

When requesting Academic Layer 3 access, requester should provide the following data:

- An IPv4/IPv6 prefix to establish the BGP session
- The Endpoint Autonomous System Number (ASN)
    - When the Endpoint does not have a ASN, a private ASN could be used
- If ToS/DSCP should be acted upon;
- The routing protocol between Endpoint and Edge Domain: static or MP-BGP
- Maximum number of routes the Endpoint plans to advertise;
- Prefixes to be advertised.

Example: ***A 10 Gbit/s connection to a new research institute that was recently created, supplying IPv4 and IPv6 connectivity to the Internet, by an R&E Network.***

Use-case #3b: Private Layer 3 VPN

Endpoints can request a private Layer 3 VPN to create a separated private Layer 3 environment for a specific purpose. When requesting a Layer 3 VPN, requester should provide the following data:

- List of Endpoints;
- If ToS/DSCP should be acted upon;
- Which IP version will be used: 4, 6 or both;
- Routing protocol: Static, MP-BGP, RIP or OSPF (not all Provider Edges will support all protocols);
- Maximum number of routes expected;
- Prefixes to be used.

Example: *the LHCONE L3VPN, connecting participants interested in access LHCONE data. The routing protocol is BGP, where IPv4 and IPv6 are in use*.

Use-case #4: Any Layer or Layer Independent Transport

In specific use cases, forwarding based on destination MAC address or IP address is not enough. Interfaces such as OpenFlow and P4 provide the flexibility of matching multiple frame's fields for traffic forwarding. Endpoints might request traffic forwarding based on specific transport protocol and ports, not just based on IP address. When supported, the Edge Domain could offer different approaches for traffic forward, and requester should provide the following data:

- List of Endpoints;
- Matching fields and wildcards;
- Bandwidth requirements.

The list of Endpoints will be used to identify if all Edges and Transport Domains will be able to support such forwarding mechanism.

Example: *a researcher needs traffic IP RTP (source or destination) to be forwarded separated from academic all-purpose IP traffic. Matching will be based on UDP ports, with two Endpoints.*

Operational Information Requirements

Due to the complexity involved in multi-domain provisioning and operation, the following items are required for all services listed in this document:

a) Service ID for common identification of a service: Each Service provisioned should have a Service ID associated that could be used by Endpoints, Edge and Transport Domains to correlate to the physical infrastructure. *This document does not state if a single Service ID should be used by all networks or if each domain will assign its own Service ID*;

b) Provisioning of the service: it is not required that all participating Transport Domains support all technologies universally, including support for static and dynamic provisioning. A mechanism may be needed to exchange information regarding supported technologies and their attributes among Transport Domains so that the Service may be appropriately provisioned (or rejected). *The mechanism to exchange information between Transport Domains regarding supported technologies is outside of the scope of this document*;

c) Acceptance testing: before considering provisioned, Endpoints or Edge Domains should be able to verify end-to-end performance of Service, including throughput of each Endpoint and failover approached in place;

d) Troubleshooting: Mechanisms should be provided to allow Endpoints and domains to debug individual segments. Issues, such as poor performance, MAC learning limitations, packet drops or inconsistent latency/jitter, are examples of issues to debug;

e) Monitoring: Both Edge and Transport Domains should monitor the status of their networks and be able to correlate the physical network with the provided Transport Service. Additionally, domains should be able to notify Endpoints of network issues that impact the Transport Service provided;

f) Visibility: Each domain should provide mechanisms for sharing operational information of each service, such as throughput, packet drops, link state, errors and RTT;

g) Service Lifecycle: At any moment, users might request changes to the Transport Service or attributes, for example, addition/removal of an Endpoint, changes to the QoS profile associated with the Service, etc. *It is outside of the scope of this document to define how changes and interaction between users and domains should happen.*

Appendix

This appendix provides more details to the definitions of Point-to-Point and Multi-Point Transport Services described in the GNA Glossary document.

Point-to-Point Transport Service

A Point-to-Point service transports traffic between two Endpoints. Point-to-Point transportation is performed accordingly to the following approach:

1. All compliant traffic sent by the source Endpoint is encapsulated/forwarded by its Edge Domain;
2. The Edge Domain forwards traffic to its next domain, which could be another Edge or a Transport Domain;
3. Traffic is forwarded to the remote Edge Domain;
4. Remote Edge Domain receives the traffic from the Transport Domain and forwards to the remote Endpoint.

Figure 1 represents a multi-domain point-to-point service and represents:

- Only two Endpoints: **1** and **2**
- Two Edge Domains: **A** and **D**
- Two Transport Domains: **B** and **C**
- A point-to-point **Transport Service**
- Four **Segments**
- Three different examples of technologies (**a**, **b** and **c**).

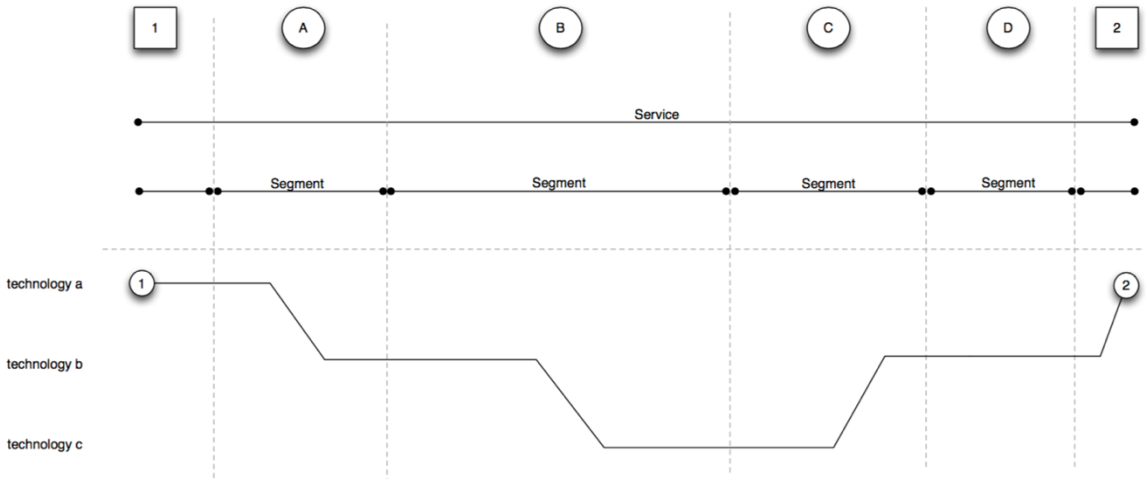

Figure 1. Representation of a Point-to-Point service

Multi-Point Service

The Multi-Point service transports traffic between multiple Endpoints. Multi-Point transportation is performed accordingly to the following approach:

1. All compliant traffic sent by the source Endpoint is encapsulated by its Edge Domain;
2. Edge Domain performs lookups for the destination Edge Domains;
3. Edge Domain forwards traffic to its next-hop domains;
4. Traffic is forwarded to remote Edge Domains that are addressed (can be one or more);
5. Remote Edge Domains decapsulate the traffic and forward to the remote Endpoints

Figures 2 and 3 are provided to represent Multi-Point service.

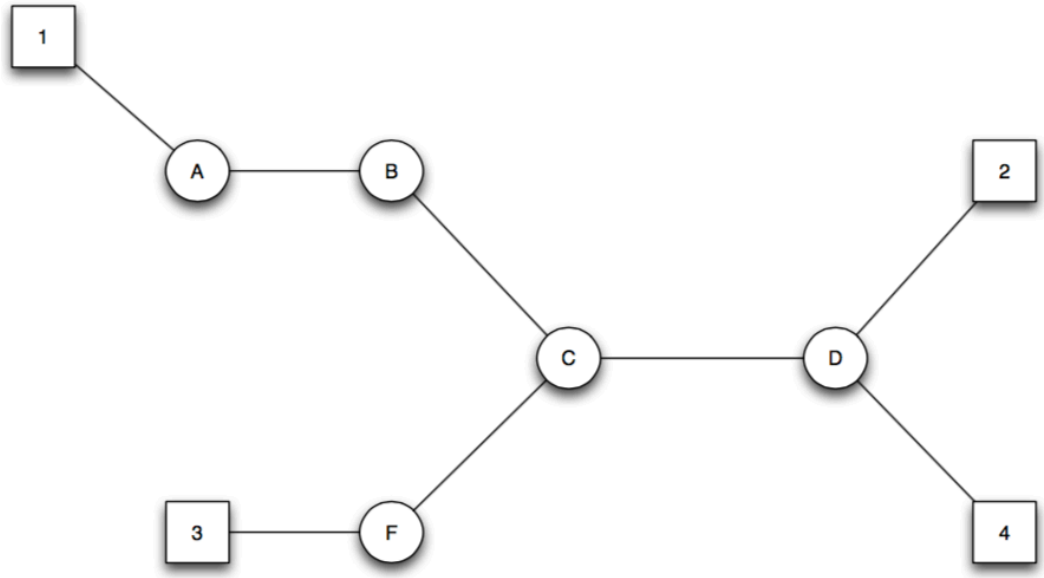

Figure 2. Multi-Point Service with four Endpoints

It is possible to see on Figure 2 Endpoints **1**, **2**, **3** and **4**, connected to Edge Domains **A**, **D** and **F** respectively. There is only one Transport Domain in the path without an Endpoint: **C**. Figure 3 represents the Transport Service, Segments and technologies.

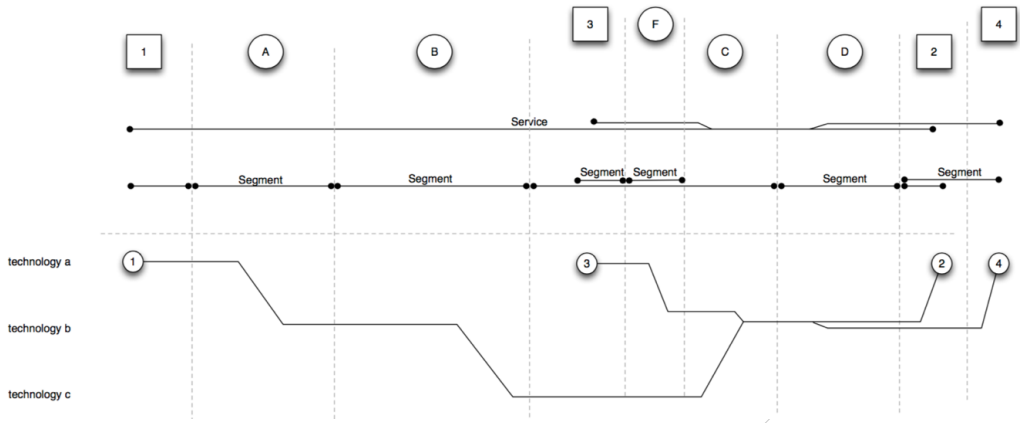

Figure 3. Multi-Point Service with four Endpoints Detailed