METRO-HAUL

METRO High bandwidth, 5G Application-aware optical network, with edge storage, compute and low latency

Grant No. 761727

Deliverable D2.3

Network Architecture Definition, Design Methods and Performance Evaluation

Editor:	António Eira, João Pedro (Coriant Portugal)
Deliverable nature:	Report (R)
Dissemination level: (Confidentiality)	Public (PU)
Contractual delivery date:	March 2019
Actual delivery date:	
Suggested readers:	Example: Executives in entertainment companies and banks, investors
Version:	1.00
Total number of pages:	119
Keywords:	

Abstract

The overall Metro-Haul objective is to architect and design cost-effective, energy-efficient, agile and programmable metro networks that are scalable for 5G access and future requirements, encompassing the design of all-optical metro nodes (including full compute and storage capabilities), which interface effectively with both 5G access and multi-Tbit/s elastic core networks.

This document provides the detailed candidate Metro-Haul architectures defined to comply with the functional architecture requirements in D2.2. Additionally, it describes the network design methods developed to plan and optimize metro networks based on these architectures, and provides early comparative performance evaluations of the proposed architectures.

Disclaimer

The research leading to these results has received funding from the European Union Horizon 2020 Programme under grant agreement number METRO-HAUL 761727

Impressum

[Full project title] METRO High bandwidth, 5G Application-aware optical network, with edge storage, compute and low latency

[Short project title] METRO-HAUL

[Number and title of work-package] WP2 - Use Cases, Service Requirements, and Network Architecture Definition

[Number and title of task] Task 2.3 – Network Architecture Design and Evaluation

[Document title] Network architecture definition, design methods and performance evaluation

[Editor: Name, company] João Pedro, António Eira, Coriant Portugal

[Work-package leader: Name, company] Albert Rafel, British Telecom

[Estimation of PM spent on the Deliverable]

Executive Summary

The overall target of the Metro-Haul project is the creation of a metro transport network technologically capable of supporting 5G services, by handling higher traffic volumes without scaling the fiber footprint accordingly, and addressing application-specific requirements through a multi-layer and multi-domain provisioning of network and compute/storage resources. Previously, a set of 5G service use cases were identified and characterized, leading to a functional definition of the Metro-Haul node requirements, for both the Access Metro Edge Nodes (AMENs), and the Metro Core Edge Nodes (MCENs). Simultaneously, other work packages (WPs) are developing the data and control plane solutions that constitute the building blocks for these nodes.

In this deliverable, these contributions are aggregated to define candidate Metro-Haul architectures, addressing various network scenarios. These architectures are to be evaluated in the scope of WP2's techno-economic analysis against reference state-of-the-art implementations. The ultimate objective of this evaluation is to demonstrate the network deployment benefits of Metro-Haul solutions with respect to key performance indicators (KPIs) regarding the capacity of the network, its cost, and its power consumption. The scope of this deliverable is to present the candidate architectures, based on the physical devices and control-plane solutions developed within the project, and provide the methodology for performing these techno-economic studies, based on network dimensioning and optimization capabilities also developed by project partners.

Given the broad scope of Metro-Haul, much of the work described here involves breaking down the end-to-end network architecture into various subdomains that can be analysed independently, and combined to form solutions addressing specific scenarios. The deliverable details the project KPIs, and specifically those that relate to techno-economic studies: network capacity, cost reduction and power consumption. Then, it defines the necessary support frameworks for performance evaluations. These include the definition of geotypes characterizing the expected coverage of AMENs/MCENs, the modelling of client traffic at a macro-level (service types, growth rates, etc.), and at the scale of hours/days/weeks (modelling tidal variations and aggregation of dynamic service flows). Additionally, the development of the open source Net2Plan planning tool to natively support scenarios of interest to Metro-Haul (such as optimized placement of virtual network functions (VNFs) over optical transport networks) is detailed, specifying the interfaces that can be used to model various dimensioning problems. Finally, the modelling of the costs and power consumption for optical, packet and data-center components is defined, leveraging where possible the data and knowledge of project partners with network deployment experience.

Afterwards, the work streams for the various architectural aspects, detailing candidate architecture and dimensioning methodologies, are presented. The first aspect is the possible architectures of the central offices being repurposed to integrate data-center capabilities, and how to best interconnect access traffic, optical transport and compute resources given the traffic requirements. Regarding the architecture of the optical transport node themselves, options are constructed for metro aggregation and metro-core segments, reusing several devices/subsystems developed in WP3. The economic impact of disaggregation at the optical layer is assessed by a specific cost model detailing potential benefits/hindrances of each option for network deployment. Finally, the efficiency gains of supporting edge computing in co-localization with optical nodes is evaluated through several network optimization frameworks, that show how 5G services can be enabled without unduly increasing network/compute resources. This analysis also attempts to quantify how Metro-Haul's control, management and orchestration platform (along with the hardware devices' flexibility) can improve network efficiency, by providing the tools to automatically reconfigure network bandwidth and VNFs across the end-to-end infrastructure in response to changing service requirements.

List of Authors

	Name	Partner
	António Eira, João Pedro	Coriant Portugal
	Marc Ruiz, Luis Velasco, Jaume Comellas, Gabriel Junyent	UPC
	Ramon Casellas, Raül Muñoz, Laia Nadal, Michela Svaluto Moreolo	CTTC
	Marco Quagliotti, Anna Chiadò Piat	TIM
	Michael Parker	LEX
	Miquel Garrich, José Luis Romero-Gázquez, Francisco-Javier Moreno-Muro, and Pablo Pavón Mariño	UPCT
	Massimo Tornatore, Francesco Musumeci, Omran Ayoub, Leila Askari	PoliMi
	José Alberto Hernández, David Larrabeiti, Ignacion Martin, David Eugui	UC3M
	Rafael López da Silva, Lourdes Luque Cantó, Oscar González de Dios	Telefónica
	Albert Rafel	BT
	Teresa Bi, Abubakar Siddique Muqaddas, Arash Farhadi Beldachi	UBRIS
	Bogdan-Mihai Andrus, Annika Dochhan	ADVA
	Francesco Fresi, Filippo Cugini	CNIT
	Antonio D'Errico	TEI
Checked by:		

Revision History

Revision	Date	Responsible	Comment
0.0	19-10-2018	Editor	Tentative ToC version

0.1	6-12-2018	Editor	Updated ToC for working streams analysis and tentative section leads.
0.2	6-02-2019	Editor	Updated ToC for KPI definition and contributors for each subsection
0.3	27-02-2019	Editor	Initial contributions integrated
0.4	01-03-2019	Editor	Further contributions aggregated.
0.5	08-03-2019	Editor	Initial M-H component descriptions integrated.
0.6	13-03-2019	Editor	Disaggregation options and methodology, and physical architectures for optical nodes included.
0.7	26-03-2019	Editor	Geotype definition, preliminary cost model components and further component descriptions added.
0.9	14-04-2019	Editor	Introduction, conclusion and executive summary added. Final contributions integrated. Initial references and list of acronyms added.
1.0	29-04-2019	Editor	Final contributions aggregated for review.
1.00	06-05-2019	M. Parker	First internal review
1.01	07-05-2019	D. King	Second internal review

Table of contents

List of Tables

List of Figures

1 Introduction

The Metro-Haul project aims to provide technological and architectural solutions to enable cost-effective, scalable and sustainable metro transport deployments capable of supporting 5G service requirements. Within the project's structure, work package (WP) WP2 is responsible for the definition of target use cases and service requirements to be supported, the definition of the functional network architecture, and finally the design/dimensioning of the network and corresponding performance evaluation. The use-case definition and functional requirements were targeted in the previous deliverables, [D2.1] and [D2.2], respectively. This deliverable D2.3 addresses the candidate network architectures are complying with the functional requirement definition and supporting the identified use-cases. Furthermore, it provides the end-to-end methodology used to assess the said candidate architectures, along with the key performance indicators (KPIs) defined at the global network level. The actual performance evaluation of these candidate architectures with respect to the KPIs, along with any refinements to the architectures themselves, will be detailed in the subsequent deliverable D2.4.

1.1 Metro-Haul Network Architecture

The geographic and hierarchical scope of the network segments targeted by Metro-Haul are shown in Figure 1. The Metro-Haul scope is bounded between heterogeneous fixed/wireless access clients and the photonic core network. On the access side, Metro-Haul nodes expand the traditional role of central offices (COs) in aggregating user/enterprise traffic, by also integrating such functions with distributed storage/compute capabilities. These expanded nodes are denoted as Access Metro Edge Nodes (AMENs), which are expected to provide an application-aware transport layer for service use-cases with specific bandwidth/IT requirements. On the interface with the core, Metro-Haul nodes are expected to assume a different configuration, being designed for higher optical throughput, and focusing their storage/compute requirements for functions with a slightly different profile, typically concentrated around caching or more centralized control/processing tasks. These nodes are termed Metro Core Edge Nodes (MCENs). The specific functional requirements for AMENs and MCENs, such as the type of service use-cases and virtual network function (VNF) profiles most suited to each of them, are presented in detail in [D2.2].

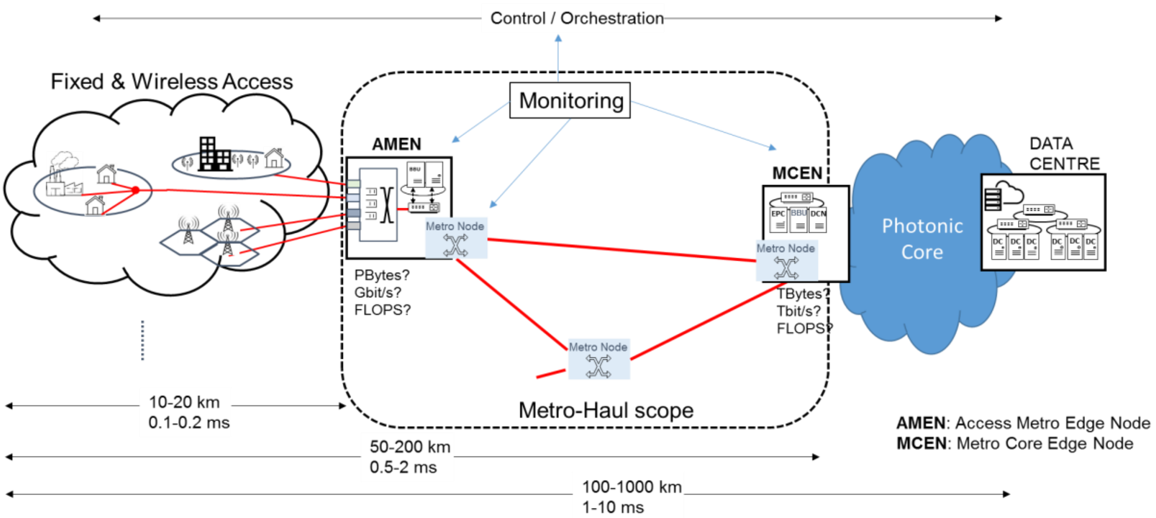

Figure 1: Metro-Haul Network Scope.

The physical structure of AMENs/MCENs can be functionally divided into a subset of different modules, as illustrated in Figure 2. It comprises an optical node component, which can be further separated into the line system and the transceivers. The former comprises the optical switching elements feeding into the wavelength division multiplexing (WDM) fiber transport network, such as reconfigurable optical add/drop multiplexers (ROADMs), filters, etc., as well as components supporting fiber transmission, such as amplifiers. The transceivers include the devices creating the optical signals that are transmitted over the fiber, based, e.g. on direct-detection or coherent transmission. The top half of Figure 2 illustrates the two other main components of an AMEN/MCEN: the storage/compute node hosting VNFs for specific services, and a layer 2/3 switching fabric that aggregates heterogeneous access traffic sources (mobile/fixed, residential and enterprise, specific vertical use-cases). This latter module also has the critical role of seamlessly interconnecting the access and optical transport networks with the compute node, thus enabling an orchestrated distribution of compute resources across a multi-layer and multi-domain infrastructure. This high-level segmentation of an optical node (valid for both AMENs and MCENs, albeit with different requirements) is also beneficial towards evaluating different candidate architectures for each scenario, as it enables performance evaluations over multiple possible configurations for the different building blocks.

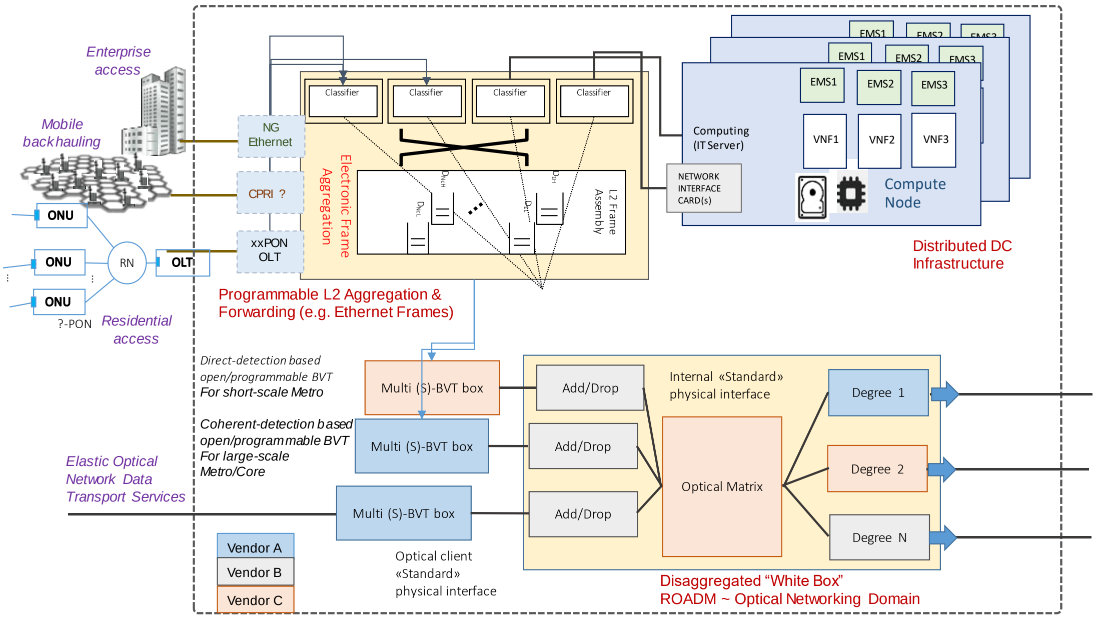

Figure 2: Metro-Haul Node Architecture.

Building upon the new device technologies developed towards building AMENs and MCENs, Metro-Haul is also highly focused on integrating these nodes into a single orchestration domain. Specifically, Metro-Haul aims to automate the optimized deployment of network services across multiple layers and multiple domains. This entails integration with the control plane that enables services to be provisioned across packet/optical networks, and that also extends to the compute resources at distributed ETSI NFVI PoPs. The scope of these orchestration capabilities is outlined in Figure 3. They are mostly envisioned towards integrating packet and optical network control with distributed VIMs at remote compute nodes, under the umbrella of the MANO layer. Ultimately, this architecture leads to the support of network slicing across the infrastructure, wherein a slice consists of a set of (possibly distributed) VNFs and the associated logical/physical network resources (across packet/optical layers and domains).

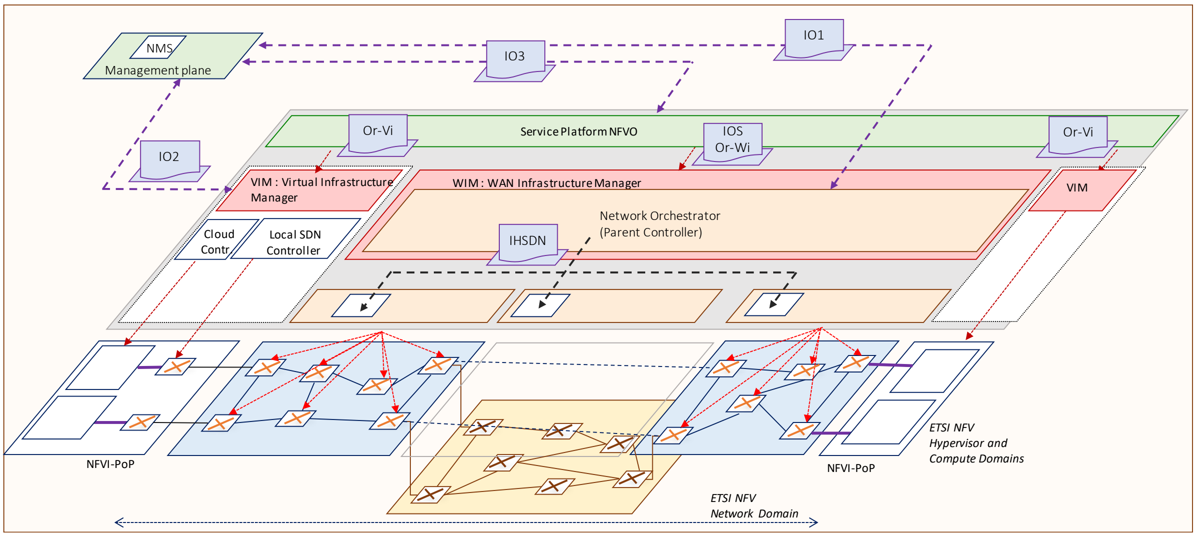

Figure 3: Metro-Haul Service Platform.

The last pillar of the Metro-Haul network vision is the support of open and disaggregated solutions that facilitate network and service deployment. This objective ties both the data- and control-plane solutions together, requiring devices or subsystems to easily interoperate by adhering to common models (e.g. based on NETCONF/YANG). This enables device-specific details to be abstracted away from the network-level software defined network (SDN) control tasks, and also facilitates the deployment of third-party applications interfacing at the device or network level, such as monitoring and data analytics frameworks pulling data directly from devices/subsystems and interfacing with the control plane for improved network operation.

1.2 Deliverable Structure

The main target of this deliverable is to identify candidate network architectures meeting the previously identified functional requirements, and outlining a methodology for the evaluation of such architectures, leading to the analysis of whether the project-defined KPIs are met. With this in mind, D2.3 is organized around evaluation streams regarding concrete aspects of the overall Metro-Haul architecture. This breakdown is necessary given the scope of Metro-Haul, which does not propose a one-size-fits-all architecture, but rather a technology toolbox that can be used selectively for different application scenarios. In the scope of Metro-Haul, these scenarios relate mostly to the geographic/demographic setting of the nodes/network (e.g., rural/urban), which influences the traffic volume/pattern, and also, in the specific case of the optical transport network, the distances to be covered by optical interfaces.

Additionally, the traffic itself requires modeling suited towards the 5G service use-cases identified in [D2.1] and [D2.2]. Hence, Task 2.3 of the project is organized around three pillars:

- **Project KPIs** to be measured in the scope of techno-economic network evaluations: network capacity, cost and power consumption.
- **Support tasks** that feed evaluation models with realistic scenario use-cases, and provide ease of access to data sets and dimensioning tools for results dissemination and replication.
- **Working streams** that focus on particular architectural aspects of the overall Metro-Haul ecosystem, and that can be combined to form a specific end-to-end (E2E) solution most suited to each scenario or use-case. The project KPIs are expected to be measured against the combined analysis done for each working stream analysis.

Being centered around architecture definition and methodology, D2.3 incorporates this tiered structure. Section 2 presents the project-wide KPIs in greater detail, paying particular attention to the ones that are expected to be measured in terms of a techno-economic network dimensioning (and not in the scope of, e.g., a demonstration). The breakdown of Task 2.3 concerning evaluating these KPIs is then detailed, identifying the specific support tasks required, as well as the working streams identified within the Metro-Haul environment. Section 3 is devoted to an in-depth definition for each support task, covering aspects such as geotype definition, traffic modeling, planning tool interworking, and cost or power consumption modeling. Finally, Section 4 delves into the working streams, logically breaking down the Metro-Haul architecture into subcomponents that can be studied individually in depth, and later combined to form an E2E architecture per reference scenario. The focus of the working streams in this deliverable is to: 1) identify suitable candidate architectures given the scope and use-cases of previous deliverables; and 2) define methodologies for evaluations of the said architectures, in terms of modeling frameworks and optimization algorithms for network dimensioning.

2 Metro-Haul KPI Definition and Evaluation

Within the scope of the 5G-PPP Programme Performance KPIs [5GPPP-KPI], Metro-Haul has defined a set of project-wide macro KPIs that map to one or more of these reference targets. This analysis was conducted across the various WPs, taking into account what new contributions Metro-Haul is providing, and which among those can be effectively measured and quantified. The goal of this section is to frame the Metro-Haul general KPIs within the context of 5G-PPP performance KPIs, elaborate on the subcomponents required to evaluate such KPIs, and establish a methodology for their evaluation (e.g. in the scope of a demonstration, techno-economic evaluation, etc.). Particular attention is given to techno-economic related KPIs, which form the backbone of the evaluation frameworks described in this deliverable.

2.1 Project KPIs

The 5G-PPP performance KPIs are listed in *Table 1*, providing a reference target scope for increasing capacity, lowering the energy footprint, enabling faster service setup and ensuring reliability/security for ubiquitous new service deployments.

Table 1: 5G-PPP Performance KPIs

	Performance KPIs
P1	Providing 1000 times higher wireless area capacity and more varied service capabilities compared to 2010.
P2	Saving up to 90% of energy per service provided.
P3	Reducing the average service creation time cycle from 90 hours to 90 minutes.
P4	Creating a secure, reliable and dependable Internet with a “zero perceived” downtime for services provision.
P5	Facilitating very dense deployments of wireless communication links to connect over 7 trillion wireless devices serving over 7 billion people.
P6	Enabling advanced user-controlled privacy.

In very broad terms, the scope of Metro-Haul is to provide a high capacity and cost-efficient optical transport infrastructure, able to cope with much higher traffic volumes, while integrating it with application-aware edge computing nodes to facilitate service deployment and optimize transport bandwidth requirements. Given these guidelines, nine project-wide macro KPIs were defined and mapped onto one or more 5G-PPP performance KPIs, as shown in Table 2.

Table 2: Metro-Haul general KPIs and mapping to 5G-PPP performance KPIs.

Metro-Haul KPIs	5G PPP Performance KPIs					
	P1	P2	P3	P4	P5	P6
MH1. Optical PtP connection set-up time			X	X	X	
MH2. Metro-Haul E2E PtP connection set-up time			X	X	X	
MH3. Set-up time of network service slice across Metro-Haul	X		X		X	
MH4. Capacity of Metro-Haul controller	X				X	
MH5. Fault/degradation detection time				X		
MH6. Capacity of Metro-Haul infrastructure	X			X	X	
MH7. New optical components/systems	X	X	X			
MH8. CapEx reduction	X				X	
MH9. Energy consumption		X				

The following is a more detailed description of each general KPI: what it measures, its overall relevance, and how it is intended to be measured/evaluated in Metro-Haul.

2.1.1 MH1. Optical PtP connection set-up time

Table 3: MH1 definition.

KPI	MH1. Optical PtP connection set-up time
Definition	SDN-based management framework enabling fast configuration time in the Optical Layer to set up or reconfigure services handling 5G applications. This KPI is composed of the following elements: - Control plane latency and optical node reconfiguration delay - Time required to instantiate a network connection through the optical layer
Context/Use case	Each demo will define the specific context and Use Case - The hardware is already in place as well as the links (different connection scenarios are possible) - All the hardware have the right control plane interfaces and everything is configured in the Orchestrator
Where to measure	In the 3 major project demos involving the Optical Layer [TIM based demo; UK based demo; Berlin based demo]
How to measure	Traces from the SDN controller, time stamps and protocol analysers.
How to evaluate	Target: ≤ 1 min, describing the detailed connection scenario where this is achieved

2.1.2 MH2. Metro-Haul E2E PtP connection set-up time

Table 4: MH2 definition.

KPI	MH2. Metro-Haul E2E PtP connection set-up time
Definition	Connection set-up time of a multilayer metro connection between two MH nodes, using the SDN-based management framework. It includes Packet over an Optical PtP Connection. The time between when the connection request arrives at the parent controller, to the time the connection is up and running, and a positive answer is returned by the parent SDN controller to the request.
Context/Use case	The SDN controllers and hardware elements are already in place.
Where to measure	Final demos of the Project (UK & Berlin)
How to measure	Parent Controller (end-to-end request and overall control)
How to evaluate	Target: ≤ 2 min

2.1.3 Set-up time of network service slice across Metro-Haul

Table 5: MH3 definition.

KPI	MH2. Set-up time of a network service slice across Metro-Haul
Definition	Instantiation time of a network slice, taking the form of a network service through an orchestrator (Open Source MANO). The slice involves a given number of VNFs, and multilayer metro connections between them in the MH network. This is the time between the moment the slice is requested until all the VNFs and connections are fully configured, and usable for application traffic. In this respect, it includes the set-up times of the VNFs, service and optical connections.
Context/Use case	This KPI is composed of the following elements: - Completion time of the resource allocation algorithms in the planning module - Control plane latency, of the interactions among orchestrator (OSM), server farms (OpenStack clusters), SDN controllers.

	This KPI also aggregates the set-up times mentioned in MH1 and MH2.
Where to measure	Demo in Bristol, demo in Berlin, control plane demo
How to measure	From the time the request has arrived to the planning tool until the full slice is up and running
How to evaluate	Target: ≤ 1hr (depending on the #VNFs, #metro connections to automatically configure)

2.1.4 Capacity of Metro-Haul controller

Table 6: MH4 definition.

KPI	MH4. Capacity of Metro-Haul Controller
Definition	The KPI defines the number of supported optical devices controlled by a single SDN Controller instance, which is related to the number of NETCONF sessions that can be managed by such a controller. This KPI is directly related to the processing capacity of the controller. It is assumed that the number of managed devices has a direct impact on relevant metrics such as the latency in configuring operations across the network, or the control plane overhead associated with managing such numbers of devices.
Context/Use case	The following assumptions are part of the context in which the KPI is assessed: - A single controller ONOS instance is running in a Linux Server with medium to high level hardware (i7, 32 Gb RAM, …); - Virtual Containers are deployed emulating hardware devices; - Devices will model OpenROADM nodes and AMEN/MCEN nodes with the same conditions as with real hardware; - The SDN Controller uses a NETCONF session over SSH transport towards the devices; - A dedicated network supports communication between the Controller and the devices, with a minimum of 10 or 100 Gbps Ethernet. The KPI applies across multiple use cases.
Where to measure	This KPI will be measured in a Control Plane testbed: - A server running the ONOS controller; - Multiple servers supporting instantiated containers.

How to measure	Control Plane only, single controller running increasing numbers of virtual containers, carried out in the lab developing the network controller.
How to evaluate	Target: control of 10 – 100 nodes (AMENs/MCENs, i.e. basically Open Disaggregated ROADMs)

2.1.5 Fault/degradation detection time

Table 7: MH5 definition.

KPI	Fault/degradation detection time
Definition	Time from when a fault/degradation actually happens (e.g., some threshold is violated or unexpected trend is confirmed) until it is detected.
Context/Use case	Quality-of-Transmission (QoT) of optical connections
Where to measure	At an agent node or the network controller
How to measure	Monitoring data from network and other devices that are analysed
How to evaluate	Target: 5 consecutive monitoring/telemetry measurements

2.1.6 Capacity of Metro-Haul infrastructure

Table 8: MH6 definition.

KPI	Capacity of Metro-Haul infrastructure
Definition	Aggregate amount [Tb/s] of service instances (background traffic + 5G-specific) that can be supported by a Metro-Haul domain (i.e. set of AMENs and MCENs interconnected by optical fibre).
Context/Use case	Support an increasing amount of traffic with application-specific requirements reusing the existing fibre infrastructure.
Where to measure	Capacity of the components being developed within Metro-Haul (WP3) Algorithms for resource planning/provisioning in joint IT/Optical service deployments (WP2/WP4)
How to measure	Techno-economic network dimensioning Scale input traffic for each geo-type (e.g. 2019-2025, different traffic CAGR assumptions)

How to evaluate	100x more capacity for 5G-specific services over the same fibre infrastructure

2.1.7 New optical components/systems

Table 9: MH7 definition.

KPI	New optical components/systems
Definition	Per-component KPIs applying to the optical node that advance the state-of-the-art in one or more of the following aspects: - Transmission capacity of an optical interface [Gb/s] - Transparent reach of an optical interface [km] - Cost of an optical component/system [cost units] - Power consumption of an optical component/system [W] Footprint of an optical component/system [RU]
Context/Use case	Developing hardware that meets the specific needs of 5G-enabled Metro networks, by advancing the state-of-the-art in aspects such as cost, power consumption, physical footprint, and capacity.
Where to measure	Theoretical modelling support, and validation in experimental dedicated testbeds, and within the project demos.
How to measure	Successful transmission within the identified parameter ranges (e.g. capacity/reach). Prototypes meeting the required specifications.
How to evaluate	Specific capacity/reach/cost/footprint targets per component/system (per component details to be provided in Deliverable 3.2). Input modelling data for techno-economic analysis (see MH.8 and MH.9)

2.1.8 CapEx reduction

Table 10: MH8 definition.

KPI	CapEx reduction
Definition	Relative investment cost for deploying a joint IT/optical infrastructure across a Metro-Haul domain (metro aggregation chain, metro-core, or combination of a metro-core with several metro aggregation chains), compared to a reference baseline architecture.

Context/Use case	Support an infrastructure owner in deploying a cost-effective 5G-enabled network that supports competitive access to bandwidth and deployment of vertical services.
Where to measure	Per-Component cost savings (WP3) Architectural cost savings (WP2)
How to measure	Techno-economic network dimensioning Scale input traffic for each geo-type (e.g. 2019-2025, different traffic CAGR assumptions) For the different traffic/network profiles, dimension the overall network cost vs. a reference baseline implementation, subdivided across different network layers/hierarchies
How to evaluate	Aggregate CapEx reduction target of 50% (in optical node, CO/DC hardware and software) vs. a baseline metro network implementation to support the same set of services

2.1.9 Energy Consumption

Table 11: MH9 definition.

KPI	Energy consumption
Definition	Relative power consumption requirements of deploying a joint IT/optical infrastructure across a Metro-Haul domain (metro aggregation chain, metro-core, or combination of a metro-core with several metro aggregation chains), compared to a reference baseline architecture.
Context/Use case	Reduce the overall footprint for establishing network services, lowering operational costs and ensuring sustainable development.
Where to measure	Per-component consumption savings (WP3) Architectural power consumption savings (WP2)
How to measure	Techno-economic network dimensioning. Scale input traffic for each geo-type (e.g. 2019-2025, different traffic CAGR assumptions). For the different traffic/network profiles, dimension the overall network cost vs. a reference baseline implementation, subdivided across different network layers/hierarchies.
How to evaluate	Aggregate power consumption reduction of 50% (in the optical node and DC) vs a baseline metro network implementation to support the same set of services.

2.2 KPIs directly evaluated in the scope of WP2/T2.3

The project KPIs presented in the previous subsection can be loosely divided into the main methods for measuring them. Some of the KPIs are naturally defined for experimental measurement, mostly in the scope of the project's demonstrations. This is the case for connection/service setup time and fault degradation KPIs (MH1, MH2, MH3 and MH5). Other indicators, such as the capacity of the Metro-Haul controller (MH4) or new optical components/systems (MH7), refer to properties of specific Metro-Haul devices/components, and how they relate to the overarching project goals (e.g. how many devices can a single controller instance support). These KPIs are mostly measured through dedicated experimental and emulation work on specific components, to ascertain their performance (i.e., not necessarily integrated within a broader E2E system demonstration). Finally, the remaining KPIs (MH6, MH8 and MH9) are related to techno-economic evaluations on the various architectural aspects for Metro-Haul solutions. In the scope of WP2 and Task 2.3, these are the most relevant KPIs to consider, as they intend to provide a global view regarding the suitability of the Metro-Haul ecosystem to aspects such as capacity, cost and power consumption. The unifying goal of the architectures and evaluations described in this deliverable is thus to contribute towards evaluating these three KPIs. As such, it is worth delving into greater detail on what constitutes each of them.

2.2.1 MH6 – Capacity of the Metro-Haul Infrastructure

This KPI intends to demonstrate how the technologies and architectures envisioned in Metro-Haul contribute to achieving a much higher throughput in the metro optical transport network, while basically reusing existing infrastructures (i.e., without simply deploying more fiber capacity).

Since the breadth of possible applications in Metro-Haul is tremendous, it is necessary to more clearly define the scenarios where the KPI applies. Firstly, referring to capacity, in the scope of Metro-Haul this implies the optical throughput, storage/compute capabilities at AMEN/MCENs, and the switching capacity needed at the nodes to interconnect transport and data-center functionalities with mobile/fixed access traffic. Thus, since different applications have different requirements according to these metrics, capacity is defined as the aggregate amount of traffic offered by service instances supported in a Metro-Haul infrastructure. This traffic per service can be mapped to required throughput, storage/compute and switching requirements. The concrete mix of service use-cases considered for the evaluation varies according to the traffic modelling described in Section 3.2.

The capacity evaluations include a set of relevant scenarios that reflect different application areas of Metro-Haul solutions. Hence, on the devices/components side, capacity improvements are measured mainly through:

- Switching capacity for optical/IT nodes
    - Storage/compute capacity per AMEN/MCEN
    - Optical express and add/drop capacity per AMEN/MCEN
- Optical transmission capacity/reach per channel (direct-detection and coherent options)
- Aggregate single-fiber capacity supported across a metro aggregation chain and metro-core domain

Furthermore, much of the technological advance in Metro-Haul is predicated on the traffic shaping advantages provided by edge computing architectures. Hence, service throughput is also achieved by bandwidth offload at the edge, be it in the southbound direction (e.g., caching) or the northbound (e.g., data processing and compression at the edge).

2.2.2 MH8 – CapEx Reduction

This KPI evaluates to what extent the technologies and architectures proposed in Metro-Haul help reduce the investment in new network deployments supporting 5G services. The reduction is measured as the expected relative CapEx savings obtained by deploying a Metro-Haul solution, compared to a baseline reference architecture supporting the same set of services. Once again, given the wide scope of possible service use-cases and scenarios (e.g., related to deployments in different geotypes), this is not expected to be a one-shot comparison, but rather a qualitative aggregation of a quantitative analysis across several representative scenarios.

The CapEx savings are expected to be obtained from two main sources:

- Per-component cost savings, related to the use of innovative devices and technologies reducing cost and footprint (e.g., photonic integrated circuits (PICs) for transceivers and switching modules), the increase in capacity per component, and the flexibility in the components enabling them to be reused in various application settings (e.g., rate-adaptive transceivers).
- Architectural cost savings, arising from lower cost node architectures (e.g. filterless or semi-filterless nodes), integrated optical/data-center nodes that reduce switching/routing elements, and the CapEx impact of supporting disaggregated solutions.

As the following sections will highlight, this analysis will resort to a careful definition of the network scenarios to consider (in terms of geographic coverage, traffic profiles, reference cost models, etc.). In each case, several possible candidate architectures are identified (e.g., for CO architectures, optical node architectures, disaggregation options), with one or more about reference/baseline network implementations, and the remaining ones based on technological solutions championed by Metro-Haul. It should be noted that KPI MH7 feeds directly into this one, providing target costs for specific components developed within Metro-Haul.

2.2.3 MH9 – Energy Consumption

The structure of this KPI is fairly similar to MH8, with a special emphasis on the power consumption savings enabled by solutions based on Metro-Haul subsystems (control and data-plane based). The energy savings are to be evaluated in several representative network deployment scenarios, comparing baseline architectures with those proposed within the project.

The energy consumption savings in Metro-Haul solutions are expected to come mainly from:

- Per-component consumption savings, due to higher integration (e.g., PICs) or use of passive components. Higher transported capacity per consumed power.
- Architectural consumption savings, due to mainly-passive node architectures (e.g., filterless nodes), lower hardware requirements in converged CO / data-center nodes, and lower transport bandwidth requirements thanks to edge computing (e.g., caching or Internet of Things (IoT) processing at the edge).

Also similarly to MH8, the overall KPI target (50% total power consumption savings) will be qualitatively analysed as the aggregation of a quantitative techno-economic study for several network deployment scenarios, comparing baseline implementations with candidate architectures outlined throughout this deliverable.

2.3 Framework definition for techno-economic assessment of KPIs

In light of the broad scope covered by the Metro-Haul ecosystem, evaluating the KPIs outlined in the previous subsection requires a flexible approach that breaks down the multiple activities within the project into parallel contributions towards a more efficient and scalable network. In the scope of techno-economic architecture evaluations, five working streams were identified, that cover and compartmentalize the main areas of research and development within the project, and summarize the technological contributions that can be directly mapped to the project's general KPIs. These streams are:

- **WS1- Techno-economic Evaluation of DC Nodes at Central Offices**
- **WS2- Physical Architectures for Optical Nodes and Networks**
- **WS3- Economic Impact of Disaggregation Models**
- **WS4- Techno-economic Impact of Edge Computing**
- **WS5- Evaluation of Autonomous Networks and E2E Orchestration**

The mapping of these streams within the overall Metro-Haul architecture is illustrated in Figure 4. Note that while some streams overlap concerning the coverage of hardware or software components, each one entails a specific scope that makes the analysis unique and complementary. Therefore, the KPI evaluation of parameters such as cost, energy efficiency or bandwidth results from the juxtaposition of the findings from the different streams.

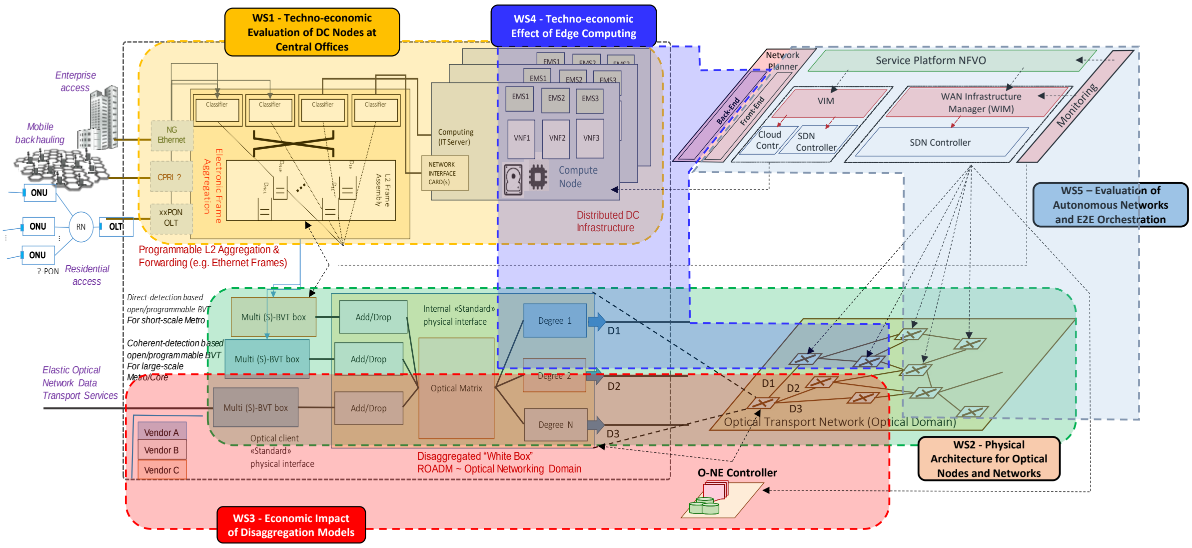

Figure 4: Mapping of working streams in the context of the Metro-Haul node and network architecture.

WS1 covers the possible architectures of COs to interconnect access networks (fixed and mobile), edge computing platforms, and metro optical transport. WS2 deals with the implementation options for the optical transport infrastructure, both on the line system (optical switching nodes, amplification, etc.) and transmission (transceiver configurations) aspects. WS3 addresses the architectural impact of disaggregation models on the optical layer, evaluating the architecture from the perspective of an infrastructure provider. Unlike WS2, which models the hardware specifications themselves in more detail, WS3 is more focused on the integration between hardware and control software, and how the different levels of disaggregation identified in [D3.1] influence network/service deployment cost and agility [ERicc18]._WS4 covers the network dimensioning aspects of enabling edge-computing platforms, on both the DC and optical transport domains. These studies leverage powerful optimization frameworks to, on the one hand, demonstrate the practical economic benefit of embedding AMEN/MCENs with edge computing capabilities (e.g. in terms of cost/footprint/bandwidth savings), and on the other hand, to identify how these architectures are essential towards guaranteeing the specific service requirements for various 5G-enabled applications, discussed at length in [D2.1] and [D2.2]. Finally, WS5 addresses how the control-plane architecture considered in Metro-Haul, aided by components to make networking more autonomous (e.g., related to fault detection and traffic prediction), can help the various stakeholders (infrastructure providers, network providers, verticals, etc.) deliver services faster and with less resource overhead.

These streams are organized to enable parallel technological deep dives into critical aspects of the metro optical architecture. In line with Metro-Haul's emphasis on supporting disaggregated/interoperable solutions, the overall architecture can ideally comprise different combinations of solutions for each building block, such as the architectures for the CO switching and optical transport node, the transmission hardware, and the control/disaggregation options running across the different hardware. In fact, given the plethora of service requirements and geographic/demographic constraints that are present in a metro transport environment, some individual solutions explored within Metro-Haul's scope will tend to be best suited to specific applications and network deployment scenarios. One key objective of the tasks outlined in this deliverable is to identify the most efficient combinations of technological building blocks according to specific scenarios' characteristics, rather than to propose a monolithic one-size-fits-all global architecture.

This parallel activity organization enables a more in-depth analysis over each stream, but on the other hand, requires the different subtasks to be carefully coordinated concerning the scenarios being addressed. To this effect, the working streams are complemented with common support tasks that provide the baseline frameworks for architectural evaluations. These form the backbone of the techno-economic analysis, providing realistic network topology and traffic assumptions for metro networks, service characterization models, open-source planning and optimization tools accessible to the scientific community for results dissemination and cross-validation, and cost/energy models for the constituent components of the converged data-center/optical-transport nodes.

Figure 5 shows the overall task organization, with the four support tasks and their interworking with the work streams, as well as the expected outcome in terms of the high-level KPI evaluations for each stream. The following sections of this deliverable provide a detailed description of the support tasks (Section 3), and the work streams (Section 4). For the support tasks, the work in this deliverable generally encompasses the definition of detailed reference network and traffic scenarios, and the methodologies to be used in each case (e.g. synthetic traffic generation, network service descriptors for planning tool purposes, etc.).

The "Network Scenarios for Architecture Evaluation" support task essentially defines the target geotypes that an AMEN/MCEN is expected to cover. Based on insourced data from partners within the consortium with access to realistic field data, as well as publicly available models, this characterization forms a discrete set of node dimensioning targets, in terms of the type and volume of services to support. Furthermore, expanding on the reference metro network topologies presented in [D2.2], the geotype parameters influencing the optical transport layout (user density and coverage area) are used to construct a standard hierarchy within the optical metro domain between access aggregation nodes and core-network interfaces. This tiered metro topology directly feeds WS2 with topology connectivity requirements, traffic volume baselines per node and expected optical fiber span ranges between nodes, all of which greatly influence the best physical architecture per geotype and tier. The reference topologies are also used as a comparison baseline for optimization algorithms and dimensioning studies within the scope of WS3, WS4 and WS5.

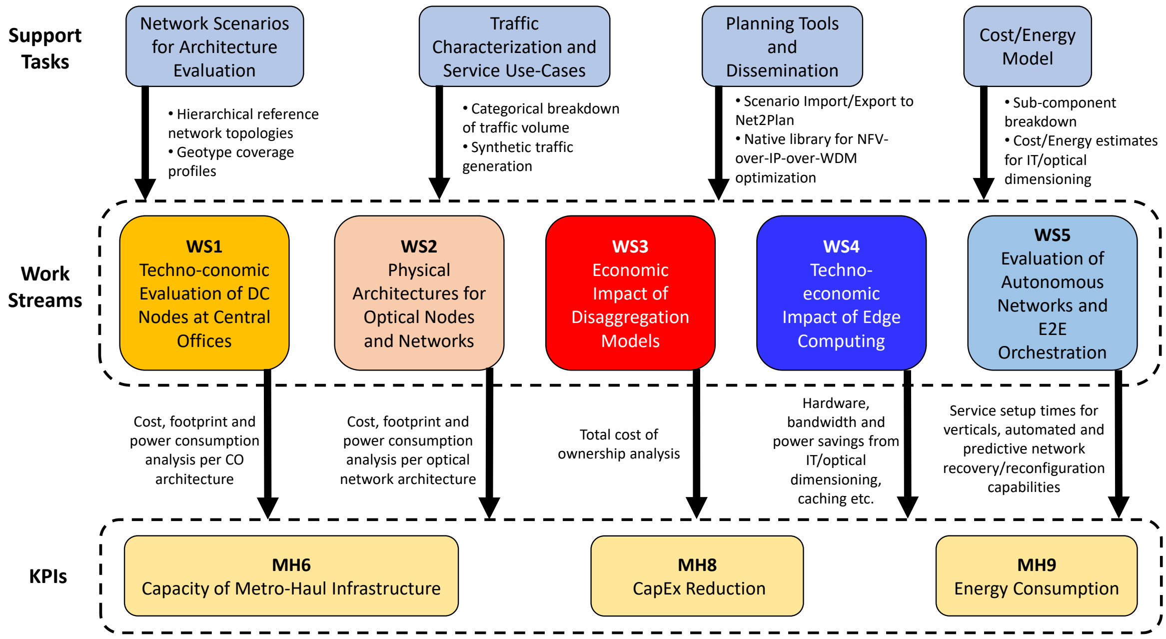

Figure 5: Sub-task organization and interworking between support tasks, WSs and task-specific KPIs

The "Traffic Characterization and Service Use-Cases" support task defines the end-to-end methodology to generate input traffic towards the various dimensioning/optimization evaluations across the various WSs. It distributes traffic into high-level categories based on the geotype analysis, and proceeds to characterize the different types in terms of (optical) connectivity required, bandwidth, variability, etc. This applies to both the traditional "background" traffic running over metro transport networks, as well as the 5G-specific service use-cases identified in [D2.1] and elaborated on in [D2.2] (including the vertical use-cases to be demonstrated within the project). Ultimately, this task provides a standard methodology to generate service-specific traffic traces with variable time granularities (depending on the dimensioning intended) that can be used throughout the various WSs.

The "Planning Tools and Dissemination" support task provides a standardized method of using the open source tool Net2Plan (which is also Metro-Haul's backend optimization tool component) [Net2Plan] to model the various network optimization and dimensioning problems for candidate architectures (particularly in WS2 and WS4). This effort covers not only providing an accessible and standardized method to load input data (topologies, traffic traces, network service descriptors, etc.) into the tool, but also the development of native libraries to aid in the development of optimization for joint IT/optical network optimization. Thus, this task fosters collaboration both within the project and with the overall research community, by providing an open-source platform to accelerate development and improve cross-validation of results.

Finally, the "Cost and Energy Model" support task details the methodology used to model the various components in the optical and data-center nodes. It aims to provide dimensioning studies with realistic cost and power consumption models that simulate the effectiveness of each proposed architecture, and ultimately validate the high-level KPI targets related to network efficiency. Much like in the "Network Scenarios for Architecture Evaluation" task, the cost/energy models leverage both in-house knowledge from consortium partners with regards to specific components developed within Metro-Haul, combined with public models for generally available components that are not a core competence in the consortium. The relevance of this task is not confined to attaining the project's goals, but can also be very useful to the broader research community. As previous projects in the optical transport space demonstrated, up-to-date cost/energy models are a key input for techno-economic research and product roadmap definition [FRamb13].

3 Frameworks for Techno-Economic Architecture Evaluation

This Section describes the support tasks created to align techno-economic evaluations under representative input data, including aspects such as network topologies, traffic modelling, planning tool interfaces and cost/power models.

3.1 Network Scenarios for Architecture Evaluation

3.1.1 Geotype definitions

The evaluation of new technologies for the deployment of next generation optical metro transport must take into account the abundance of different application scenarios concerning coverage, throughput, traffic profiles, etc. From the Metro-Haul node perspective (either AMEN or MCEN), this variability arises from the characteristics of the geographic area covered by the node, and from the service profiles themselves. While the latter determines how a Metro-Haul node handles application-specific requirements with a smaller granularity, the former shapes the aggregate volume and general pattern of traffic that must be handled. Given the broad scope of technologies addressed across Metro-Haul, it is not to be expected that a single architecture or component set meets all such high-level requirements optimally in terms of cost, footprint, scalability, etc. Hence, it is useful to categorize geographic scenarios into a discrete but representative set of geotypes, characterizing the areas covered by an AMEN/MCEN. This approach is well established when dimensioning COs (on which the physical locations of AMEN/MCENs may be based) for fixed and mobile aggregation [ComboD3.3].

Based on the combined data provided by network operators within the consortium (British Telecom, Telecom Italia and Telefonica), node coverage areas were subdivided into rural, suburban, urban and dense urban geotypes. For each of them, according to Table 12, a range for the physical coverage area, household density and mobile base station density is defined. For both the area and household density, an average value is also provided, since the distribution within each interval is not uniform. For the synthetic generation of geotypes used for the analysis in Section 3.1.2, an exponential distribution was assumed in each interval for both the area and households.

Table 12. Geotype attributes.

GeoType	Area [km^2]	Households per km^2	Household Range	Average Area [km^2]	Average Households	Share of Business Lines	Mobile Base Stations per km^2 [ComboD3.3]
Rural	50-200+	28-133	1680-7980	60	4800	0.1%	0.05
Suburban	13-75	120-1033	3600-30990	30	9600	1%	0.2
Urban	5-21	500-4000	9000-72000	18	18540	5%	1.5
Dense Urban	1-5	3500-5500	10500-16500	3	13500	10%	4

The coverage area determines the fiber distances required to interconnect nodes, which have a significant impact on the optical performance requirements for channels (e.g., transceiver specifications, allowable losses per node, etc.). The density of each area determines the total number of residential households (fixed lines). Furthermore, based on the geotype, a fixed number of business lines about residential ones is assumed, to differentiate traffic service profiles in urban/rural areas. The number of mobile base stations within the node's coverage is assumed to be area-dependent for each geotype profile, as assumed in [ComboD3.3]. As Section 3.1.2 will elaborate, the fiber distances can be used to characterize a set of nodes within a single metro aggregation chain (i.e., a single domain) per geotype. As a result, the geotype defines the raw volume and type of traffic per node, as well as the physical network characteristics of lower-tier metro networks (i.e., aggregation chains). This input can be used to obtain actual volumes by combining it with parameters that may vary over time, such as the number of fixed lines per household or the average traffic per line, as will be reported through Section 3.2.1. This input will be critical to understanding the concrete use-cases for different proposed optical layer technologies, such as for transceivers and switching node architectures.

3.1.2 Reference Topologies

In this subsection, the topologies that are considered for studies within the Metro-Haul project, specifically in WP2 activities of simulation, dimensioning and techno-economic evaluations, are specified.

The first version of topologies and network scenarios provided by operators was reported in [D3.1], and they are briefly summarized hereafter. They reflect the actual deployed operator networks or networks planned to be deployed in the short term. These operator topologies and network scenarios have been used to derive the unified Metro-Haul reference topology presented in this subsection.

The two main guidelines driving the definition of the reference topologies from specific operator scenarios have been the following: 1) tosynthesize the different scenarios capturing their unifying features as much as possible; and 2) to target the provision of a metro infrastructure supporting a heterogeneous fixed/mobile access, including future 5G RAN architectures. This infrastructure must include the typical range of distances to be covered within the metro segment, which is clearly stated in the project proposal.

Regarding guideline no.2, Metro-Haul assumes that the range of distances of the metro network segment is between 50 and 100 km, which fall within the typical transmission distances within a mix of urban and dense urban areas. Certain suburban or rural areas may require fibre distances longer than 100 km and up to 200 km.

In addition, two main types of "node functionality" (often referred to as "nodes" themselves), are defined in Metro-Haul: the AMEN, which gathers traffic from access networks and hosts a limited set of network services; and the MCEN, which assures metro domain internetworking, provides an extended set of centralised services and allows the interconnection with the backbone. Topologies are strictly connected to such functionalities, because each topological node can host one of the two, or both, such functionalities. This general and high-level framework on how access, aggregation and core functionalities can be put together to build an AMEN or an MCEN is described in Section 3.2 "Guidelines for METRO-HAUL Central Office Design and Assembly" of [D3.1].

Looking at the reference networks presented in Section 2 of [D3.1], the three scenarios are quite different and are briefly summarized here below.

The network proposed by BT counts 150 metro core nodes (they could host MCENs in the Metro Haul framework) plus about 1000 central offices performing the role of aggregation nodes (they could host AMENs). The architecture is organized in two levels: a higher tier of meshed Metro Core nodes plus a lower tier made of aggregation networks collecting each a subset of aggregation nodes. Metro Core nodes are connected through a mesh of logical connections (no details about physical topology are provided). The topology at the aggregation level is a horseshoe with 8 to 12 nodes between two metro core nodes (see Figure 3 of [D3.1]). Horseshoe lengths are generally within 50 km, but with the extension of this length in rural areas to up to 100 km and sometimes longer. There is no information on the percentage of chain distances lower than 50 km, between 50 and 100 km, and greater than 100 km.

Telefonica describes a network deployed in a metropolitan area of Latin America serving 5 million fixed and mobile users. The network is depicted in Figure 4 of [D3.1] and is structured in four levels organized using ring-star topologies (ring within the same layer and star when aggregating the lower layer). Looking at the Metro-Haul architecture and trying to establish a possible association between the Latin America Telefonica example network and a future network organized according to the Metro-Haul paradigm, the AMENs can match Layer 1 (L1) aggregation nodes (380 nodes organized in 19 clusters of 20 nodes each) while MCENs can coincide with L2 aggregation nodes (19 nodes with connectivity to L1 clusters, plus 14 pure transit nodes: 33 in total). Telefonica core nodes (6 locations in L2aggregation layer) can be the point of interconnection between MCENs and backbone nodes, while the access nodes can be simply considered as small access nodes consolidated into AMENs (6 small access nodes per AMEN on average). Regarding distances, and assuming a horseshoe topology for the aggregation level, the L1 clusters of potential AMENs aggregation towards a couple of MCENs can have a max path distance of the order of 170 km on average (as the mean link length of L1 nodes is 8.6 km) with the current size of clusters of 20 nodes. This distance is too much for the standard Metro-Haul assumption on horseshoe length in the aggregation segment. Then, in such a context, the number of nodes belonging to each cluster should be reduced to less than 10 on average (but it depends on the specific cluster). Distances between MCENs are an average of 13 km. No details are provided about the actual topology of the Telefonica network, but if we assume as a reference a square lattice regular topology of 36 nodes with links of 13 km, the longest paths between the farthest couples of nodes are of the order of 130 km, which is a typical mid-haul regional network diameter.

The TIM topologies described in Section 2.1 of [D3.1] are the current metro regional WDM transport network serving about 1,500 central offices out of a total of 3,700 (such a number is the target after the completion of an ongoing consolidation plan to reduce the central office number from about 10,000 to 3,700). The metro-regional network of TIM is composed of 14 separated macro-regional networks (some networks include more than one administrative Italian region) structured into two tiers: a meshed core and an extension (or aggregation) part connected using rings or weakly meshes to the core. Figure 1 of [D3.1] depicts such a structure. Each WDM metro-regional network is a transparent island made of a single flat network of ROADMs (in core nodes), and FOADMs (in aggregation nodes) and a lightpath can be established between any couple of nodes (belonging to either core or aggregation) if its optical feasibility results are satisfied. According to the Metro-Haul definitions, TIM metro core nodes are the candidate to become MCEN nodes, while the aggregation nodes will assume the role of AMEN nodes.

The remaining 2,200 central offices, not yet served by the metro, regional WDM network, are currently connected with other transport solutions like packet transport rings or dark fiber point-to-point connections. These 2,200 central offices not already covered by the WDM metro-regional network, in Metro-Haul framework could also be candidates to became AMEN nodes, and collected using horseshoes according to the topological model described in the following.

Considering the characteristics of the operator networks described above, together with the envisioned Metro-Haul network architecture, the reference network topology structure is therefore depicted in Figure 6. The network is organized into two tiers: an aggregation tier composed of horseshoe (or ring) subnetworks made of aggregation nodes and two core nodes (the last ones constitute the end-tail nodes of the horseshoe), and a core tier made of core nodes interconnected by a mesh of links.

According to the geotypes specified in subsection 3.2.1, the fiber length of a horseshoe depends on the specific geotype. Hereafter, *L* is the symbol used to identify the total horseshoe fiber length, i.e. the total length of the chain of AMEN nodes in km. Dense urban environments should have a total fiber length shorter than 25 km for the minority part of their aggregation networks, while the majority will have lengths between 50 and 100 km. Urban areas should be characterized by the majority of horseshoe lengths being less than 50 km, with most of the remaining part consisting of lengths of up to 100 km. Horseshoes in Sub-urban areas should be a balanced mix of cases with lengths less than 50 km, and lengths between 50 and 100 km, with some few exceptions where the length is equal to or greater than 100 km. Finally, rural districts are expected to show longer horseshoes, but with a minority of them still less than 100 km, the majority being with lengths ranging from 100 to 200 km and, another minority set that is greater than or equal to 200 km.

This model and the associated ranges of distances reproduce, with some adaptations, the real network scenarios proposed by operators, and can be applied also to other geographical and operator specific contexts.

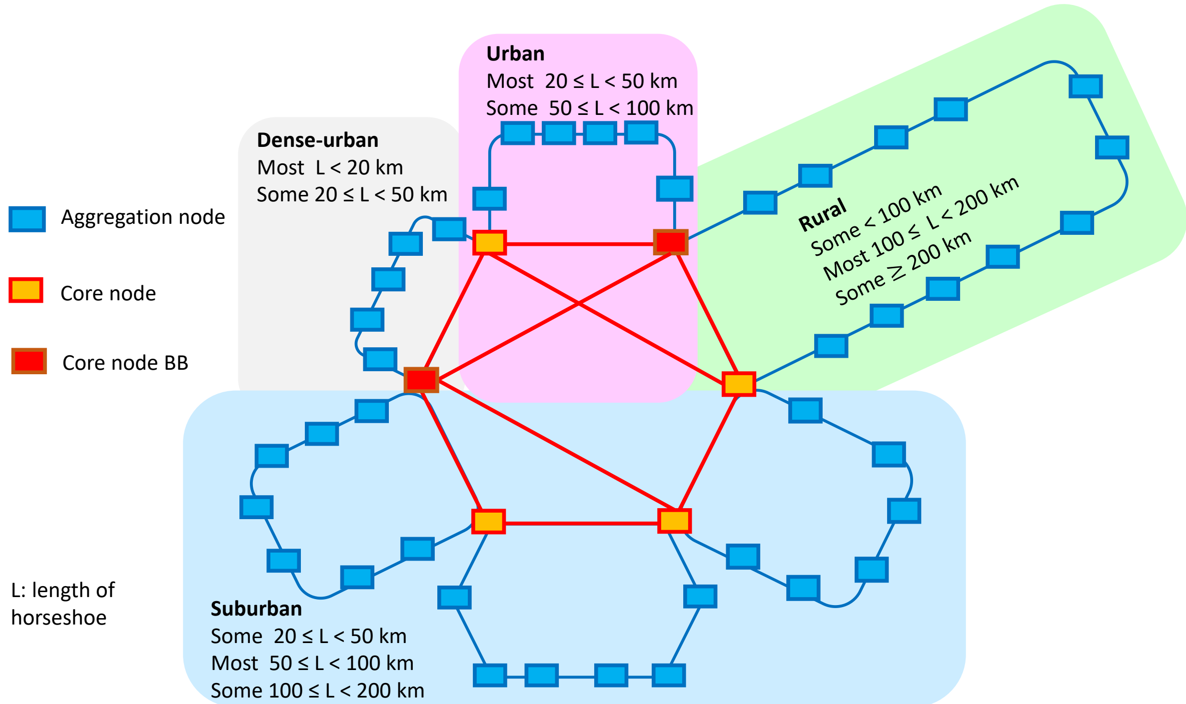

Figure 6. Reference Metro-Haul topology (L is the total horseshoe fiber length)

Table 13 reports the ranges of distances for the aggregation networks with horseshoe topology assuming that each horseshoe covers a homogeneous set of AMEN nodes. The exact share of each distance range per geotype context, as well as the percentage of horseshoes belonging to each geotype context about all of the horseshoes in the network, are very dependent on the specific geographic and demographic characteristics of each country and operator. Hence, the values in Table 13 should be taken as indicative, comprising averaged educated guesses from the available operator data.

Table 13. Horseshoe length ranges for homogeneous geo-types areas.

Geotype context (homogeneous)	Total distances range for horseshoe topology				
	L <20 km	20 ≤ L < 50 km	50 ≤ L < 100 km	100 ≤ L < 200 km	L ≥ 200 km
Dense urban	YES (majority)	YES (minority)	NO	NO	NO
Urban	NO	YES (majority)	YES (minority)	NO	NO
Suburban	NO	YES (minority)	YES (majority)	YES (rare)	NO
Rural	NO	NO	YES (minority)	YES (majority)	YES (rare)

Concerning the meshed network interconnecting the core nodes, the degree of mesh will depend on the specific implementation, but in general, one can expect a slightly meshed network with node degree values between 3 and 6, depending of course on the number of core nodes belonging to the metro core network. Regarding the role of core nodes within the Metro-Haul framework, they can exploit in general both AMEN and MCEN functions (and so they can be labelled as MCEN nodes) with only a subset of them that could be connected with the long-distance backbone. In Figure 6, the MCEN nodes interconnected with the backbone are the two red ones, while the remaining four orange ones form the central metro core mesh. Regarding the number of nodes in the metro core meshed network, it can range from a few nodes (4) to tens of nodes (40), depending on both the area covered and the traffic collected from the aggregation subnetworks.

Typical link distances and network diameter of the metro core mesh depend on the geographical context covered by the network. *Table* 14 shows typical node number, nodal degree, link lengths and network diameters of metro networks covering the two contexts of a city (the city area and its close surroundings only) and a wide metro-regional area.

Table 14. Length ranges in a Slightly Meshed Network for homogeneous geo-types areas.

Feature	Metropolitan Core			Regional Core		
	Minimum	Average	Maximum	Minimum	Average	Maximum
Node number	4	12	20	8	20	40
Node degree	2	3	5	2	4	6
Link distance	1 km	5 km	20 km	1 km	70 km	120 km
Net. diameter	10 km	20 km	50 km	100 km	200 km	400 km

3.2 Traffic Characterization and Service Use-Cases

In this section, the traffic characterization methodology is presented at two different scopes. On the one hand, *macro traffic* characterization aims to model the traffic at a high level. To this aim, a methodology to calculate the background traffic generated by the mass-market users collected on a node is described. Also, the expected evolution in the coming years of this mass-market traffic, as well as the traffic injected by the services behind the vertical use cases identified in previous deliverables, is briefly presented.

On the other hand, modelling aspects for fine granular (i.e. sub-hour) traffic are presented in the *micro traffic* characterization section. In this regard, the CURSA-SQ methodology to generate synthetic traffic flows is presented. To illustrate how macro and micro traffic characterization can be used to generate synthetic traffic traces for multiple purposes (network planning, simulation, machine learning training and validation, etc.), an example is provided.

3.2.1 Macro Traffic Characterization

In this subsection, macro traffic characterization is given in terms of data rate (b/s) in the busy period (usually the busy period can be of 15 minutes or 1 hour; in the last case it is called busy hour). Both directions, the downlink component (stream from the network to the user) and the uplink component (from the user to the network), are considered. For each traffic type defined, the downlink component (DL) is explicitly calculated, while the uplink component (UL) is derived through a specific parameter, as a percentage of the downlink component. In the following, only the data plane traffic is considered, as it the most relevant component in terms of volume, and because it has an impact on the metro architecture and dimensioning.

The macro traffic types to be carried by the metro network infrastructure are summarized in AMEN (UPF type 2).

Table 15.

Background traffic (type T1) is the traffic component due to residential and business mass-market customers and constitutes the evolution in the 5G perspective of the analogue component present in current networks. Traffic is collected from all types of access networks connected to the metro infrastructure: fixed residential and business access, fixed wireless access and all types of mobile radio access (as far as relevance of traffic volume generated is concerned, they are limited in practice to 4G/LTE and 5G NR).

The other type of traffic is the components due to the **vertical use cases (type T2)**, which stem from the specific analysis performed in [D2.1] and [D2.2]. For these types of traffic, a use-case specific modeling is performed, to emulate their unique requirements (e.g. specific VNF chain profiles) and assess their deployment over a Metro-Haul network.

Background traffic is in turn categorized with three main types of traffic subtypes, each one depending on the end points involved (terminal to terminal, or content provider to the terminal) and on the user plane function used to access the 5G core network (a centralized or a distributed user plane function). It is out of the scope of this model to deepen the discussion of the implications of 5G core architecture and its functional blocks on the traffic flow model described here. In the following, according to the 5G system architecture depicted in a simplified version in *Figure 7*, it is assumed that traffic from the access networks of any type enters in the 5G Core through a User Plane Function (UPF). UPF is a function that can be specialized per type of services and located at different points of the network. In the traffic model, two types of UPF are assumed to be present within the metro scope: a centralized UPF placed only at the MCEN connected with the backbone (UPF type 1), and a distributed UPF assumed localized at the edge at a AMEN (UPF type 2).

Table 15: Traffic types to be carried in the metro network

Traffic type	*Traffic subtype*	*Source*	*Destination*
T1 – Background Mass market Residential and business	**T1.1 Point to point (P2P)** Communication between user devices Transiting on UPF Type 1	AMEN	Metro Core (AMEN) Backbone
	T1.2 Heterogeneous server-mediated exchanged with a remote DC. Transiting on UPF Type 1	AMEN	Metro Core (MCEN BB) Backbone
	T1.3 Cacheable services Cache on local or remote DC Mainly, but not only, video Transiting on UPF Type 2	AMEN	Metro Core (AMEN or MCEN)
T2 - Vertical Use-Cases	**Subtypes from use cases defined in [D2.1]**	AMEN	MCEN Backbone

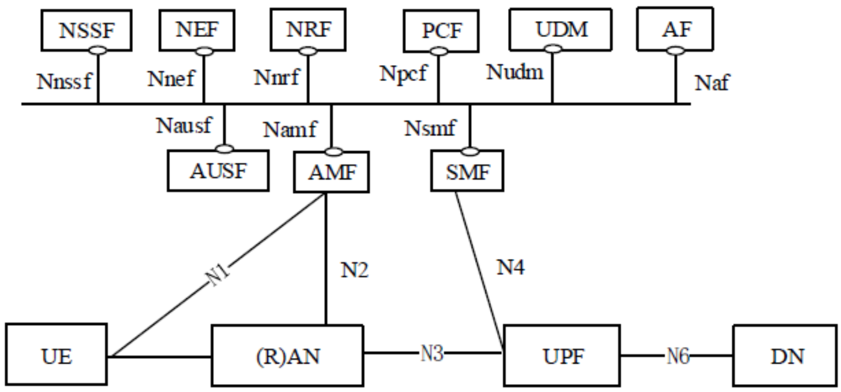

Figure 7: 5G System architecture interface representation (top: control plane, bottom: user plane) [3GPP-TS-22.261]

The three subtypes of background traffic T1 are:

- **T1.1 – Point-to-point traffic (P2P)** which involves two user terminals (fixed or wireless mobile) and is assumed to transit on a centralized UPF of type 1. The source of this traffic is one AMEN and the destination can be at another AMEN belonging to the same metro network, or another point outside the metro network reached through the interconnection gateway with the backbone (located in an MCEN BB).
- **T1.2 - Heterogeneous server-mediated traffic**, which involves a user terminal connected to an AMEN node and a server located in a centralised metro POP (an MCEN interconnected

with the backbone, i.e., an MCEN BB) or in a server outside the metro. As for T1.1, this traffic subtype also transits through a centralized UPF of type 1.

- **T1.3 - Cacheable services,** which, concerning their nature being of a high data rate (or high data volume) and also very popular (i.e., the same content is enjoyed by many users in the same period) can be effectively cached next to the user. The location of the cache can be extremely decentralized in the AMEN where the user is connected, or it can be placed in an intermediate point (in the near MCEN) or even in a more centralized point next to the backbone (MCEN BB). This traffic type relies necessarily on the UPF of type 2 because the content can be downloaded from the cache located at the AMEN where the user accesses the network.

The way in which the traffic subtypes generate flows in the metro network and towards the backbone is depicted in *Figure 8*. The network is schematically represented with three domains: Metro aggregation, Metro Core, and Backbone. Topologies are not relevant and can be horseshoes or rings in the aggregation and a mesh in the Core, as illustrated in section 3.1.2.

In the top-left part of *Figure 8* a legend is given for the symbols used for nodes and for the UPF function. MCENs are distinguished between the ones having a connection with the Backbone (MCEN BB, in red) and ones that do not have such a connection (simple MCEN, in orange). Flows for T1.1 are of two types: between devices within the metro domain (in green) and between a device in the metro and another outside (in grey). In both cases, traffic transits in a UPF of type 1 located in an MCEN BB. In the middle of *Figure 8* showing the T1.2 traffic type flows, the exchange of traffic of the device can be done with a source in the MCEN BB or with a source in the backbone. Also, in this case, the transit relies on a UPF of type 1 in the MCEN BB. Finally, in the bottom of *Figure 8*, the case of subtype T1.3 of cacheable services is given. Sources of service content can be located in four different points: in an AMEN (in blue), in an MCEN (in orange), in an MCEN BB (in red) or the backbone (in grey). Traffic is assumed to always transit in a UPF function of type 2 located in the AMEN where the user accesses the network because this is necessarily the case when the content source is in the AMEN.

The methodology to find out the background traffic from node and traffic subtype parameters is given in *Figure 9*. The methodology is based on three main inputs (I1, I2 and I3), is applied in a cascade of two models (M1 and M2), and obtains an intermediate output (O1, the traffic volume generated by a node) and the final output (O2, the traffic flows obtained for each traffic subtype).

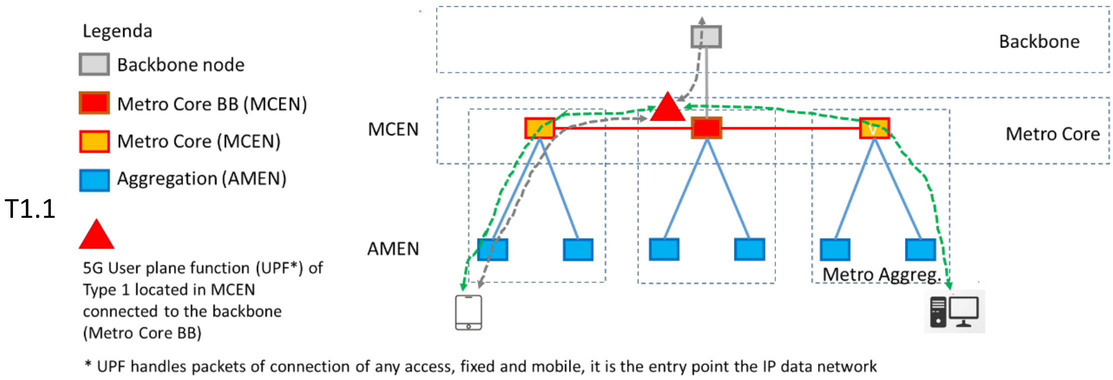

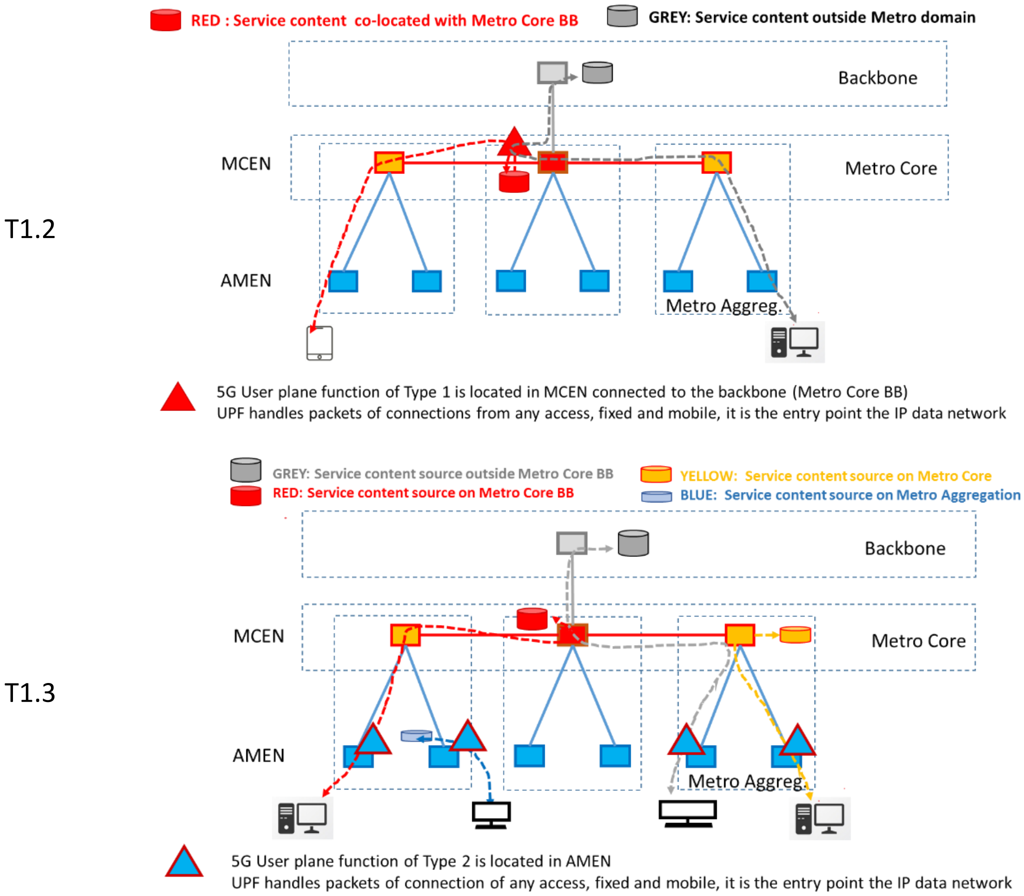

Figure 8: Traffic flows generated by subtypes of T1 background mass-market traffic.

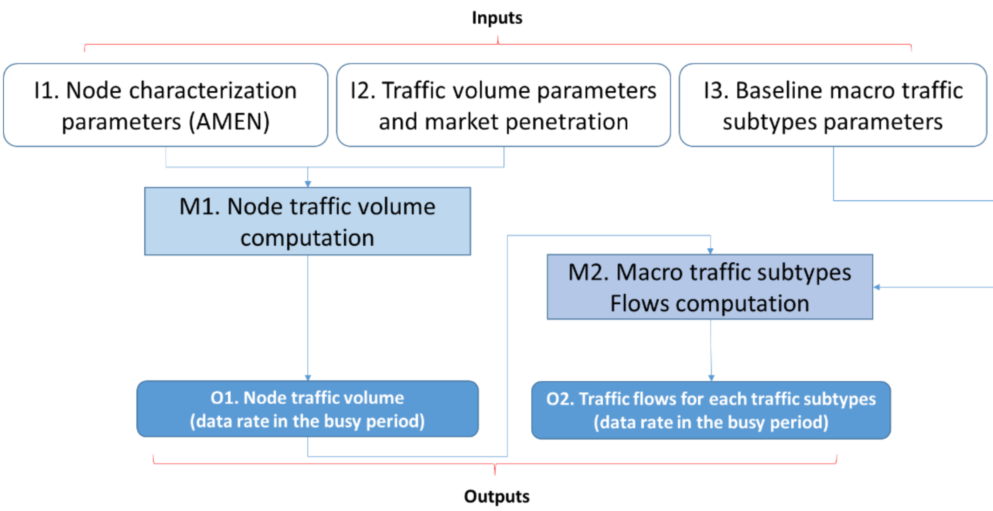

Figure 9: Methodology to compute the traffic offered to the node and flows associated with subtypes for background mass-market traffic (type T1 of AMEN (UPF type 2).

Table 15).

Inputs I1 and I2

Node characterization parameters (an example is reported in *Table* 16) are: the node name, the area covered, the number of households, the collected mobile sites (current situation regarding 4G deployment), and the NR Remote Radio Units (RRUs, not yet installed, but forecast under the assumption that 5G NR will have achieved an advanced stage of deployment). Concerning the mobile coverage, a typical site is currently organized to cover the whole 360 degrees space with three sectors (120 degrees each), with 3G and 4G usually coexisting, possibly exploiting more than one frequency band per each Radio Access Technology (RAT). In future, with the densification of the radio coverage of the space enabled by 5G NR, 3G will disappear, and the actual 4G site will be the point of collection of many NR RRUs. According to the figures in Table 16, the average number of NR RRUs per current 4G site is 5.

Table 16: Node parameters with values for the example “Suburban Alpha” (I1)

Node Name	*Area covered*	*Households*	*Mobile sites (4G)*	*NR RRU (5G)*
Suburban Alpha	25 km^2	20000	10	50

Table 17 reports market penetration and customer traffic parameters. The parameters “lines per household”, residential and business, account for the fact that the specific operator has a share of the market and this is expressed in terms of the number of lines retailed per household. These parameters can change with time depending on the combination of the trend of fixed lines (which is expected to continuously increase for many years) and the evolution of the market share of the operator. The values of average traffic per residential line (downlink value) are taken from [JAHer19] and are substantially in line with the data of operators participating in the project, and only slightly lower than the one reported in Table 4 of [D2.2] (10 Mb/s) which is targeted for a rather challenging network. The fixed business lines per household may either be generally defined by the geotype (i.e., as a function of the fixed lines as reported in Table 12), or through a specific parameter when more accurate data for a specific coverage area is available. The value of the average traffic per business line reported in Table 4 of [D2.2] (100 Mb/s) results in a very high and probably nota very realistic scenario. For this reason, traffic per business line is assumed to be twice the value assigned to residential users. Concerning the busy period peak factor (a multiplying factor for obtaining the traffic in the busy period from the average traffic) this has been extrapolated from [CiscoVNI17], where it emerges that the peak factor also appears to have a tendency to increase in the future.

Table 17: Market penetration and customer traffic parameters (I2, fixed part)

Year	*Residential fixed lines per household*	*Business fixed lines per household*	*Average DL traffic per residential line [Mb/s]*	*Average DL traffic per business line [Mb/s]*	*Busy period peak factor*
2019	0.50	0.05	1.8	3.6	4
2025	0.55	0.06	7.6	15.1	5

Table 18 presents parameters for two alternative options in mobile traffic calculations. Option 1 is based on the traffic generated by mobile devices. Starting from the number of mobile devices per fixed access line (it can be extrapolated, for instance, from documents like [AGCOM18], which reports market data about fixed and mobile lines for Italian telco operators), the number of fixed access lines on a node (sum of residential and business) and the volume of traffic exchanged by a mobile device (reported in [EricssonMobRep18]), it is possible to obtain the DL whole mobile traffic offered to a node in the busy period.

Option 2 is based on the traffic per 4G site and per NR RRU in the busy period. The total mobile traffic collected by a node can be obtained by multiplying these two parameters by the number of 4G sites and the number of NR RRUs connected to the node, respectively, and then summing them up. 4G DL traffic per node site of option 2 in Table 18 has been taken from Figure 8 of [NGMNWPLTE11] assuming that the 4G-system site operates with one band in 2019 and is expanded to two bands by 2025 (each band tricell generates around 40 Mb/s at a site, maximum number of bands for 4G are four). Concerning the DL traffic at the busy period per NR RRU, its value is not easy to estimate because the 5G radio access systems are still under development. However, a likely value is the one from Table 4 of [D2.2], from which the value reported in Table 18 is taken from.

A third option could be based on specific traffic per area (in Pbyte per year per km^2, for instance using numbers shown in the report [McKinseyRoadto5G18], which gives data, current values and 2024 projections, for some big cities worldwide). This approach, which would use the area covered by the node, is more critical to apply, because the specific traffic is strongly dependent on the geographic and socio-economic context covered by the node: a node-specific characterization would be needed.

Table 18: Parameters for mobile traffic calculation (I2, mobile part: Option 1 is based on traffic generated by mobile devices, and Option 2 is based on mobile station aggregated traffic)

Year	*Option 1 – device based*		*Option 2 - site/RRU based*	
	Mobile devices per fixed access (residential plus business)	*The volume of DL traffic exchanged by a mobile device [Gbyte per month]*	*DL Traffic per 4G site at busy period [Gb/s]*	*DL traffic per 5G RRU at busy period [Gb/s]*
2019	2	7.5	0.04	0
2025	3	40	0.08	0.6

The values assigned to the traffic given in *Table* 17 and *Table* 18 (which are realistic for the year 2019) are shown as a mere indicative example, and can thus be modified for use in the Metro-Haul evaluations should a more reliable, and updated reference data set to become available.

Model M1 and Output O1

The background traffic at an AMEN due to mass-market residential and business customers can be directly obtained by model M1 using input data I1 (*Table* 17) and I2 (*Table* 17 and *Table* 18). The results obtained for the node whose feature is collected in *Table* 16 (the node called *Suburban Alpha*) applying model M1 are reported in *Table* 19 and

Table 20 for fixed access and mobile access components respectively. All traffic values express the downlink (DL) component. The two methods used for mobile traffic result in different values, and this is due to the different node parameters and source of data used by the two methods. For 2019 the traffic of the mobile components of Option 2 is clearly overvalued (a factor of x5 comparing it with the corresponding value obtained with Option 1) because all of the mobile devices are assumed to be active and generating traffic in the busy period, while an important fraction of them may not be active. For 2025 the values for traffic generated by 5G NR of the two options are not so different.

Option 2 appears as the most reliable one in terms of realism, as it is linked with the actual potentiality of the infrastructure technologies and it does not depend on the combination of market data and customer profiles, which are intrinsically more uncertain, especially in future projections.

Table 19: Traffic collected in the busy period by the node of the example from fixed lines (O1)

Year	*Residential lines*	*Business lines*	*Fixed residential DL traffic [Gbit/s]*	*Fixed business DL traffic [Gbit/s]*	*Total fixed DL traffic [Gbit/s]*
2019	10000	1000	72	14.4	86.4
2025	11000	1200	418	90.6	508.6

Table 20: Traffic in the busy period collected by the node of the example from mobile access (O1)

Year	*Mobile DL traffic - Option 1*			*Mobile DL traffic - Option 2*		
	Mobile devices	*DL Traffic per device [Mb/s]*	*Total mobile traffic [Gb/s]*	*LTE traffic (4G) [Gb/s]*	*NR traffic (5G) [Gb/s]*	*Total mobile traffic [Gb/s]*
2019	22000	0.088	1.9	0.4	0	0.4
2025	36600	0.585	21.4	0.8	30.0	30.8

Input I3, Model M2 and Output O2

Table 21 shows in a single view the input I3, the model M2 and the output O2 of the methodology indicated in Figure 9. The traffic flows in the metro network depicted in Figure 8 associated with the background mass market traffic subtypes of Table 34 are obtained simply by multiplying the total traffic collected by the node by the I3 parameter coefficients (percentages in red bold in Table 21) which account for sharing of the whole traffic between subtypes and, within each subtype, for the percentage of traffic exchanged locally (within the AMEN) or remotely at different level (MCEN, MCEN BB or in a place outside the metro reached by the backbone). Values of the resulting traffic of Table 21 (DL and UL traffic flows in Gb/s) are obtained by applying the sharing percentages of the model I3 to the total traffic of 539.6 Gb/s which is the total traffic in the Year 2025 of fixed and mobile users as reported in the example of *Table* 19 (fixed users, 508.6 Gb/s) and

Table 20 (mobile users, Option 2, 30.8 Gb/s). The values in the bottom of Table 21 refer to the total T1 traffic and include the total traffic (in black bold), which is 100% of the AMEN traffic, and the other percentages and traffic values of the example, which results in *a posteriori* from the application of the model M2.

Table 21: Parameters I3 (percentages in bold red) for traffic subtype flows computation, and resulting values of traffic flow (output O2) for an example with total AMEN downlink traffic of 539 Gb/s as a result of the application of the model M2.

Subtype component	*AMEN downlink traffic*	*Uplink traffic w.r.t downlink traffic*	*Traffic going outside the metro area*	*Traffic exch. locally in the AMEN*	*Traffic exch. with the*	*Traffic exch. with*

					nearest MCEN	*the nearest MCEN BB*
Subtype T1.1 P2P	**10%**	**100%**	**33%**	-	-	**67%**
Resulting traffic of example in [Gb/s]	53.9 DL	53.9 UL	17.8 DL 17.8 UL	-	-	36.1 DL 36.1 UL
Subtype T1.2 Heterogeneous	**30%**	**30%**	**60%**	-	-	**40%**
Resulting traffic of example in [Gb/s]	161.8 DL	48.5 UL	97.1 DL 29.1 UL	-	-	64.7 DL 19.4 UL
Subtype T1.3 Cacheable	**60%**	**10%**	**15%**	**30%**	**20%**	**35%**
Resulting traffic of example in [Gb/s]	323.6 DL	32.4 UL	48.5 DL 4.8 UL	97.1 DL 9.7 UL	64.7 DL 6.5 UL	113.2 DL 11.3 UL
Total T1 traffic	100%	25.0%	30.3%	18,0%	12.0%	39.7%
Traffic of example in [Gb/s]	**539.4 DL (fixed + mob.)**	134.8 UL	163.4 DL 51.8 UL	97.1 DL 9.7 UL	64.7 DL 6.5 UL	214.0 DL 66.8 UL

3.2.2 Macro Traffic Evolution Assumptions

5G technologies and their deployment are acting as an on-going driver for network traffic growth and evolution of services. Even within this still current 4G era of network deployment, network traffic is continuing to grow exponentially, for example at a compound annual growth rate (CAGR) of 25% per annum (slightly reducing from 27% of the previous year), as shown in the Cisco CGI graph of Figure 10. Such strong growth rates as highlighted in the recent Cisco CGI 2016-2021 [CiscoGCI16] report continues to astonish, more than three decades after the initial advent of the Internet. Indeed, recent statistics from the BT network for a central core router are indicating even higher growth trends in peak usage bandwidths, with a recent trend of 45% CAGR growth per annum; albeit, somewhat lower than the historic CAGR trend of 65% (see Figure 11). Overall, it is clear that on-going research, development and deployment of optical technologies into 5G networks continue to be an important feature, to support these highly exponential traffic growth rate trends, and provide the necessary bandwidth capacities with agility and programmability.

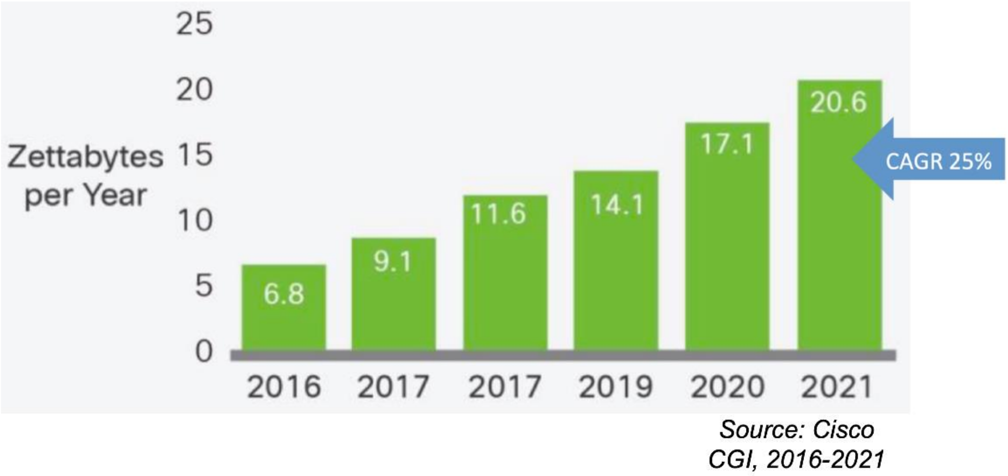

Figure 10: Data Centre IP traffic CAGR growth rate of 25% (2016-2021)

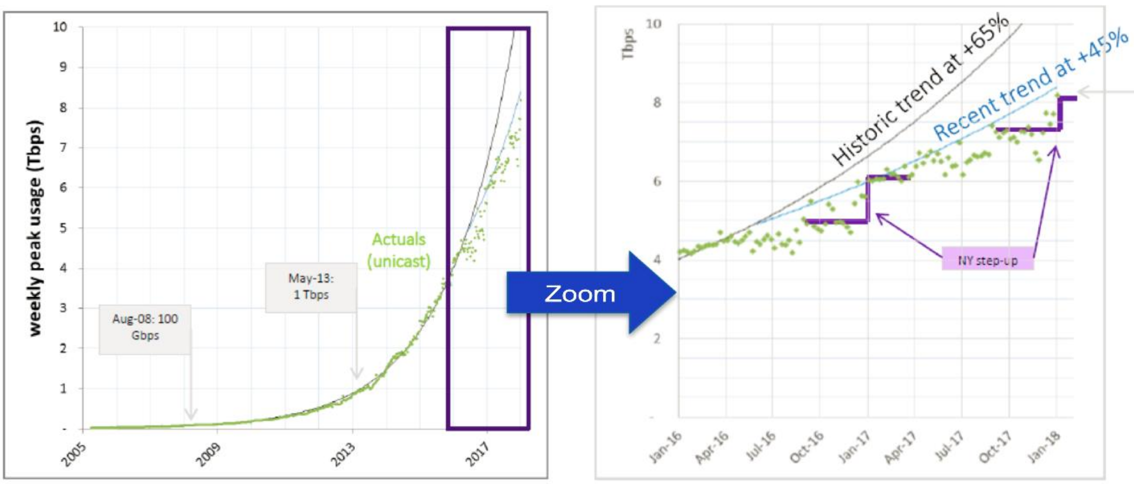

Figure 11: BT central core router CAGR growth trends, 2005-2018.

As a technology, 5G distinguishes itself from the previous 4G and 3G generations, since it represents an integrated and seamless networking (both hardware and software) and services technologies solution, with the new high-functionality services, therefore, representing a particularly important driver for the expected 5G deployment and take-up. Whereas the earlier generations technologies concentrated on the enhanced network performance from a physical layer (PHY) perspective (e.g. higher bandwidths, spectral efficiency, densification, and seamless mobility), 5G represents a symbiotic development of the key services and vertical use cases that will additionally help drive its successful deployment. Thus in the Metro-Haul project, in addition to the physical layer flexible and cost-effective optical technologies that are being developed, equally important attention is being paid to understanding the 5G services and vertical use cases that are exploiting the enhanced networking capabilities of 5G. Such new services and use cases will, therefore, tend to also act as drivers for the increasing traffic volumes and enhanced KPI specifications that are required to deliver the required quality of service/experience QoS/E parameters that must necessarily accompany the new 5G services.

Within Metro-Haul we have identified a selection of 5G use cases that we will be demonstrating in showcases at the end of the project, as well as other important 5G use cases also important in driving the on-going evolution of 5G technologies. In the previous deliverable [D2.2], we identified the anticipated bandwidth throughputs required by the various use cases at the initial deployment of 5G; however as indicated in the Cisco and BT growth trends above, it is clear that these use-case data throughputs will also grow substantially over time, such that the Metro-Haul network capacities will need to be designed to be able to grow alongside the increased demands of these services, and also be appropriately future-proofed.

In Table 22, we display the 5G use cases already discussed in [D2.2], but showing a progression of their aggregated bandwidth throughputs over a 15-year timeframe. The most bandwidth-hungry applications are the video-based services, such as LiveTV and 6DoF VR, which start at an aggregated 600 Gb/s and 250 Gb/s throughputs respectively, but both growing to 17 Tb/s and 7 Tb/s, respectively after 15 years. Indeed, for a horseshoe topology, with 10 AMEN's connected to an MCEN hub at the center, will require overall aggregated throughput bandwidths passing an AMEN/MCEN a further factor x10 higher. In which case, a single AMEN or MCEN can expect to see a rise in throughput capacities from 16 Tb/s through to 455 Tb/s after 15 years. Assuming 1000G (i.e. 1T) capacities per wavelength, an AMEN/MCEN featuring 455 Tb/s throughput capacities will therefore require 455 such 1T wavelengths.

Table 22: Aggregated 5G service requirements for a single AMEN/MCEN for different use cases featuring 25% CAGR.

	Throughput (Gb/s) @ 25% CAGR			
Use Case	*Year 0*	*Year 5*	*Year 10*	*Year 15*
mIoT Utility Metering	20	61.0	186.3	568.4
CDN (eMBB)	72	219.7	670.6	2046.4
LiveTV (eMBB+URLLC)	600	1,831	5,588	17,053
Service Robotics (eMBB+URLLC)	100	305.2	931.3	2,842
Smart Factory (eMBB+URLLC+mIoT)	50	152.6	465.7	1,421
6DoF VR (eMBB+URLLC)	250	762.9	2,328	7,105
ITS (eMBB+URLLC)	100	305.2	931.3	2,842
Crowdsourced Video (URLLC)	100	305.2	931.3	2,842
Secure SDN Control Video Distribution	150	457.8	1,397	4,263
Fixed residential	100	305.2	931.3	2,842
Business connections	100	305.2	931.3	2,842
Mobile broadband	100	305.2	931.3	2,842
TOTAL	**1,600**	**4,900**	**14,900**	**45,500**

3.2.3 Micro Traffic Characterization

In contrast to the macro traffic characterization in Sections 3.2.1 & 3.2.2, micro traffic characterization targets at modelling traffic flows with granularity below 1 hour, e.g. 1 minute. To this aim, relevant characteristics of two main categories need to be identified, namely: *i*) tidal traffic variations and *ii*) per-user service-related features of the distinct types of traffic. In the Metro-Haul project, we are using well-known studies in the literature to provide the key inputs for micro traffic characterization.

Concerning tidal variations, there are several research articles which have already pointed out daily variations, e.g. of Internet traffic over the time of day and actually every week [SAlba17]. This same behavior has also been observed of the Internet traffic of a University Campus [PVela16], at Internet Service Providers [AMozo18] and in the cellular towers of a mobile operator in Shanghai [FXu17]. *Figure 12* below shows the profiles collected from references [SAlba17, PVela16, AMozo18, FXu17]). Essentially, average (aggregated) traffic patterns typically exhibit two “camel humps”, with the busy hour occurring before lunchtime and a second hump right after lunch-time in business environments or, vice versa, at residential households. However, there are other patterns belonging to specific services that could be conveniently used to independently model the traffic types and subtypes in Table 15, i.e. video-on-demand traffic pattern experiencing three “camel humps” with incremental magnitudes from morning to evening [MRuiz16].

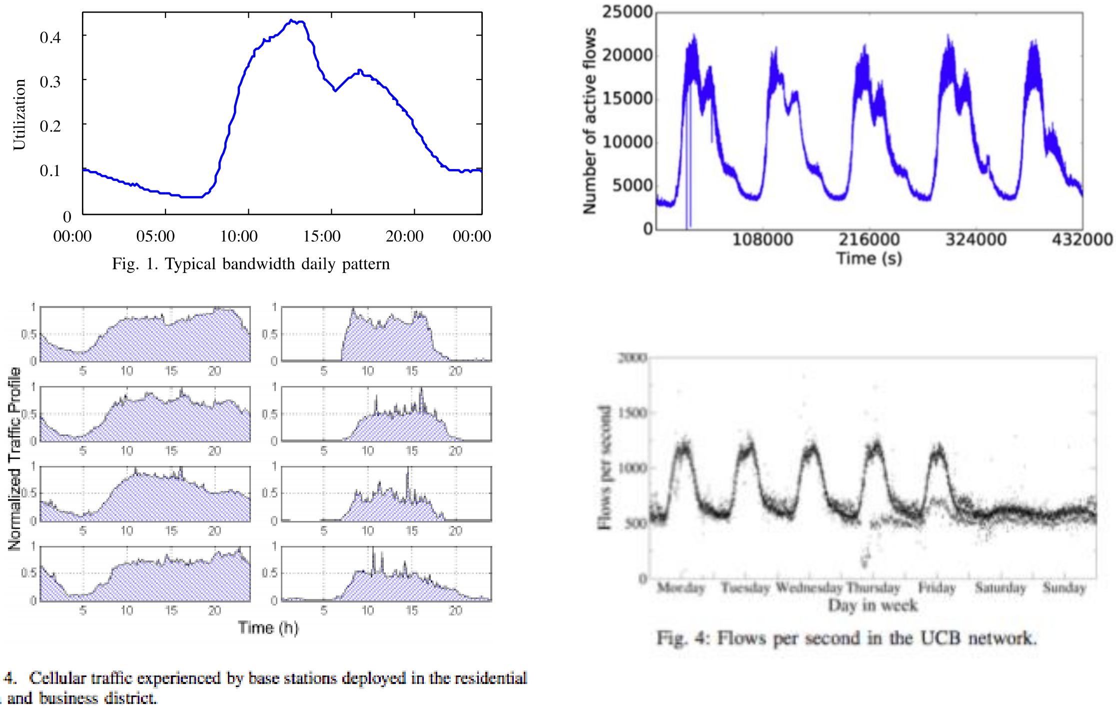

Figure 12: Reference traffic patterns (reproduced from [Albandea17, Mozo18, Xu17, Velan16])

For illustrative purposes, *Figure 13* shows tidal traffic patterns based on the references above that could be used to model business and residential contributions to T1 traffic type identified in Table 15. These tidal variations, normalized between 0 and 1, and detailed in Table 23, exhibit a peak-to-average ratio of about 2.5.

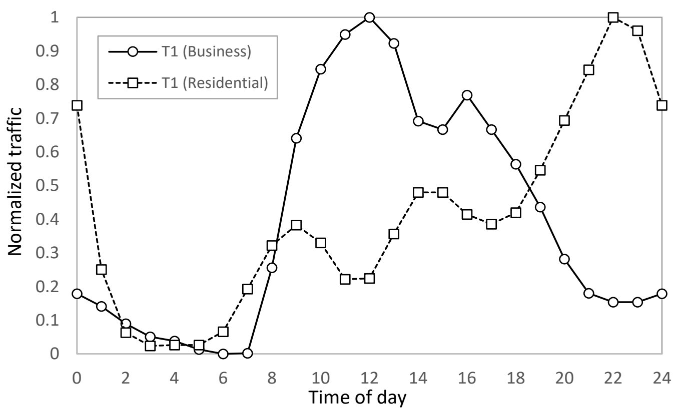

Figure 13: Examples of reference tidal variation.

Table 23: Tidal variation values

	12am	*1am*	*2am*	*3am*	*4am*	*5am*	*6am*	*7am*	*8am*	*9am*	*10am*	*11am*
T1 (Business)	0.179	0.141	0.09	0.051	0.038	0.013	0.000	0.002	0.256	0.641	0.846	0.949
T1 (Residential)	0.739	0.251	0.063	0.024	0.027	0.026	0.066	0.193	0.322	0.383	0.330	0.222
	12pm	*1pm*	*2pm*	*3pm*	*4pm*	*5pm*	*6pm*	*7pm*	*8pm*	*9pm*	*10pm*	*11pm*
T1 (Business)	1.000	0.923	0.692	0.667	0.769	0.667	0.564	0.436	0.282	0.180	0.154	0.154
T1 (Residential)	0.224	0.356	0.480	0.480	0.414	0.385	0.420	0.546	0.694	0.844	1.000	0.960

Regarding the characterization of services, several research works have provided analytical and numerical analysis of several services, mainly those behind traffic types T1 and T2. The following groups of characteristics can be distinguished:

1. *User activity behavior*: these characteristics capture the behavior of the users of a specific service, e.g. duration of video reproductions [LHuan17] or P2P conversations [JGome08].

2. *Data exchange*: these characteristics focus on how the service generates the data to be transferred according to users' activity. For instance, when a user requests content reproduction, a certain amount of audio and video (media) is sent to the user to fill an initial buffer. After that, media segments of a given short duration (e.g., 5 sec.) are regularly sent following a typical ON/OFF pattern until the content finishes or the reproduction is stopped [ARao11].

3. *Users infrastructure*: these characteristics allow adapting the data exchange to packet traffic since network infrastructure can impact the service. For instance, in a VoD service, video quality is adapted as a function of the throughput [HAzwa14]. This could impact the size of media segments and consequently, the packet traffic characteristics.

The above service-related characteristics are not deterministic, but they follow statistical distributions. Therefore, by analyzing them, traffic that every individual user introduces in the network can be modelled in terms of a few *random variables* capturing how bursts (and even packets) are generated by a single active user. The most relevant random variables are: *i*) *inter-arrival burst rate*, defined as the rate between consecutive bursts; *ii*) *burst size*, defined as the number of bytes transmitted in a burst; *iii*) *inter-arrival packet rate*, defined as the rate between consecutive packets in a burst; and *iv*) *packet size*, which is defined as the total number of bytes (headers included) in a packet.

3.2.4 Synthetic Traffic Generation Methodology

To generate traffic traces, we will use CURSA-SQ, a methodology to analyze network behavior when the specific traffic that would be generated by groups of service consumers is injected [MRuiz18]. CURSA-SQ includes input traffic flow modelling with second [s] and sub-second granularity based on specific *service* and *user* (*consumer* in terms of CURSA-SQ notation) behaviors, as well as a continuous G/G/1/k queue model based on the logistic function. The methodology allows us to accurately study traffic flows at the input and outputs of complex scenarios with multiple queuing systems, as well as other metrics such as delays, while showing noticeable scalability.

A general overview of the CURSA-SQ methodology is presented next. Without loss of generality, let us consider a scenario where a network operator provides connectivity between service consumers and service providers. This can be extended to other scenarios like machine-to-machine communications, etc. according to the needs of the vertical use cases behind the traffic type T3. Figure 14 illustrates the scenario where a service is requested by the consumers; UL traffic arrives from service consumers in a network node that aggregates and forwards it toward the selected service provider, whereas in the DL, such a node forwards the traffic coming from a service provider (in response to service requests) to the specific service consumer.

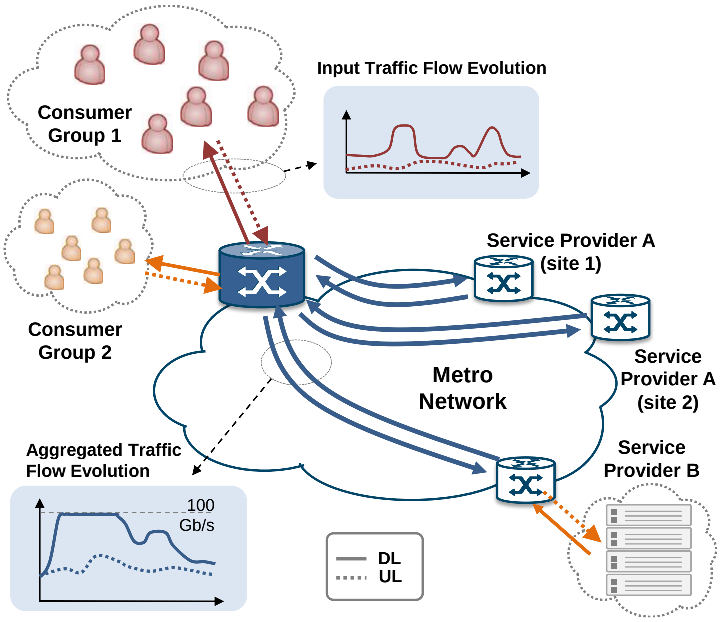

Figure 14: General overview of targeted scenarios.

The target of the proposed methodology is to study and generate traces of the aggregated traffic flows as a function of consumer traffic flows (hereafter, *input traffic*), as well as the characteristics of the network infrastructure. To reduce the number of input traffic flows, consumers of the same type of service and with the same characteristics can be grouped. Finally, a consumer group can be served from one (e.g. MCEN) or more locations (e.g. AMEN and MCEN) of the same provider.

We will use different traffic flow generators for the UL and DL. Those generators will generate traffic flows, in terms of bitrate (b/s), with a granularity *T* fine enough to study flows (in the order of hundreds of milliseconds) but several orders of magnitude higher than those typical times and sizes of packet-based traffic generation (Figure 15a). In the UL, one single flow generator per consumer group will be used to produce the traffic flow for all the active consumers in the group; this flow generator will be located at the consumer group location and will target one or more service provider's sites. In the DL, each service provider's site will contain a flow generator to produce the traffic flows toward the consumer groups.

The generation process is summarized in Figure 15b; it is based on first characterizing each service (labeled 1 in Figure 15b) to finding the UL and DL traffic characteristics (2) for a single service consumer. Then, the traffic flow bitrate is generated by scaling the traffic characteristics to the number of active consumers forecast for a given time period (3), while transforming the characteristics from the discrete into the continuous domain (4).

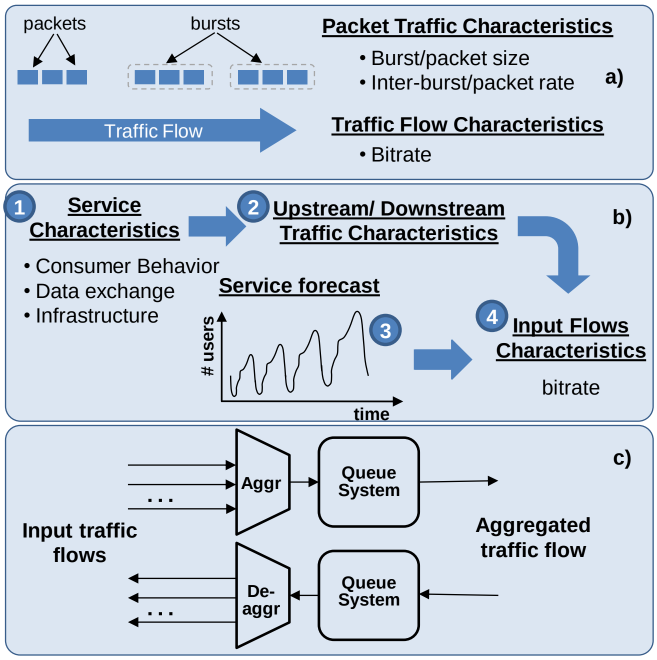

Figure 15: Overview of the CURSA-SQ Methodology.

Once input traffic flows are generated in terms of a bitrate for every period and every direction, they are used to generate aggregated traffic flows. To this end, several upstream input upstream traffic flows are aggregated, and the resulting flow feeds a queue system (Figure 15c). The reverse process is followed in the DL direction (for the sake of simplicity, not shown in the figure).

To illustrate how to use the above-mentioned methodology for fine granular synthetic traffic generation in conjunction with the macro and micro traffic characteristics detailed in this section, an example is provided in the following. Specifically, a synthetic trace of day 1 of the year 2019 containing traffic measurements every minute is generated for the subtype T1.1 (P2P) and the AMEN downlink traffic example i*n Table 21*. For tidal traffic, the business profile in Figure 13 is adopted. Also, based on the references [JGome08] and [RBoll08], the following variables and parameters allow us to characterize both the user's activity behavior and the data exchange behavior of P2P users:

Variable	*Distribution*	*Mean*	*Std*
Session interarrival time (s)	Gaussian (truncated)	10.61	573.93
Session duration (s)	Gaussian (truncated)	575.98	1461.66

Finally, the expected number of users and their access throughput limitations are based on the data provided in *Table 19* andTable 20, for fixed and mobile access, respectively. The result of applying the synthetic traffic generation methodology on the abovementioned numerical example is depicted in Figure 16. Note that the peak traffic approaches 54 Gb/s, which is satisfactorily in line with the macro traffic characterization of *Table 21*.

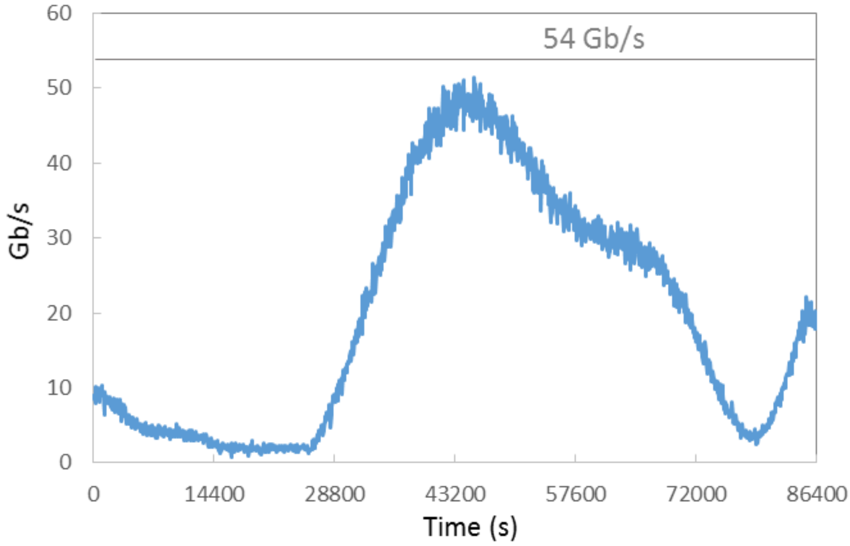

Figure 16: Example of synthetic traffic generation.

3.3 Planning Tools and Dissemination

This Section describes the interfaces developed using the open source tool Net2Plan, to natively support joint IT/Optical network dimensioning. A brief overview of the network planner component within Metro-Haul is provided, followed by the description of the optimization library support.

3.3.1 Planning/Placement/Reconfiguration Subsystem

The Placement, Planning, and Reconfiguration Subsystem (namely, Network Planner) aims to optimize the network resources from two different perspectives: Off-line network design algorithms are mainly devoted to capacity planning both for green-field scenarios and partially deployed (i.e. brown field) deployments. Once network infrastructure is in the production stages and operational, on-line resource allocation takes into account flows generated by end-user-oriented services that have different requirements in terms of bandwidth, delay and QoS. In this regard, the Network Planner enables the optimization of the resource allocation in the optical metro network to effectively provision VNFs in specific computing nodes considering heterogeneous requirements. The Network Planner architecture, introduced in Section 3.6 of [D4.1], is divided into front-end and back-end subcomponents, as illustrated in Figure 17.

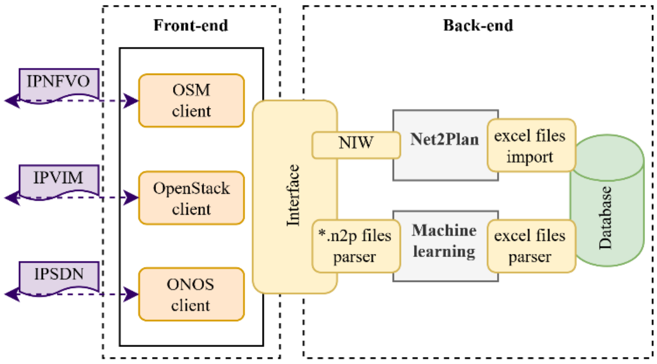

Figure 17: "NetPlanArch": Planning and reconfiguration system.

3.3.1.1 Front-end

The Front-end contains the interfaces that allow the exchange of information necessary for the planning tool (the Back-end) to carry out its operations (Figure 56 in Section [D4.1]) where the Network Planner interacts with the Metro-Haul Control, Orchestration and Management (COM) system. The COM, responsible for the dynamic provisioning of services, interacts with the Network Planner via the IPNFVO, IPVIM, IPSDN interfaces that provide the planning, the capability to query the Service Platform NFVO, the VIMs and the WIM in order to plan the network properly and provide network resource allocation and capacity planning solutions. Additional details are provided below. In line with the COM description in Section 4.5.1 of [D4.1], the Network Planner works across-layers interfacing with SDN controllers at the (Network) Control Layer and WIM/VIM/NFVO elements at the MANO Layer. For instance, the Network Planner can assist the SDN controllers on path computation while providing indications to the VIM/OSM elements for the placement of VNFs.

The Front-end includes a series of client modules to exchange information with the software components of the Control, Orchestration and Management (COM) architecture and consume it in the Network Planner framework. This information involves the network topology in terms of network nodes, links and traffic engineering attributes, such as occupied optical frequency slots or additive metrics as well as the status of the different hypervisors in computed nodes of the infrastructure. In particular, the exchange of information is performed exploiting three interfaces:

- Interface Planner – NFV-O (IPNFVO) aims to provision network services composed as an ordered sequence of VNFs through network service descriptors (NSD), which are the information model that represent network services in Open Source MANO (OSM) [1] terminology. OSM is an open-source NFV Management and Orchestration (MANO) implementation under the umbrella of ETSI, in charge of orchestrating the VNF placement and life cycle in one or more data centers, potentially connected by a transport network.
- Interface Planner – SDN controller(s) (IPSDN) aims to gather information from the transport infrastructure and provides transport network resource allocation in line with the requirements of the network services.

Relevant for IPNFVO and IPSDN, recent demonstrations reported the Planning Tool Net2Plan assisting an OSM instance in the optimal allocation and instantiation of network services (SCs) in a simulated transport network [FJMor18]. Subsequently, [FJMor18-2] demonstrates an ONOS controller emulating the packet-layer network and considered the end-to-end latency requirements to perform flow allocation jointly with the SC allocation and VNF instantiation via OSM.

- Interface Planner – VIM (IPVIM) interacts with the IT resource manager to enable multiple functionalities inherent of the NFV technology. In particular, a recent demonstration reported an open-source Net2Plan extension for interfacing multiple OpenStack instances, which enables multi-tenant slicing, IT resource visualization and VM migration [MGarr19].

3.3.1.2 Back-end

The Back-end takes in the information gathered from the Front-end to execute its algorithms with the eventual support/assistance of ML. The Back-end block has been conceived as a very flexible

[1] Website: https://osm.etsi.org/

environment which allows one to easily plug-in a variety of computation algorithms implemented in heterogeneous software environments, provided they are compatible with the overall architecture. Specifically, the Back-end is currently composed of two modules.

The first module is populated by algorithms natively developed in the open-source Java-based Net2Plan tool environment [Net2Plan]. Net2Plan has been chosen by the Metro-Haul consortium as a planning tool / Back-end module and as a *common framework* for collaboration among partners, exploiting its multiple functionalities for algorithm development, capacity planning, resource provisioning and automatic report generation. To this end, a specific library is available in the current Net2Plan version 0.6 in a Javadoc format, namely the NFV-over-IP-over-WDM (NIW) library described in the next subsection (3.3.2). Net2Plan will also help in disseminating the results in public repositories, for public validation and inspection. Algorithms hosted in this module include: a) a network service /service chain allocation algorithm; b) a VNF placement algorithm; c) a network Resource and Wavelength Allocation (RWA) algorithm. The algorithm (a) is in an advanced stage of development and already included in the Net2Plan framework, (b) and (c) is currently under development and integration. All these algorithms adopt network optimization techniques, are oriented to connection and service-chain real-time provisioning and exchange information with the Front-end using the NIW library integrated into Net2Plan. They are implemented in Java. The second module incorporates algorithms oriented to a predictive periodical or off-line network re-optimization; they make use of ML techniques and are developed in Python.

The front- and back-end modules exchange information using native Net2Plan *.n2p files, which are XML-based representations of network status. They include: d) an ML-based RWA algorithm; e) an ML-based VNF-placement algorithm; f) an ML-based algorithm integrating VNF placement and RWA. The algorithm (d) is in an advanced stage of development, (e) will be developed in the next months, while the development of (f) is targeted for the end of the project. The Back-end module also includes a database to store any information consumed and produced by the algorithms (see right-hand side of Figure 17). This permits the usage of historical data-sets that are particularly useful for training the ML algorithms. A database solution based on MS Excel files can also be used to import data to the NIW library as described in subsection 3.3.2.2.

3.3.1.3 Integration

Figure 18 shows the workflow involving the integrated components that summarize the process to, e.g. optimize network resources. It is divided into two different phases: training and real-time. In the training stage, the excel file containing the historical data regarding nodes, links, types of VNF executed, service chains and demands is provided to the back-end to train the machine learning algorithm. During the real-time phase, the interface between the front-end and back-end is responsible for providing the current network status via *.n2p files, both to a simple use-case that is computed using NIW, and to the machine learning module, which calculates the optimal solution according to its training. The result including the topology and allocated demands is then transmitted to Net2Plan. One of the main advantages of the approach we have followed in defining the architecture of the Metro Haul Planning Tool is the possibility of having a wide and expandable library of algorithms that are ready to be tested in the MH network architecture. Therefore, the user has several degrees of freedom in choosing the algorithms that will be used for a specific demo or a specific application. Moreover, it will be possible to compare the results of different algorithms on the same problem, to facilitate the evolution of the computation methods. For instance, we will be able to compare the single-step approach (integrated VNF placement + RWA) to the probably less well performing, but simpler, two-steps approach (first VNF placement, then RWA).

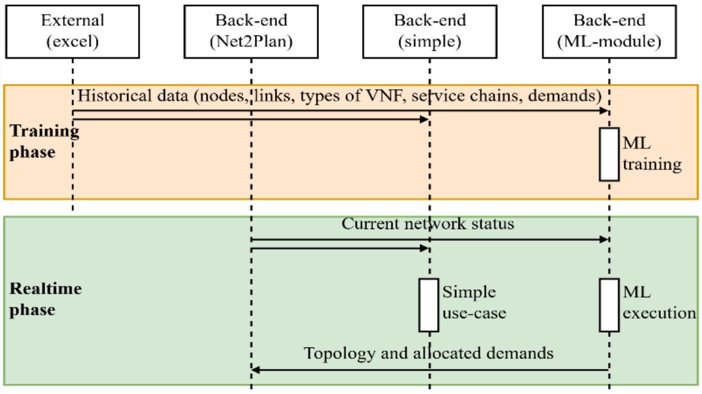

Figure 18: Workflow showing the integrated components that shows both the training and real-time phases.

We consider 4 possible attributes of implemented algorithms, namely VNF Reactive, Network Reactive, VNF Proactive and Network Proactive, see Figure 19. A given algorithm, e.g. #4, performs functions associated with VNF and Network planning in a pro-active way. The algorithms provided by the partners of the project can be split into six different functionalities, according to the type of optimization they implement, *i.e.* they compute the configuration either in reactive or proactive modes in terms of VNFs, network resources or both.

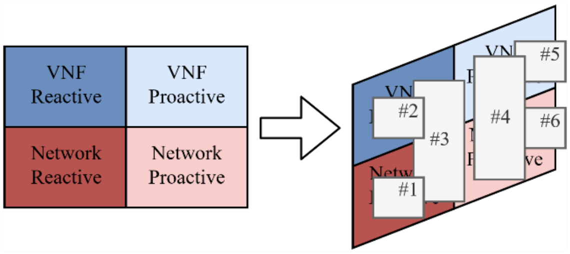

Figure 19: "AlgOptions": Placement, planning and reconfiguration algorithmic view.

3.3.2 Common Input/Output for Techno-Economic Evaluations

The NFV-over-IP-over-WDM (NIW) Net2Plan open-source library was created in the context of the Metro-Haul project. NIW includes detailed documentation in Javadoc format and a sample algorithm. NIW significant goals include: (i) ease algorithm development based on multiple available methods; (ii) provide a versatile framework which, based on its Excel spreadsheet importer and reporting functionalities, permits the abstraction of the underlying representations for rapid network engineering (network capacity planning and dimensioning); and (iii) foster collaborative works inside and outside the Metro-Haul consortium following its open-source philosophy. In the following, we describe the NIW library through its two major components. First, we detail the NIW structure and its implementation, listing the elements and methods that compose it. Second, we overview its functionality for importing data to Net2Plan via Excel spreadsheets.

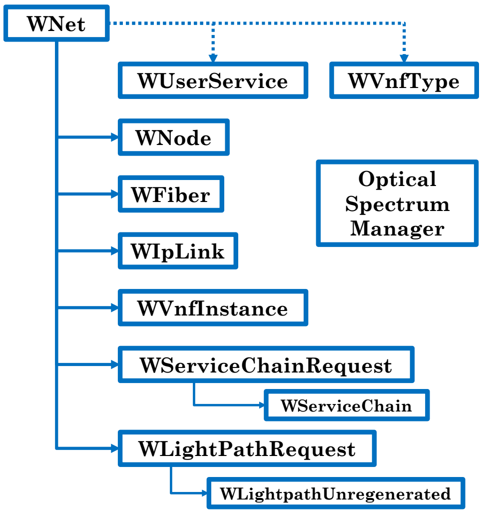

Figure 20: NIW library elements and their tree-based class hierarchy.

3.3.2.1 NIW structure and implementation

The NIW library contains the elements illustrated in Figure 20, which follow a tree-based class hierarchy:

- *WNet* represents a full IP-over-WDM network, with potential IT capabilities for VNF instantiation in the nodes. It gives access to all the network information and inherits all the methods and characteristics of all the other elements listed in Figure 20 except the Optical Spectrum Analyzer. Useful methods to import and export full designs are saveTo and loadFrom File.

- *WNode*, represents a node in the metro network. It contains network node definitions including name, type, position, population (i.e. the number of users served by this node) and hardware resources available for instantiating VNFs: the total amount of CPUs, RAM and HD space in the node. Nodes can be the end points of WDM links, lightpaths and IP links.

- *WFiber*, represents a unidirectional WDM optical fiber between two nodes. It is defined by the origin and destination nodes and optical fiber characteristics: length (km), propagation speed, valid optical slots that the fiber propagates, and physical-layer parameters (attenuation coefficient, chromatic dispersion coefficient, polarization mode dispersion PMD_Q factor). Optical slots are of 12.5 GHz width, and are identified by an integer. Slot 0 has a central frequency of 193.1 THz, and other slots are referenced from that baseline. The user can also specify the position of the optical line amplifiers and their gains, noise factors and PMD. WFiber elements are represented by links in the WDM layer of Net2Plan GUI.

- *WIpLink* represents the unidirectional part of a bidirectional IP link. It is defined by its origin and destination nodes, user-defined length (km), capacity and optionally, the lightpath that realizes the IP link. The IP link length is considered to be the physical length of the underlying lightpath, if it is realized by one, or the user-defined length mentioned above if not. WIpLink is represented by Link in the IP layer GUI of Net2Plan.

- *WLightpathRequest* represents the *intent* or request to establish a lightpath in the network. It is defined by the origin and destination nodes, line rate (Gbps), and the requisite or not to be realized with 1+1 optical protection. WLightpathRequest elements are represented as Demand elements at the WDM layer of Net2Plan GUI.

- *WLightPathUnregenerated*, represents a lightpath satisfying a particular lightpath request, without OEO regeneration in any intermediate node. It is defined by the sequence of fibers (WFiber) traversed, that make a path between the lightpath request end nodes. It is also characterized by the optical spectrum occupied in the traversed fibers, given as the identifiers of the occupied optical slots. Note that the model permits flexi-grid and mixed-line rate networks. WLightPathUnregenerated elements are represented as Route elements in the WDM layer of the Net2Plan GUI.

- *WServiceChainRequest,* represents the request for deploying an *anycast* service chain. A service chain request is characterized by a set of potential origin nodes, a set of potential destination nodes, the sequence of *types* of VNFs to traverse, and the amount of traffic injected in the origin of the chain. This traffic can be different for different time slots, thus representing multi-hour traffic scenarios where, e.g. the traffic intensities change along the day. The chain is also characterized by latency constraints imposed to the service chains satisfying this request, given by the maximum latency values from the origin node to the entry of each VNF in the chain. Also, it is possible to specify compression factors of the traffic when traversing each VNF. For instance, if a VNF in the chain is expected to compress the traversing IP traffic, its compressing factor can be introduced as a parameter in the service chain request, and automatically applied by the library. Service chain requests appear as Demand objects at the IP layer of the Net2Plan GUI.

- *WServiceChain* represents a service chain realizing a service chain request. The service chain is a sequence of IP links and VNF instances that, (i) start in one of the request origin nodes, (ii) end in one of the request end nodes, (iii) traverse VNF instances of the types and order mandated by the chain request. The initially injected traffic can be user defined, so the occupied resources at the traversed IP links and VNF instances are automatically updated.

- *WVnfType* is a user-defined type of VNF that can be later used as part of service chains. Examples include: transcoder, firewall, NAT, DHCP, etc. A VNF type is characterized by processing time, a maximum amount of IP traffic that it can process, and the amount of CPUs, RAM and HD space consumed. The user can also add VNF instantiation constraints, by specifying the set of nodes where VNFs of this type can be instantiated.

- *WVnfInstance* represents a particular instantiation in a particular node, of a VNF of a particular type. VNF instances can be traversed by service chains, summing up as much IP traffic as the one specified by its type.

- *WUserService* represents information of a user service that can be defined in the network and used to automatically populate the network with service chain requests. A user service can be, e.g., “videoconferencing”, or “Web surfing”. User services are bidirectional. Upstream traffic originates in the user, and ends in the network, and downstream traffic ends in the user, and is originated in the network. The characteristics of upstream and downstream traffic can be defined differently.

- *OpticalSpectrumManager* is used to account for the occupation of the optical spectrum in the network, e.g. informing of spectrum clashes, and providing useful methods for solving several routing and spectrum assignment problem variants.

Table 24. Excel file importer five-tab structure and attributes.

Tab	Parameters
Nodes	Name, Longitude, Latitude,Type, isConnectedToCoreNode, NodeBasePopulation, TotalNumCPUs, TotalRAM_GB, TotalHD_GB, ArbitraryParams*
Fibers	Origin, Destination, Length, Bidirectional?*, Slot Ranges, Attenuation Coeff, Chromatic dispersion Coeff, PMD design value, Amplifier Positions, Amplifier Gains, Amplifier Noise dB, Amplifier Pmd, ArbitraryParams*
Vnf Types	VnfType (name), VnfInstanceCapacity_Gbps, OccupCPU, OccupRAM_GB, OccupHD_GB, Processing_time_inMs, Constrained to some nodes?*, Nodes valid for instantiation, ArbitraryParams*
User Services	Service Name, List VNF upstream, List VNF downstream, TrafficExpansion_UP, TrafficExpansion_DOWN, MaxLatency_UP, MaxLatency_DOWN, InjectionDownExpansion, Ending in core node?*, ArbitraryParams*
PerNodeAndService TimeIntensity	ServiceChainInjectionNodeUniqueName, ServiceUniqueId, Morning, Afternoon Evening

* Optional attribute.

3.3.2.2 Excel file importer

The NIW library includes an Excel file importer functionality. As listed in Table 24, the Excel file is based on a five-tab structure with a specific scheme to be consumed and processed by NIW. The intention of the Excel file is to include enough information for (i) defining a WDM network topology with IT resources in the nodes, (ii) defining a set of user services, their VNF types to traverse, and their traffic intensity, potentially with time-varying intensities.

The NIW library will be able to convert such an Excel-based human-readable scenario, into a Net2Plan network instance NIW-compatible, with (i) the given WDM topology, and (ii) with *per-node* service chain requests generated according to the user services defined, and the node population.

The concept of *user service* is the key to defining the service chain requests in the network in a simple form. It represents the upstream and downstream flow of traffic that *each user* in *each node* of the network will request.

Adding a new user service to a design requires adding just one more row in a particular Excel tab. A user service definition includes:

- User service name. (“VoIP”, “Gaming”, …)
- List of VNF types to traverse upstream (user to a network) and downstream (network to the user).
- Traffic compression/expansion coefficients when traversing each VNF (1.0 if no VNF compresses or expands traversing IP traffic), and the ratio between the upstream and downstream traffic. That is the ratio between the average traffic that a user sends to the network, and the average traffic the network sends to the user in this service.
- Maximum latency values to enforce in the upstream and the downstream, measured in milliseconds to reach each VNF entry, counting from the service chain start.

- Per-user injection traffic, at each time slot. By defining more than one-time slot, it is possible to specify multi-hour traffic profiles.
- An indication if the service should reach the network core or not.

When importing an Excel file definition, the NIW library creates two service chain requests (upstream and downstream) starting in each network node, ending in (i) a metro node tagged as "core" if the user service is requested to reach the core, (ii) any network node if not.

The resulting design is an input to the LA-ML-SCA algorithm, which receives the service chain requests, and allocate them by (i) instantiating the required VNF instances, (ii) creating the needed IP links and lightpaths.

Additional information can be found in an exemplary Excel file [NIW].

3.4 Cost and Energy Model

This section is concerned with the cost and power consumption definition of the optical, packet-switched and data center building blocks to be used to design Metro-Haul's AMEN and MCENs.

3.4.1 Methodology and Inputs

Essentially, there are already several sources where both cost estimates and power consumption values of different types of equipment can be obtained. In particular, both previous and ongoing EU projects which have already conducted techno-economic studies and reported lists of components with their associated costs and power consumptions in publicly available deliverables. These projects and deliverables are:

- IDEALIST: D1.5. Evaluation of Flexgrid Technologies.

- 5G NORMA: Deliverable D2.3 Evaluation architecture design and socio-economic analysis[2]

- COMBO, D3.4 Assessment of candidate transport network architectures for structural convergence[3]

- ACINO, MS12 Cost Model for network scenarios

- ORCHESTRA, Initial report on market analysis, competitive analysis, standardization activities, dissemination and exploitation plans

- 5G-CROSSHAUL: D1.2: Final 5G-Crosshaul system design and economic analysis[4]

---

[2] Publicly available at (last access March 2019):
https://www.it.uc3m.es/wnl/5gnorma/pdf/5g_norma_d2-3.pdf

[3] Publicly available at (last access March 2019):
http://www.ict-combo.eu/data/uploads/deliverables/combo_d3.4_wp3_17june2016_v1.0.pdf

[4] Publicly available at (last access March 2019):
http://5g-crosshaul.eu/wp-content/uploads/2018/01/5G-CROSSHAUL_D1.2.pdf

While other past projects have techno-economic related reports, the equipment under study on those projects is out of scope for Metro-Hhaul. Examples are METIS II and CHARISMA which focus on costs of Antennas and Radio-related equipment, and are not critical for the design of MAN nodes. Other on-going projects with ongoing techno-economic analyses are, however, not available at the time of writing: 5G-CITY, 5G-TRANSFORMER, ONE5G, SAT5G.

That said, after a deep analysis on these deliverables and projects, we have realized that many of them rely on data collected during STRONGEST, which was summarized and published in an article in the Journal of Optical Communications and Networking back in 2013 [FRamb13]. Both equipment settings, cost and power consumption values reported in [FRamb13] and [WHedd12] are already too old for use today (2019), so the Metro-Haul consortium has decided to investigate new hardware equipment, cost and power consumption with present features and values for the realization of up-to-date AMENs and MCENs, and only use data from past projects for those hardware components where an accurate cost estimate could not be found.

Initially, a list of components to realize Metro-Haul's AMENs and MCENs was generated and circulated among the partners. This list tried to consider not only existing hardware already available in the market, but also experimental components under research from Metro-Haul's partners. For these experimental components, cost and power consumption estimates were generated by those responsible partners.

In a second stage, this list was circulated among all partners to collect a survey. Partners, especially those closer to the market like operators (BT, TID and TIM), manufacturers (NOKIA, ADVA, ERICSSON, CORIANT) and research centers with experimental groups (CNIT, CTTC, TUE, HHI) filled the table with a list of potential components to be considered for building metro nodes. Partners then provided cost and power consumption estimates along with a number between 1 and 5 which refers to the level of confidence of this partner on the numbers provided (being 1 low confidence and 5 highest confidence).

In a third and final stage, a final list was elaborated combining the cost and power consumption estimates for all components. Concerning costs of components, we did not use absolute prices since these are highly variable over time and can change dramatically over the years; instead, we decided to use relative prices of components concerning some baseline since such values tend to be more stable in the long term. A similar approach was used by the authors of [FRamb13, MBala18]

This final list includes post-processing of the data provided by all partners, weighted by their confidence levels and after removing outlier estimates. In general, values with highest confidences weight more to the final numbers than estimates with low confidence.

As a result, a final list of cost and power components was provided with three main weighted estimates: average weighted value, optimistic value and pessimistic. This final list will be used in the techno-economic analysis to be conducted and reported in deliverable D2.4.

3.4.2 Model Components

The realization of Metro-Haul's AMEN and MCEN nodes will require a long list of hardware components, including not only optical and packet switching equipment, but also computing and storage hardware boxes. In this light, a preliminary list of components that will become the basic building blocks of MH nodes is listed below, split per functionality. It is worth noticing that only generic descriptions are available, while the full cost and power consumption data will be delivered in D2.4.

Table 25. Optical components of the cost model.

OPTICAL COMPONENTS	
Name	**Description**
GENERIC OPTICAL EQUIPMENT	
Dispersion Compensation Module (DCM)	
Optical splitter 1x2	Splitter module to be used with wavelength blocker, different split ratios
Optical splitter 1x4	
Optical splitter 1x8	
EDFA dual stage variable gain (booster)	splitter and amplifier terminal for WL blocker (1) with booter and pre-amp. (2) only with pre-amp
Wavelength blocker	only blocker device
WSS 1x4 module fixgrid (one direction)	Fixed-grid Wavelength Selective Switches with different input/output port combinations
WSS 1x9 module fixgrid (one direction)	
WSS 1x20 module fixgrid (one direction)	
WSS 1x4 module flexgrid (one direction)	Flexgrid WSS with different input/output port combinations
WSS 1x9 module flexgrid (one direction)	
WSS 1x20 module flexgrid (one direction)	
Fixed-grid Mux/Demux	Fixed-configuration (50/75/100 GHz) A/D for 40/48/64/80/96 channels
COLORED WDM TRANSPONDERS (fixed-grid)	
10G	
25G/50G	PAM4 IM/DD
25G/50G	Coherent with integrated DCM
100G	
FLEXIBLE TRANSPONDERS	
BVT Transponder 100G	
BVT Transponder 400G	2x 200G line side, 4x QSFP28 at client side
S-BVT transceiver 2x50G (MCM/DD)	
BVT Transponder 100-600G (2020/21 onwards)	CFP2-ACO with variable baud-rate/mod-format (32/64Gbaud and QPSK through 64-QAM)

Table 26. Packet-Switching components of the cost model.

PACKET-SWITCHING COMPONENTS	
Name	**Description**
TELCO-LIKE SWITCHES Description: N interfaces at 10 Gb/s + M interfaces at 40 Gb/s + L interfaces at 100 Gb/s	
Small	matrix 1.6 Tb/s (N,M,L) such that 10*N+40*M+100*L smaller than 1600
Medium	matrix 4.8 Tb/s (N,M,L) such that 10*N+40*M+100*L smaller than 4800
Large	matrix 6.4 Tb/s (N,M,L) such that 10*N+40*M+100*L smaller than 6400
ExtraLarge	matrix 12.8 Tb/s (N,M,L) such that 10*N+40*M+100*L smaller than 12800

L2 DC-LIKE SWITCHES	
Leaves will be TOR-like switches aggregating traffic from servers computing/storage nodes towards an upper switch/router, Spines can also be used to interface the Optical Nodes (Nx10G,Mx40G,Lx100G)	
Small	Only 10G interfaces, eg (N=32 or 48, M=0, L=0)
Medium	10G and 40G interfaces, i.e. (N=48,M=4 or 24, L=0)
Large	40G and 100G interfaces, i.e. (0, M=36, 6)
ExtraLarge	40G and 100G interfaces with large port density and Deep Buffer Jericho 2 (0, M=40 or 80, L=40 or 80)
ROUTER CHASSIS + LINECARDS	
Legacy, Juniper MX, Nokia 7750, Huawei NE40, etc	
Small chassis	10 slots x LineCards (below)
Medium chassis	16 slots x LineCards (below)
Large chassis	32 slots x LineCards (below)
LineCard Nx10G	LineCards with different port density and capacity per port e.g. 14, 40 or 48 ports at 10G,
LineCard Mx40G	e.g. 3, 10 or 25 ports at 40G,
LineCard Lx100G	e.g. 1, 4 or 10 ports at 100G
PLUGGABLES	
10G SFP+ copper	10 – 100 Gb/s, below 5 meters reach, server to leave (i.e. TOR) switch
40G QSFP copper	
100G QSFP copper	
10G SFP+ optical	10 – 100 Gb/s, optical SMF long reach
40G QSFP optical	
100G QSFP optical	
100G CFP optical	

Table 27. Data-center components of the cost model.

DATA-CENTER COMPONENTS	
Name	**Description**
COMPUTE NODES	
Examples: HP Proliant DL 360, Huawei 1288-v5 or 2288H-V5, Dell R740 or R940	
Small	Intel Xeon Gold 6134 with 8 cores, 64 GB RAM, 600 GB HDD
Medium	Intel Xeon Gold 6140 with 16 cores, 128 GB RAM, 1.2 TBB HDD
Large	Intel Xeon Platinum 8160 with 2 CPUs x 24 cores, 128 GB RAM, 3.6 TB HDD
ExtraLarge	Intel Xeon Platinum 8160 with 4 CPUs x 24 cores, 192 GB RAM, 3.6 TB HDD
STORAGE	
Examples: Samsung SSD, iXSystems TrueNAS X10 2U, HPE Nimble Storage HF40/60, Dell EMC Unity 300	
SSD Small	4 TB capacity
SSD Large	8 TB capacity
NAS Small	20 TB
NAS Medium	60 TB
NAS Large	120 TB
NAS ExtraLarge	400 TB
OTHER SPECIALIZED HARDWARE	
BRAS	
HW Firewall	
Load Balancer	Eg Fortinet FortiGate 3000D
Data Domain Backup	eg HPE OpenCall Load Balancer, Radware Alteon. F5Big-IP, etc

4 Candidate Architectures and Evaluation Methods

This Section details the ongoing efforts in each working stream to define candidate architecture and methodologies for techno-economic evaluations of the various aspects of Metro-Haul nodes and the E2E Metro-Haul network.

4.1 Techno-Economic Analysis of DC Nodes at Central Offices

4.1.1 Candidate Architectures

To host distributed computing workloads in the form of VNFs, the Metro-Haul AMEN/MCEN architecture includes Datacenter infrastructure deployed at the Central Offices alongside Metro optical equipment. These distributed workloads will be the enabler for new business models (e.g. verticals services) to achieve Metro-Haul objectives, and they can also reshape the way Metro IP networking features are implemented. Besides enabling new business opportunities, Metro-Haul nodes have to maintain a parity of features with existing deployments to account for service continuity. In that respect, it must be noted that current Metro Central Offices, in addition to optical transport capabilities, also encompass IP level functionality to inject and process current business and residential services and aggregate, at either Layer 2 or Layer 3, both fixed and mobile access networks.

Different local networking architectures for deploying DC nodes in Metro Central Offices have to be analysed to provide the most optimized solution in terms of techno-economical parameters. The baseline for this comparison would be to connect the DC nodes to existing Metro Central Office architectures. This section describes different alternative scenarios to this baseline reference to include the functionality of Data-Center nodes in a Central Office architecture.

4.1.1.1 Fabric-less scenario

In this scenario, exemplified in Figure 21, the DC nodes at the Central Office are connected directly to the Central Office Layer 3 routers, that aggregate the local fixed and mobile access nodes and that are connected to the local transport optical equipment that connects this Central Office with other Metro locations

Comparing this scenario with the other alternatives, its main characteristic is that it has no leaf and spine switching fabric, so the packet layer of the Central Office is composed of just one router (if no equipment redundancy is required) or several routers (to account for equipment redundancy or capacity growth).

The layer 3 routers present in this scenario could be based either on whiteboxes or on traditional legacy routers. The Layer 3 routers will implement most of the IP features required at the Central Office, both for IP Core network integration as well as for Data-Center Interconnect (DCI) features to connect the workloads present at this location as VNFs with other workloads in remote DCs (either centralized or distributed ones).

Depending on the GeoType considered, this scenario can be foreseen to be cost-effective (to be ratified by the subsequent Techno-economic studies) for very distributed scenarios of Central Offices that provide services to a small number of users (and hence a small number of servers to connect). That could justify its implementation with just one router in cases where the small number of subscribers served allows for no equipment redundancy.

It should be noted that the legacy variant of this kind of scenario could be used, regardless of the GeoType, as a baseline for the different scenarios considered and that will be shown hereafter, because this is the typical architecture deployed currently in Central Offices, with the exception of the DC nodes that are not present as of today.

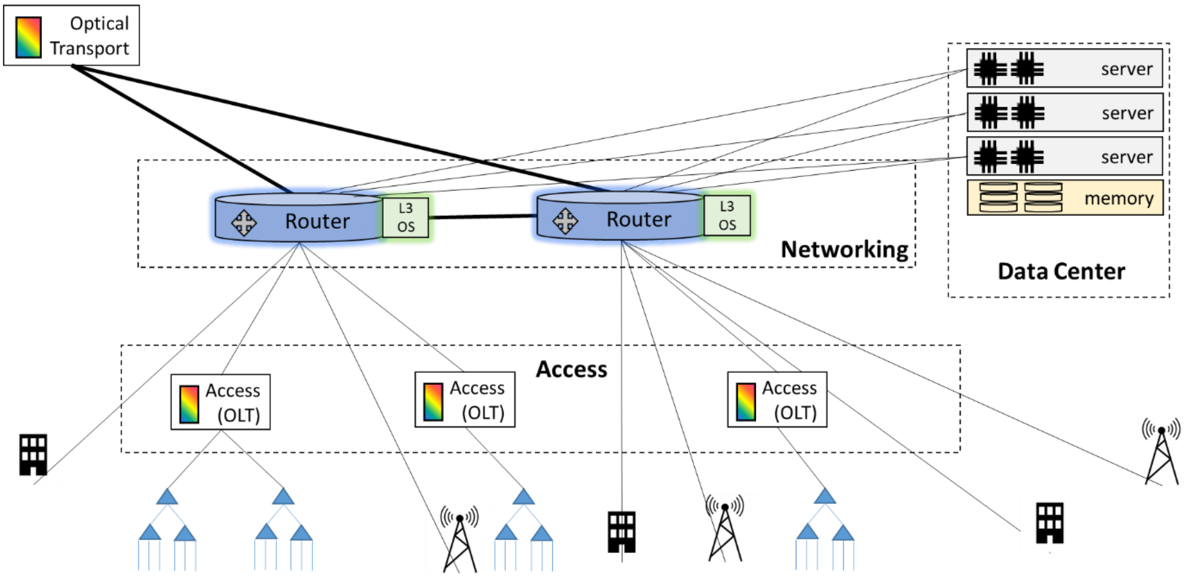

Figure 21: Fabric-less scenario.

4.1.1.2 L2 Fabric Greenfield Scenario

In contrast to the previous scenario, the L2 Fabric Greenfield Scenario shown in Figure 22 will be composed of a Layer 2 switching fabric, but without any IP router, which means that all IP functionalities, the IP Core network integration and the Data Center Interconnect, should reside on VNFs in the Central Office at the DC nodes. This type of topology is similar to the CORD (Cental Office Re-architected as a Datacenter) architecture. CORD is an ONF open source project intent on transforming the edge of the network into an agile service delivery platform. The main objective is to deliver economies of scale for data centers with the implementation of certain functions as VNFs with the SDN programmability of the Central Office leaf-and-spine switching fabric. The CORD architecture creates two control planes: one for networking, using ONOS as the controller, and one for managing VNFs, using a new CORD operating system dubbed XOS, which stands for the everything-as-a-service operating system.

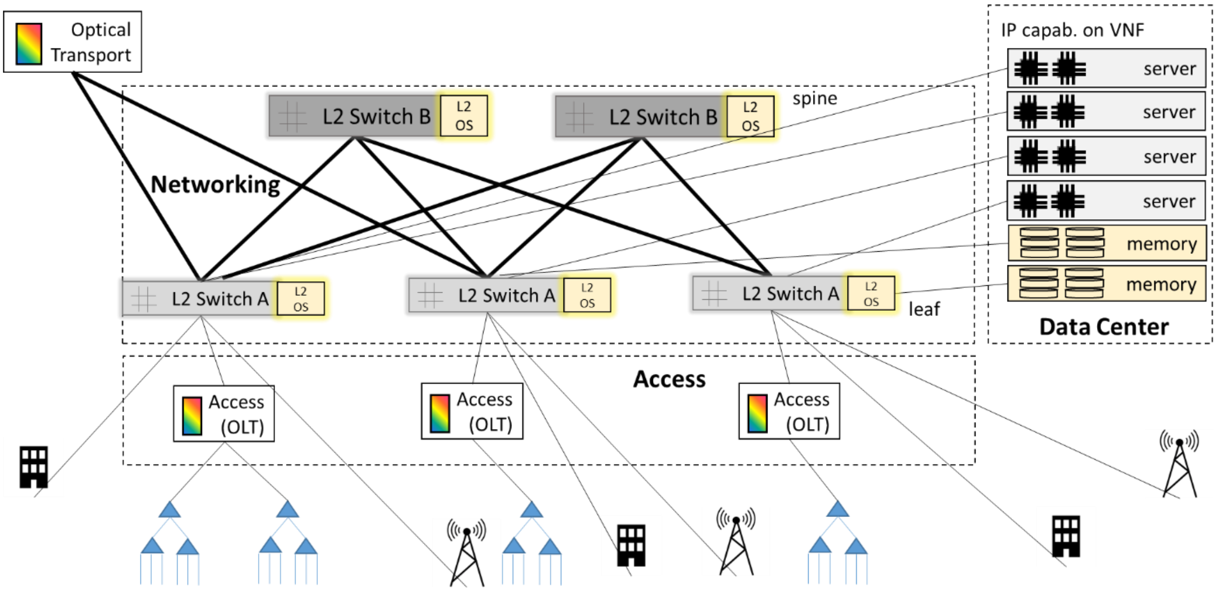

Figure 22: L2 Fabric greenfield scenario.

4.1.1.3 L2/L3 Fabric Greenfield Scenario

This scenario is an evolution of the previous one, and it is in line with the planned evolution of CORD itself. In this case, as Figure 23 shows, the Central Office packet layer includes a switching fabric that is capable of Layer 2 and Layer 3 functionalities. This enables some IP networking functionalities to reside on the switching fabric, whereas others reside on VNFs. IP functionalities with high computing requirements or that apply to a fraction of the traffic traversing the node may reside in VNFs. These IP functionalities may be in line with the functionalities that are implemented in service cards in current IP routers or with IP control functionalities of current IP routers.

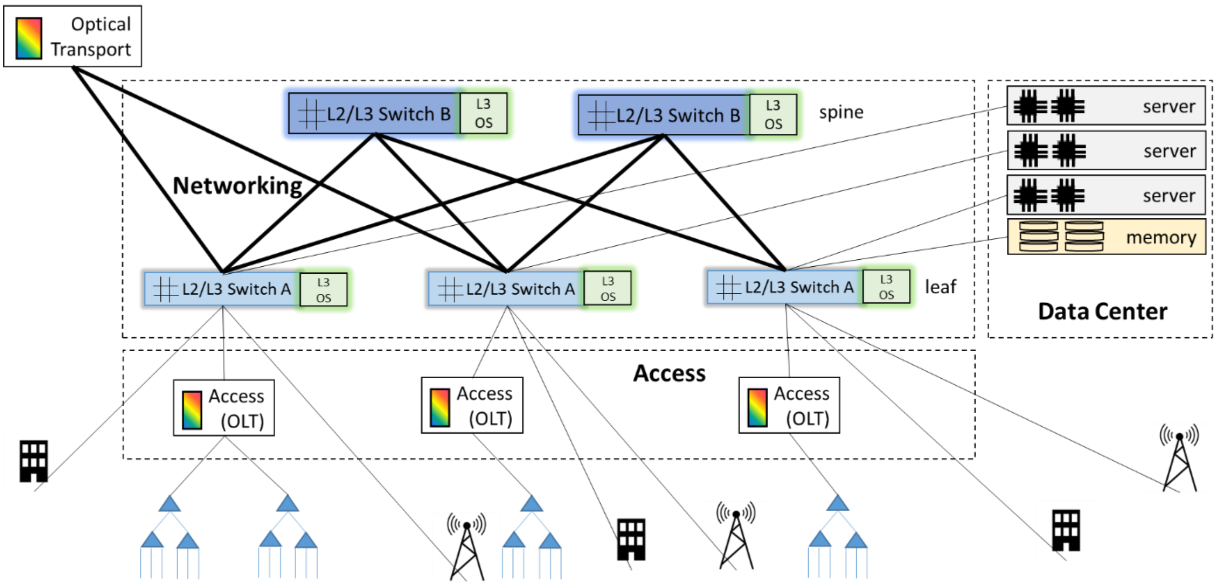

Figure 23: L2/L3 fabric greenfield scenario

4.1.1.4 L2/L3 Hybrid Fabric Scenario

This scenario, illustrated in Figure 24, could be considered a blend of L2 and L2/L3 fabric greenfield scenarios because there is no router IP considered, some part of the switching fabric stage is L2/L3 based, and some IP functionalities could reside on VNFs. Additionally, it should be taken into account, that it considers that part of the switching fabric stage can be L2 based, as used today in existing Data Centers. On the other hand, it allows for two "hemispheres" in the Central Office, the Telco-grade hemisphere devoted to replicating state of the art IP node functionalities, and the DC-grade hemisphere devoted to deploying Cloud-like services at the edge. At the control layer, there can be a separation of SDN control domains in line with the Broadband Forum TR-384 CloudCO Architecture: the DC SDN controller for the DC "hemisphere", and the PNF SDN controller for the Network "hemisphere".

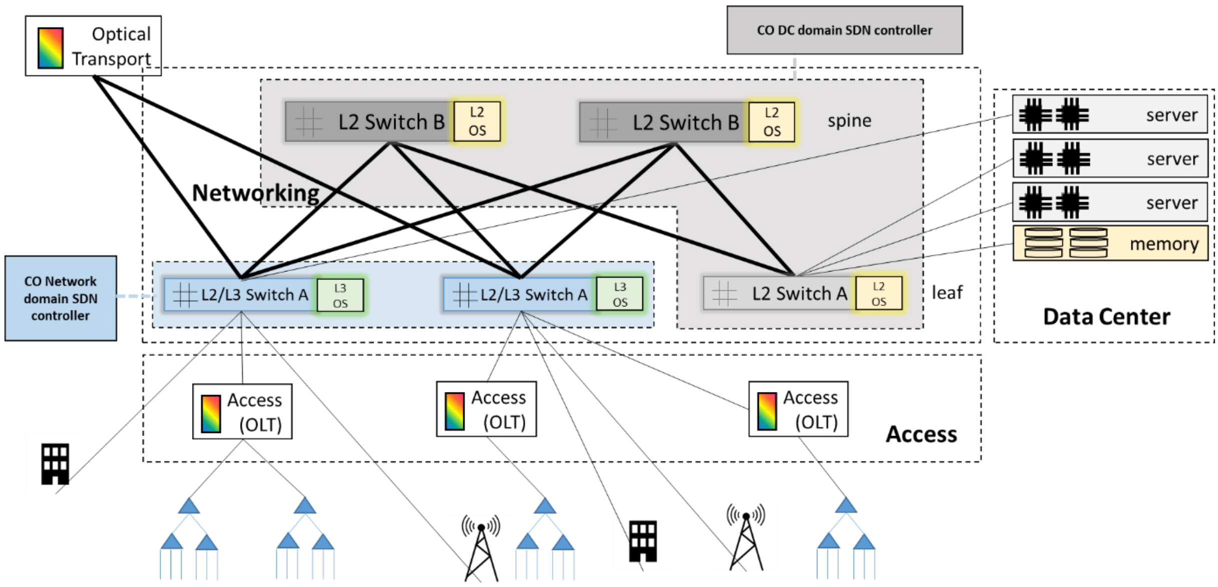

Figure 24: L2/L3 hybrid fabric scenario

4.1.1.5 L2/L3 Fabric Brownfield Scenario

In this type of scenario, as Figure 25 shows, the IP core integration functionality resides on traditional legacy routers, and the L2/L3 switching fabric takes care of the IP Data-Center local networking and the DCI features. Depending on the traffic pattern (N-S vs E-W distribution) connection to a legacy router could be at the spine or (border) leaf level.

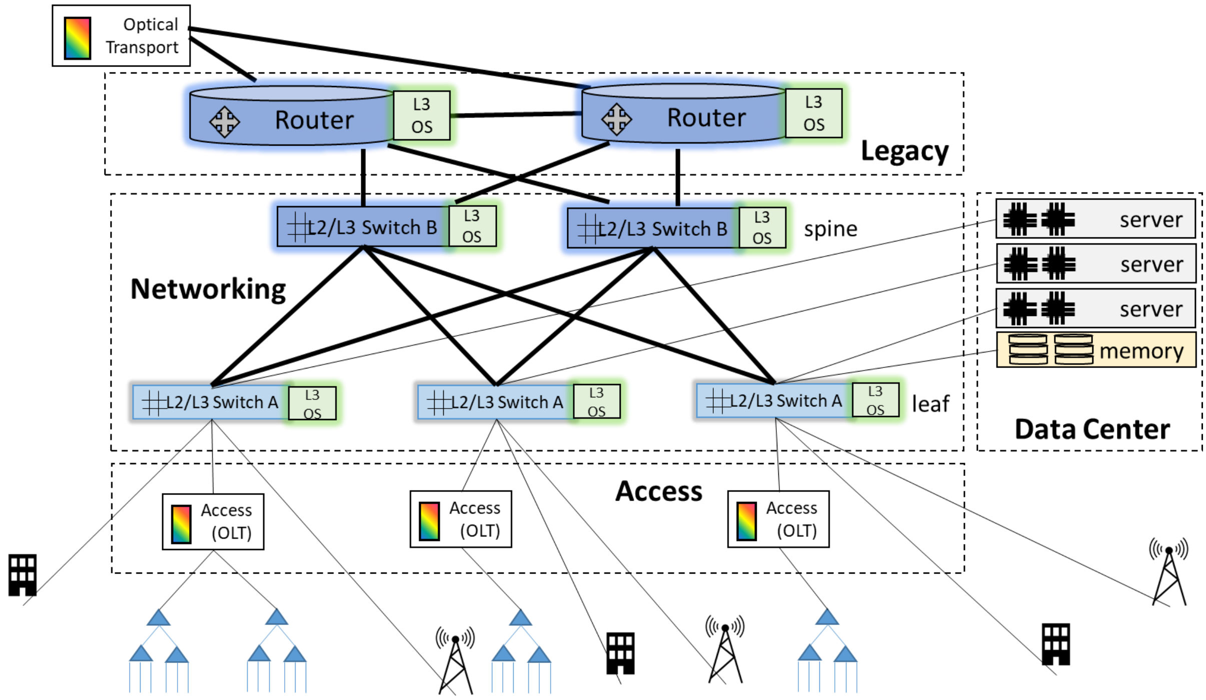

Figure 25: L2/L3 fabric brownfield scenario

4.1.1.6 L2 Fabric Brownfield Scenario

The last considered scenario, shown in Figure 26, differs from the previous one in the fact that the switching fabric (in this case) is composed by L2 switches instead of L2/L3 switches. This implies that the Data Center local networking will reside in the L2 switching fabric, and Core IP and Data Center Interconnect networking integration will reside on legacy routers. Similarly to the previous scenario, depending on the traffic pattern, connection to legacy routers could be at the spine or (border) leaf level.

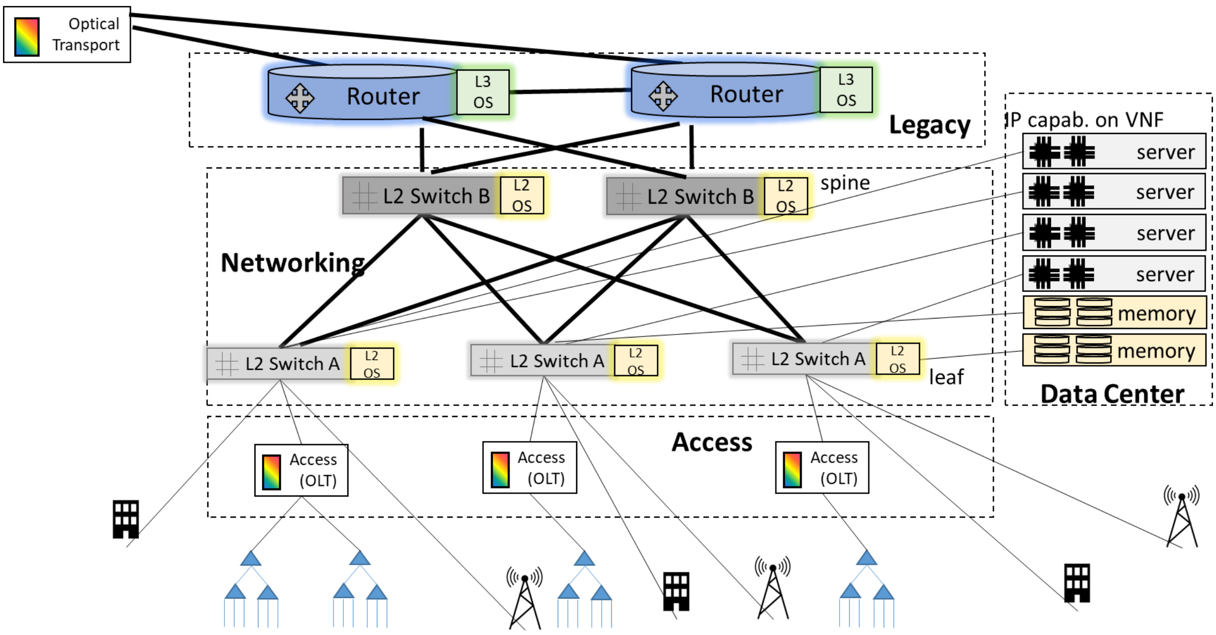

Figure 26: L2 fabric brownfield scenario

4.1.2 Metro-Haul Components

This subsection details specific devices developed within Metro-Haul that can be (wholly or partly) mapped to the candidate architectures already described. It should be noted, given that the techno-economic model is mostly founded on commercially available components, the experimental nature of the Metro-Haul specific devices may not include all the features required in such a system. However, the technical advantages that led to the development of these solutions are outlined. Further details on the specifications of these components may be found in [D3.1].

4.1.2.1 FPGA-based Programmable Switch + Bandwidth-variable transponder

The AMEN and MCEN edge nodes in the Metro Haul project connecting the access network with the metro network, and the metro network with the core network, respectively, include the programmable traffic aggregation, disaggregation and switching mechanism in their architectures for flexible network service provision. The SDN-enabled FPGA-based programmable switch is a candidate for this traffic aggregation and switching mechanism, since it aggregates data coming from heterogeneous access networks and data from local compute nodes, as well as PNFs and transmits data to other metro or core network nodes. Also, it can identify Ethernet packets with specified VLAN tags and transmit them to different destinations.

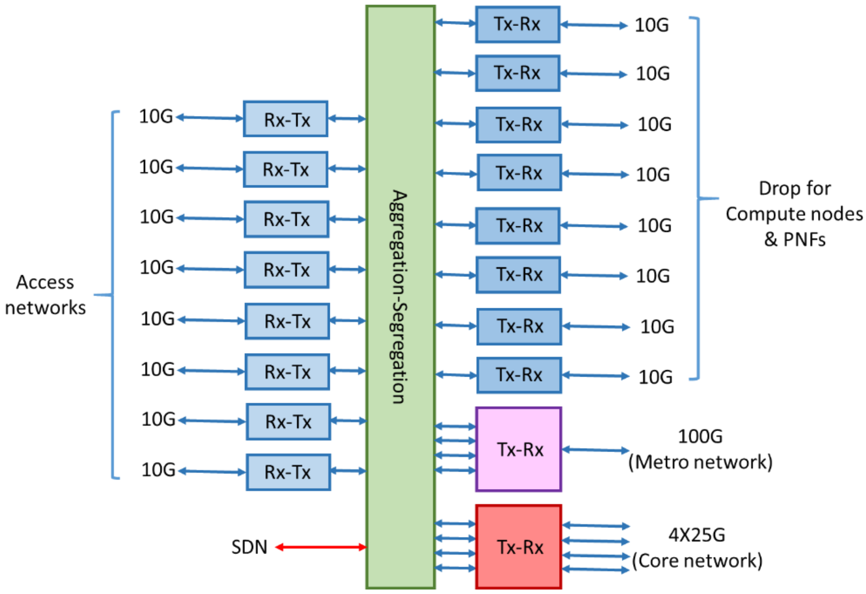

Figure 27: FPGA-based Programmable Switch Architecture

Figure 27 shows the architecture of the FPGA-based programmable switch. It provides multiple 10 Gbps ports that can interface with different access network technologies (Ethernet, PON, CPRI, LiFi, etc.) as long as the format of the received frames is Ethernet. Traffic from PONs is received as Ethernet from the OLT. LTE traffic is received from the base station in the eCPRI packetized format, the eCPRI traffic consisting of Ethernet frames carrying radio signals. Hence all the traffic received and transmitted to the access networks from the FPGA-based switch is based on Ethernet and can be segregated and identified via specific VLAN tags.

The edge node sends the traffic to either the core or the metro networks using a 100 Gbps transmitter. Additional 10Gbps ports are used to drop the data at the node towards the compute nodes and other PNFs. The ingress edge node is responsible for traffic aggregation and mapping, while the egress edge node has the reverse function. The edge node is fully integrated with an SDN control plane that can modify the properties of the edge node on the fly.

One possible mapping of this component within the Metro-Haul architecture is within the scope of the greenfield L2 fabric scenario outlined in Figure 22, where each L2 switch is hosted in a data-center. In the case of Metro-Haul, the data-center can be AMEN or an MCEN. However, the access networks (e.g., PON, LTE etc.) are aggregated on the L2 switch which is hosted in the AMEN. Here, the FPGA-based programmable switch can be mapped to the L2 switch and provide Layer 2 connectivity services. Considering the packet layer of Central Office (CO) architecture, there is a leaf and spine switching fabric which is where the FPGA-based programmable switch can fit. This switch can route traffic using VLAN tags as part of the Ethernet frame, where Layer 3 routing is handled by the VNFs themselves.

Since the FPGA is performing Layer 2 based traffic aggregation and switching, OpenFlow is a good candidate for the control of the FPGA-based programmable switch as a South-bound protocol by the packet SDN controller. Traffic between different VNFs and PNFs is identified using VLANs.

4.1.2.2 ADVA Computing Node

As a computing node for the AMEN and MCEN points of presence, the ADVA FSP 1050 ProVM platform as an L2/L3 demarcation device with an inbuilt server optimized for packet processing is ideal.

- **Hardware platform**

The FSP 150 ProVMe is a hosting platform combining both x86 server and Hardware-based acceleration and diagnostics in a carrier-grade platform. The physical data plane of the device defines the handoff of an Ethernet or IP communications service. The hardware data plane is enhanced to provide additional functions relevant to NFV. 4 GE ports are dedicated to access/client handoffs, 1 port is dedicated to the network connection, and 1 port can be configured to be either Access or Network. Single fiber (bi-directional) and dual fiber and copper Small Form-factor Pluggable (SFP) transceivers are supported. Also, the FPGA provides 4 further internal GE ports connected to the factory fitted internal server:

- 2x GE port used for customer traffic path. VLAN separation of individual EVCs addressed to the VMs (bidirectional).
- 2x GE ports used for mirrored traffic from Network Interface Device (NID) Port mirrors to the server (unidirectional).

The ProVMe provides a Xeon D-1539 8 Core 1.6 GHz (16 vCore) for NFV hosting, 16GB of Dual Channel 2133 MHz DDR4 SDRAM, and two SSD modules of 128GB and 512 GB respectively.

A common ARM-based supervisory processor monitors the complete solution, controls management access to both the server and the NID, and reports to traditional EMS/OSS systems. Figure 28 shows the generic hardware architecture:

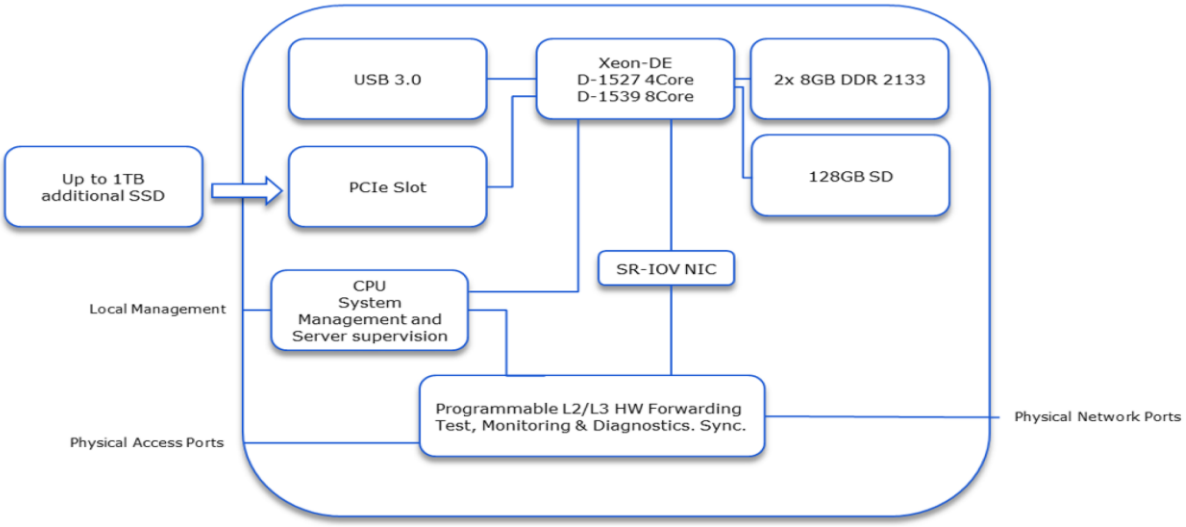

Figure 28: Compute Node Hardware Platform

- **Software Platform**

Ensemble Connector is a high-performance switching and virtualization platform for hosting multi-vendor VNFs. Because of its extensibility and modularity, specific data path functionality is supported for any deployment model at the customer premises, at the central office, or in the cloud data center.

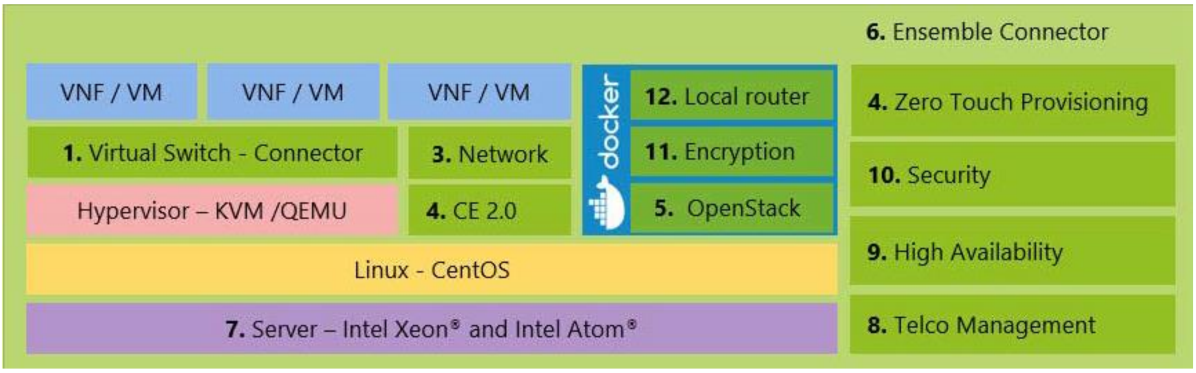

Figure 29: Compute Node Software Platform

1. Accelerated vSwitch
2. Carrier Ethernet 2.0
3. Networking, including LTE
4. Zero Touch Provisioning (ZTP)
5. Embedded Cloud (OpenStack)
6. Integrated OS with Open Interfaces
7. Device Scalability
8. Telco Management
9. High Availability (HA)
10. Platform Security
11. Encryption Engine
12. Local Router

The forwarding performance of Ensemble Connector utilizing DPDK hardware acceleration is faster, more efficient, and provides more consistent latency than open vSwitch.

Layer-3 Networking: Virtual Routing Functions (VRFs) are provided by Ensemble Connector that separate IP forwarding domains. The forwarding tables of a VRF can be built with static rules or dynamically through the border gateway protocol (BGP). Ensemble Connector natively supported local router functions. Each Ensemble Connector VRF can provide Network Address Translation (NAT) and DHCP server functions on designated VRF interfaces.

OpenStack Cloud: A self-contained embedded OpenStack cloud placed on the edge node can enable cloud-native deployments without the issues created when separating the OpenStack controller from its agents.

High Availability: compute devices running the Connector software platforms can be arranged into high availability configurations with an active/standby instance using Virtual Router Redundancy Protocol (VRRP) from hosted VNF configurations.

Virtual Machine Support: the compute nodes provide a virtualization environment with the following features:

- KVM/QEMU – KVM (Kernel Virtual Machine) that supports hardware virtualization while QEMU (Quick Emulator) emulates the target operating system.
- Standard OpenStack compute node APIs (Nova, Neutron, Ceilometer, Cinder).
- Standard OpenStack controller node APIs for:
- Glance – stores and retrieves virtual machine images.
- Swift – stores objects (optional).
- Keystone – Authenticates and authorizes service.
- Nova – Manages the life cycle of compute instances. Schedules, spawns and decommissions VMs.
- Neutron – connects networks as a service.
- Cinder – provides persistent block storage to running instances.
- Ceilometer – monitors and meters the OpenStack cloud.
- Heat – orchestrates OpenStack.
- Horizon – installs the OpenStack UI for debug usage.

4.2 Physical Network Architectures for Optical Nodes

The deployment options for metro transport networks in a 5G-enabled setting can be quite varied, given the range of service profiles and geographic coverage combinations. If the densification of fiber rollouts is reasonably understood to be necessary to increase capacity, there is much more debate on how to deploy them in a cost-efficient but scalable way. Within the Metro-Haul scope, the roles of the AMEN/MCEN nodes are sufficiently flexible to cover front-haul/back-haul applications and different functional split options for the RAN [D2.2]. The integrated optical/compute node, with joint orchestration of layers and domains, provides the control-plane flexibility necessary to flexibly implement the various options. However, the specific architecture of the optical nodes and their interconnection can also take many forms concerning cost, flexibility, scalability, resilience and several other factors. This subtask looks into greater detail at the types of architectures that can be used to deploy the optical transport part of AMENs/MCENs. This includes both the transmission technology and the optical switching architecture employed at the nodes. The main goal is to identify, based on the tiered network model in Section 3.1 and the service categorization in 3.2, the most suitable technologies for deployment in each set of scenarios.

To this effect, this section details the candidate physical architectures envisioned for metro aggregation and metro core networks. These candidate options generally present a trade-off between cost and capacity/flexibility that make each option a desirable sweet spot for particular applications. The identification of these scenarios is done by running the different optical node architectures through optimization and dimensioning frameworks aligned with the geographical and service profile assumptions outlined in this deliverable. Therefore, this section also presents, from a methodology perspective, the type of optimization tools that will be used in the techno-economic analysis of the physical network architectures, with examples of the type of results that can be expected as an outcome. Combining the node coverage geotypes, the service characterization and traffic modelling, and the cost modelling of the individual network components, these optimization tools should be able to provide an accurate assessment for deployment options concerning optical transport architectures.

Finally, this section also identifies the specific hardware components developed within Metro-Haul which advance the state-of-the-art in terms of transmission and optical switching technologies. These components are generally mapped to the candidate architectures under consideration, also indicating in what context they are expected to improve over existing solutions (e.g. lower cost or power consumption, higher capacity, etc.). It should be noted, however, that this mapping should not be expected to cover the full functionality that such components would have in a commercial production setting, given their development stage and level of maturity of the technologies involved.

4.2.1 Candidate Node Architectures for Metro Aggregation Domains

As outlined in Section 3.1.2, the reference topologies considered are hierarchically divided into metro aggregation and core domains. In the metro aggregation segment, a set of AMENs are interconnected in a chain topology, with two MCENs at the ends providing an interface to the metro-core domain. The metro aggregation domain, being more prevalent and interfacing more directly with access/end-user traffic (even though MCENs also accumulate this role), typically require a more cost-driven node design that exploits efficiencies of scale, even if at the expense of the higher performance/capacity required higher up in the hierarchy. This compromise can be achieved both in the cost/capacity of the optical interfaces used for transmission, and on the complexity of the switching architectures. Table 28 outlines the main candidate options to be explored in this context, with respect to the express architecture of the optical nodes, the add/drop architecture and the transceiver options.

Table 28. Candidate Architectures for Metro Aggregation Domains.

Express Node Architecture	Add/Drop Node Architecture	Transceiver Technology
Fixed-Frequency Filter Filterless (DuFiNet, Drop & Waste) Semi-Filterless 2-degree B&S ROADM	Splitter/Coupler Splitter/Coupler Mux/Demux	Coherent (CFP2-ACO) Direct-Detection

A full candidate architecture is a result of combining each option for express, add/drop and transceivers. Modules can be freely combined although some restrictions apply. For instance, all express architectures except fixed-filter support add/drops based on splitters or multiplexers. However, the specific maximum add/drop channel counts to be supported, together with noise accumulation and power budget constraints determine whether a splitter-based approach is still feasible or filtering in the add/drop layer is required. The transmission technology is also generally compatible with any mix of express and add/drop architecture, although the resulting performance (impacting reach, maximum channel counts, amplification strategy, etc.) is affected. As such, the functional role of the switching node is mostly defined by the express architecture, with the add/drop and transceiver technology scaling the achievable number/capacity of optical channels. Therefore, greater focus is given to express node architectures on the overview of the candidate models that follow.

4.2.1.1 Fixed-Frequency Filter Nodes

A node based on fixed-frequency filters, illustrated in Figure 30, is assumed as the baseline in our techno-economic comparisons. In this architecture, each node can only add/drop a fixed set of frequencies that were preset upon the network's deployment. Hence, the transparent connectivity in the network (i.e., the set of nodes that can establish direct connections) is fixed during its lifecycle. This type of scenario is common in legacy metro aggregation networks for simpler chain/ring topologies, where each node is only expected to communicate with the aggregation hub node.

The simplicity of the node's architecture is inversely proportional to its functionality. Concretely, such a design either does not allow direct connections between AMENs in the same chain, or enables only a very limited subset of frequencies to be reserved for such connections. In either case, any application requiring more flexible traffic patterns (e.g., cases where information handover is required between applications in adjacent AMENs) is highly constrained, as traffic must be routed through an MCEN, which increases the overall bandwidth within the chain and introduces additional latency. Moreover, the fixed network design leaves little room for the network to absorb variations in the traffic profile. One tenet of novel 5G services is localized traffic hotspots that can generate bandwidth peaks due to, e.g. events, emergencies. In this architecture, the available bandwidth must be pre-apportioned to each AMEN, reducing overall network flexibility. Finally, this node design is incompatible with flexible-grid designs, since it also defines a fixed spectrum window for each channel, making it harder for an operator to leverage adaptive transceivers that can tailor the transmission format to the characteristics of a given optical path [VLope16].

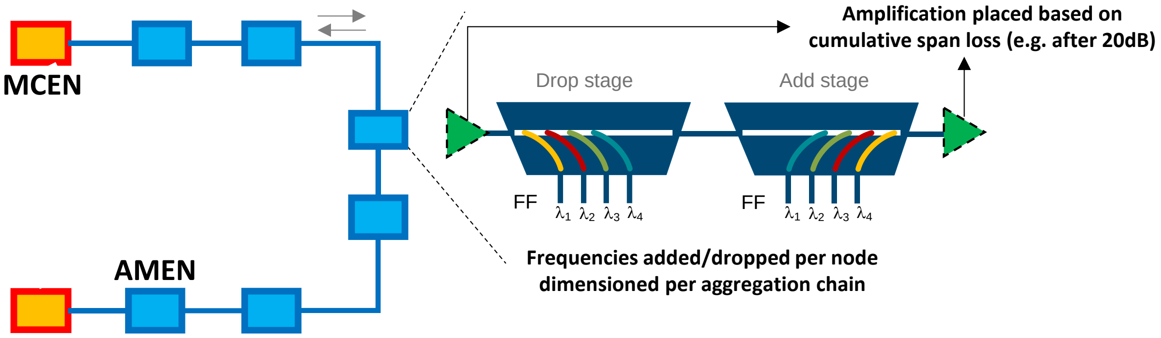

Figure 30: Fixed-frequency filter based express architecture.

4.2.1.2 Filterless Nodes

One option being championed within Metro-Haul is the possibility of using a simple splitter/coupler based express architecture. This option retains a low-cost design for optical nodes, while adding flexibility for connections that is otherwise absent from a fixed-filter configuration. In this design, all channels are expressed across the metro aggregation chain, being filtered only at the edge MCENs. Here, the connectivity problem inherent to a fixed-frequency filter design is solved, since any-to-any lightpaths can be freely established. The main drawback of this architecture lies with the inability to reuse the same frequency for different lightpaths, since interference between the unfiltered signals is not manageable.

On a high-level, the filterless node design thus improves on network flexibility, but possibly at the expense of a lower overall throughput, depending on specific traffic patterns. However, additional considerations must be taken into account for an in-depth analysis. For instance, a filterless architecture tends to introduce lower express losses for channels, which can enable the use of higher capacity channels in critical applications, or reduce the amplification requirements. Depending on the channel profiles used, this option may also reduce impairments due to filter cascading [JMFab16]. On the other hand, the lack of filtering also implies that noise levels accumulating through links and at the receiver are higher, which may offset performance in a specific scenario (e.g., longer links).

In the scope of this sub-task, two main flavors of filterless nodes are evaluated. The first is the standard drop and waste (D&W) architecture, shown in Figure 31, where splitters/couplers are used in a traditional fiber-pair configuration. A 3dB splitter/coupler is used to drop/add signals. On the add/drop, another splitting/coupling stage can be used for lower channel counts (where splitting losses do not imply drop amplifiers are needed), whereas arrayed waveguide grating (AWG)-based structures may be used for higher add/drop capacities (e.g. in smaller chains). Additional constraints may also apply to the transceivers used, due to noise accumulation at the receiver. In the case of coherent transponders, tunable broadband filtering is assumed as a compromise solution, limiting the total received power to acceptable levels without imposing an overly narrow filtering function for each channel.

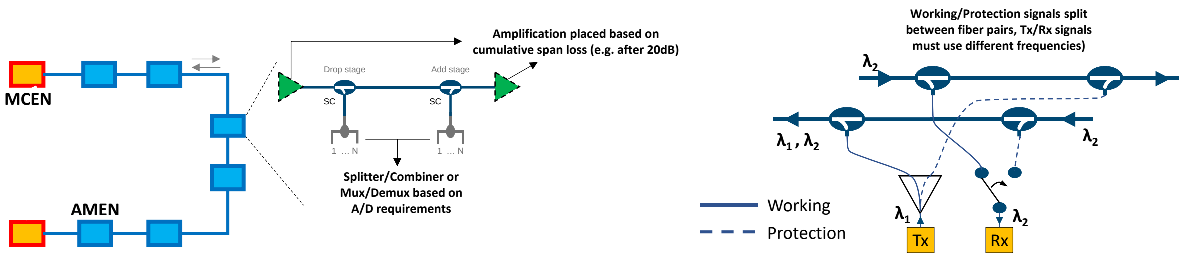

Figure 31: Filterless drop & waste express architecture and protection requirements.

The other option for a filterless deployment is the Dual Fibre Network (DuFiNet), presented in greater detail in Section 4.2.1 of [D3.1]. In this instance, each fiber is dedicated to either downstream or upstream transmission (from the MCEN hub to each AMEN). Thus, each fiber has either only a drop- or add-side, which still features the same splitter/combiner structure, as shown in Figure 32. From a functional perspective, this architecture can only provide AMEN-AMEN connectivity through an MCEN (with added latency). However, performance-wise, it has been shown to support high-capacity transmission formats and co-existence with legacy NRZ formats.

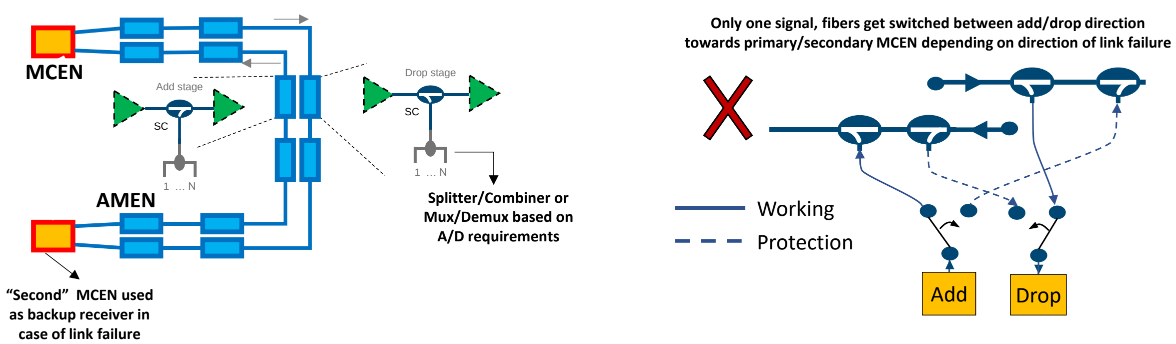

Figure 32: Filterless DuFiNet express architecture and protection requirements.

4.2.1.3 Semi-Filterless Nodes

The semi-filterless node is based on a wavelength-blocker solution, which replaces the WSS used in a more traditional Broadcast & Select (B&S) ROADM. The basic architecture for the express node is shown in Figure 33. The wavelength-blocker provides the ability to choose (at a given frequency granularity) which frequencies can be forwarded to the next span. Hence, in a 2-degree node such as those required for metro aggregation chains, the component is functionally equivalent to a WSS, since the add channels can use a coupler after the wavelength-blocker. On the other hand, relative to a B&S ROADM, there is an additional splitter/coupler in the express path, and the add/drop structure cannot be upgraded to colorless/directionless designs, given that there's no element to actively select the desired/unwanted add/drop channels at each degree.

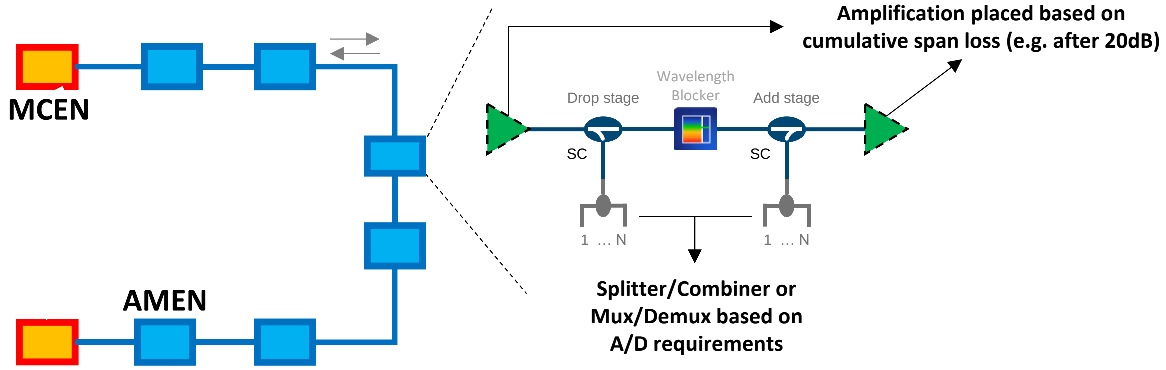

Figure 33: Semi-filterless wavelength-blocker based express architecture.

4.2.1.4 2-degree B&S ROADM

The most flexible (and complex) architecture considered for the metro aggregation domain evaluation envisions the use of a B&S structure, with a splitter connecting to the express and drop paths, while a WSS filters the appropriate channels coming from the upstream node and the add section (see Figure 34). This architecture provides any-to-any connectivity with frequency reuse capabilities, hence maximizing both capacity and flexibility. However, it requires a costlier active component, and can introduce filtering penalties for lengthy cascades of ROADMs in channels with tight filtering windows. On the other hand, compared with filterless scenarios, it limits the amount of accumulated noise through the filtering at each node. It should be noted that, although the WSS is comparatively costlier than a fixed-frequency filter, splitter/combiner or wavelength-blocker based solutions, the limitation to two-degrees in this domain allows the use of low port-count pluggable WSS modules, which bring down the overall component cost and footprint.

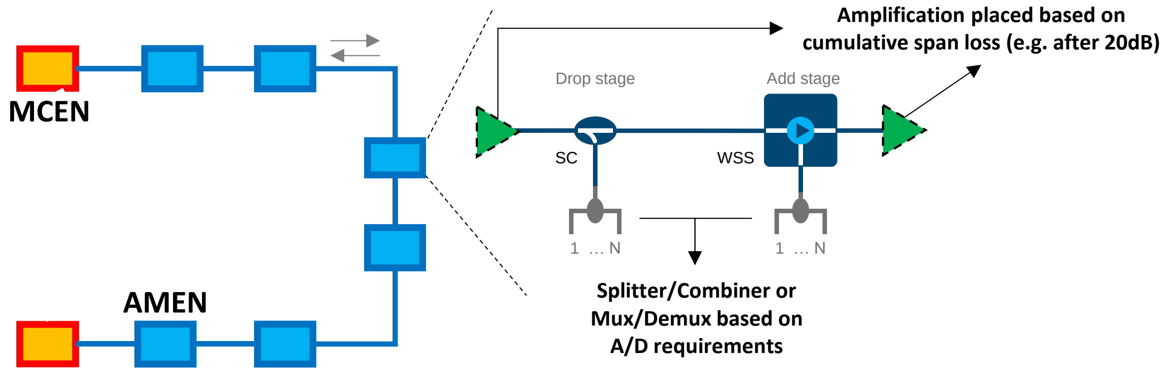

Figure 34: B&S ROADM express architecture.

4.2.1.5 Methodology and Early Assessment

Based on the possible node architectures described in the previous subsections, an initial techno-economic evaluation considers the hardware requirements needed to support various traffic levels. The main vectors of analysis in this framework are the capacity of the infrastructure, its flexibility and its cost. Capacity is simply evaluated as the overall throughput/bandwidth supported by a metro aggregation domain as a whole (in AMEN-MCEN or AMEN-AMEN connections). Flexibility refers to how efficiently the architecture the reconfiguration of traffic profiles and bandwidth fluctuations. Cost results from these two aspects, indicating the amount of hardware needed not only by the switching nodes themselves (splitters, WSS, mux/demux, etc.), but also the transceiver costs incurred by each option (which tend to vary also as a result of the network's flexibility).

This evaluation is comprised of two main components: a multi-period optimization framework, which calculates the optimal number of transceivers required to serve a given traffic matrix (reusing whatever hardware was already previously deployed at a node), and an optical performance model which estimates the achievable lightpath capacity given the characteristics of a path in the metro aggregation chain (number of hops, span lengths, insertion losses of components, amplifier noise figures, etc.). The traffic matrix is a result of combining background traffic (typically hubbed between AMEN-MCEN), with more dynamic patterns arising from chaining of VNFs within the metro aggregation domain, as exemplified in Figure 35. Further details on the operation of this framework can be found in [AEira18, JPedr18].

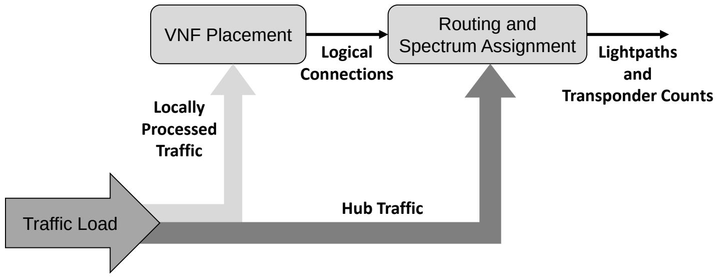

Figure 35: Workflow for evaluation of line system requirements per express architecture.

Figure 36a shows example results for the transponder counts (assuming coherent interfaces) for express architectures based on B&S ROADM, filterless D&W and fixed-frequency filters. The analysis is made for varying chain profiles (total length and number of nodes), and also for varying traffic profiles, where the share of traffic related to VNF chaining (and hence more meshed in the pattern) progressively increases. These early results confirm the high-level analysis of each architecture, suggesting the ROADM-based designs are the benchmark in terms of capacity and flexibility (although the architecture cost, which is not accounted for yet, is higher). The filterless designs are limited mostly by raw capacity due to frequency reuse, particularly when there are more nodes in the chain. The fixed-frequency filters, on the other hand, are mostly limited in terms of flexibility, being unable to either serve meshed traffic altogether, or being forced to route it through multiple hops. As a consequence, illustrated in Figure 36b, the architecture also plays a role in the achievable service latency, as the more constrained architectures are forced to use longer paths (longer propagation latency) and additional lightpath hops (more optical/electrical/optical conversion latency).

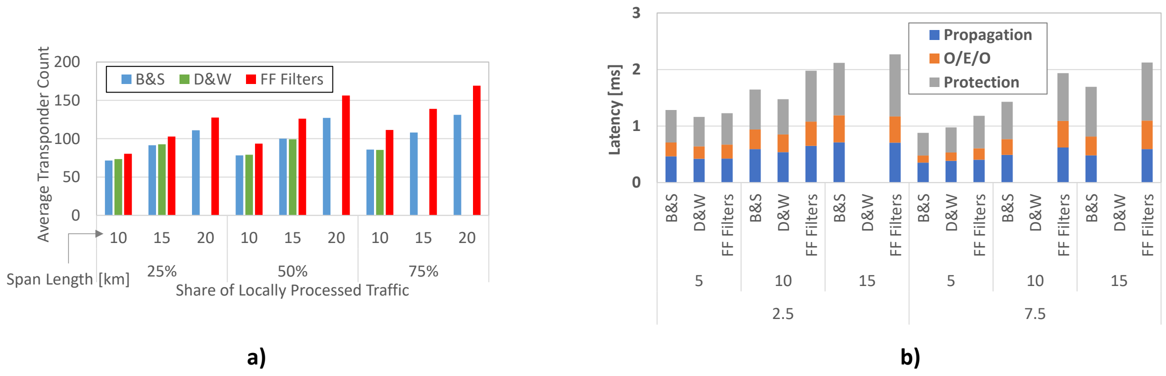

Figure 36: Evaluation of metro aggregation requirements per express architecture: a) transponders (10 node chain); b) latency.

In the future scope of this subtask, this framework is to be harmonized with the support tasks in Section 3. Namely, the metro aggregation chain profiles need to be based on the particular geotype definition considered. The traffic assumptions (low-granularity fluctuation and categorization) must be reworked to consider the model in Section 3.2.3 and the evolution in Section 3.2.2. Additionally, the cost of the transceivers must be weighed against that of the express (and add/drop) architectures, based on the optical cost model in Section 3.4. Finally, additional architectural options are to be considered, namely including filterless DuFiNet flavors, semi-filterless wavelength-blocker based design, and options with direct-detection transceivers.

4.2.2 Candidate Node Architectures for Metro Core Domains

The Metro-Core domains, as described in Section 3.1.2, comprise several MCENs interconnected in a mesh pattern. Some of these MCENs (though not necessarily all) serve a dual role as hubs for metro aggregation chains and switching nodes in the metro-core. Additionally, some of these MCENs may also act as AMENs in the sense that they provide multi-access traffic aggregation and local computing functions within their area of coverage.

In this scope, metro-core nodes tend to have higher traffic volume requirements and more complex traffic patterns than those needed for metro aggregation domains. This is expected to be reflected in the architectures that would be suitable for deployment in the optical nodes, as well as in the capacity requirements for transceivers. As such, analogous to the options in Table 28, the candidate building blocks for optical nodes in the metro core domain are outlined in Table 29. There are generally fewer options than in the metro aggregation domain, since the requirements tend to be more stringent and there is less variability as one moves up the network hierarchy (i.e., additional levels of aggregation tend to harmonize the per-node requirements).

Table 29. Candidate Architectures for Metro Core Domains.

Express Node Architecture	Add/Drop Node Architecture	Transceiver Technology
Multi-Degree B&S ROADM (low-cost LCoS WSS or SOA-based ROADM) Semi-Filterless	Fixed (AWG-based, fixed-grid)	Coherent

	Colorless / Directionless (fixed and flexible-grid)	

Support for meshed topologies requires that either the nodes support advanced (WSS-based) channel switching, or the network is optimized to provide switching functions only at select nodes, enabling network subdomains to operate with a restricted topology (e.g. based on filterless nodes). This hierarchical dimensioning is further described in Section 4.2.3. The add/drop structure must also be considerably more advanced and complex than in the metro aggregation domain, as it is expected that more advanced protection/restoration techniques are employed, requiring flexibility to use add/drop ports to address different frequencies/directions. In this scope, it is also envisioned that the higher traffic requirements will necessitate the move to a flexible-grid scenario where highly granular channel formats trading reach for spectral efficiency can be leveraged for maximum benefit [FGuio17]. On the transceiver side, given the higher capacity per interface relative to metro aggregation, as well as the comparatively longer transparent connection requirements, it is assumed that deployments will use coherent transponders (e.g. based on CFP2-ACO pluggable formats).

Within the option for multi-degree B&S ROADMs, the scenarios can comprise different deployment options. Recent approaches focus on Liquid Crystal on Silicon (LCoS) WSSs, which tend to be suited for flexible channel spacings. Within Metro-Haul, alternative approaches are also being addressed, such as a B&S degree where the egress is based on combining AWGs with SOAs, as an integrated switching and power equalization solution [NCala17]. This approach may provide benefits in terms of reduced cost and power consumption in select applications, although specific limitations do apply (e.g., the AWG-based structure assumes a fixed-grid channel width must be used).

4.2.3 Hierarchical dimensioning of semi-filterless networks

In this section we describe the design approach used to perform comparison of filterless vs. semi-filterless solutions in a hierarchical metro-area network. The design approach is structured in two steps. As a first step, we use a heuristic algorithm for the design of filterless optical networks based on dividing the network into edge-disjoint fiber trees, where each fiber tree denotes a set of interconnected optical fiber links and defines the interconnection between the passive splitters and combiners. As a second step, we perform Modulation format, Routing and Spectrum Assignment (MRSA) of traffic demands with the objective of minimizing the overall network spectrum using an Integer Linear Programming (ILP) model (overall framework is shown in Figure 37).

First step: Given the network topology (with the location of filters and filterless nodes in the network) and the traffic demands, our proposed heuristic algorithm establishes edge-disjoint fiber trees taking into consideration the “laser loop” and the “connectivity” constraints. The laser loop constraint guarantees that no closed loops are present in the interconnection of the fibers, while the connectivity constraint guarantees all traffic demands can be accommodated. The output of the heuristic approach consists of:

i) Established fiber trees
ii) Possible routes (fiber trees) each traffic demand may be assigned

and is then passed as an input to an Integer Linear Programming (ILP)-based module (as shown in

Figure 37).

Second step: An ILP model is used to solve the MRSA problem, which is formulated as follows: Given the set of fiber trees, the traffic demands, the possible route assignment per traffic demand and the set of modulation formats, we assign a modulation format and a route for each traffic demand and allocate the required frequency slot units with the objective of minimizing the maximum frequency slot unit (FSU) utilized among all links and/or the overall network spectrum utilization. The MRSA in filterless networks is subject to the following constraints:

i) Each traffic demand is assigned exactly one route and one modulation format.
ii) Each traffic demand is assigned a sufficient number of FSUs, with a starting FSU and an ending FSU, guaranteeing FSU's contiguity constraint (consecutiveness of FSUs) and continuity constraint (an FSU on a link is assigned to only one demand). Note that the number of FSUs per demand is a variable of the problem as it depends on a demand's bandwidth request and the modulation format assigned and thus,
iii) if an FSU is occupied on a link of a given fiber tree to accommodate a demand, the FSU is considered reserved on all links of the fiber tree, due to the broadcast nature of filterless networks, and cannot be assigned to any other demand passing through the links of the fiber tree.

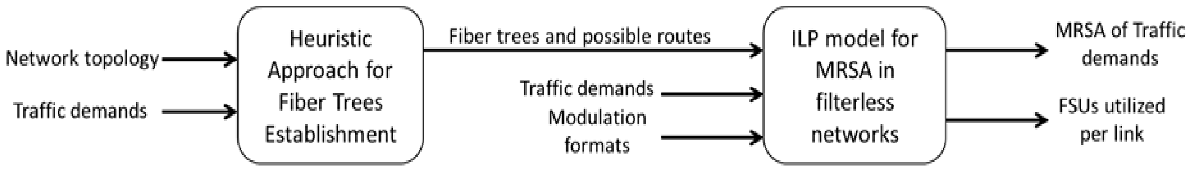

Figure 37: The overall framework of the filterless network design and resource allocation.

4.2.3.1 Network and Traffic Models

To perform the evaluations, we consider a metro network topology spanning over various hierarchical levels similar to the one in Figure 38 consisting of 4 categories of nodes: AMENs, MCENs, Metro Nodes (MN), and Access Nodes (ANs). The AN is where traffic of fixed and mobile end-users is aggregated. The AMEN represents a central office where the access head-ends and the metro network interfaces are located. We assume the AMEN supports multiple access technologies and aggregates traffic from ANs. The MN represents a pure transport metro node. The MCEN is where the metro head-ends and the core network interfaces are located.

The AMENs and MCENs serve as mini and regional data centers, respectively, i.e., they are equipped with computing and storage capabilities and provide latency-stringent and bandwidth-hungry services (e.g., 4K Video-on-Demand delivery, Virtual Reality) requested by end-users. We consider 40 ANs, 20 AMENs, 6 MNs and 2 MCENs. The ANs are interconnected to AMENs in tree-like topologies. The AMENs are interconnected in a ring topology to MNs, which are interconnected in a meshed topology to MCENs. We consider a single bidirectional fiber where each fiber direction is used either for upstream or downstream traffic. Moreover, we consider 4 different network deployments, namely;

1) ***Active-Photonic***, where all nodes are equipped with WSS-based ROADMs.
2) ***Filterless***, a fully-filterless solution where both AMENs and MNs are assumed to be filterless.
3) ***Filterless MNs***, a semi-filterless where only the MNs are considered to be filterless nodes whereas other network nodes are active.
4) ***Filterless AMENs***, a semi-filterless solution where only the AMENs are filterless nodes.

We compare the proposed network deployments in the context of a Content Delivery Network use case (i.e., Video-on-Demand service) as described in [D2.1]. Specifically, we assume AMENs and MCENs are equipped with caches storing video contents, where first-tier popular contents are stored at AMENs and second-tier popular contents are stored at MCENs, and then generate traffic accordingly. We assume each AN requests 150 Gbps of traffic on average (equivalent to 10,000 video content requests with an average bit-rate of 15 Mbps per request), summing up to a total of 6 Tbps of overall network traffic. We also assume content is delivered to ANs from the nearest location that it is cached, i.e., either from the nearest AMEN or the nearest MCEN.

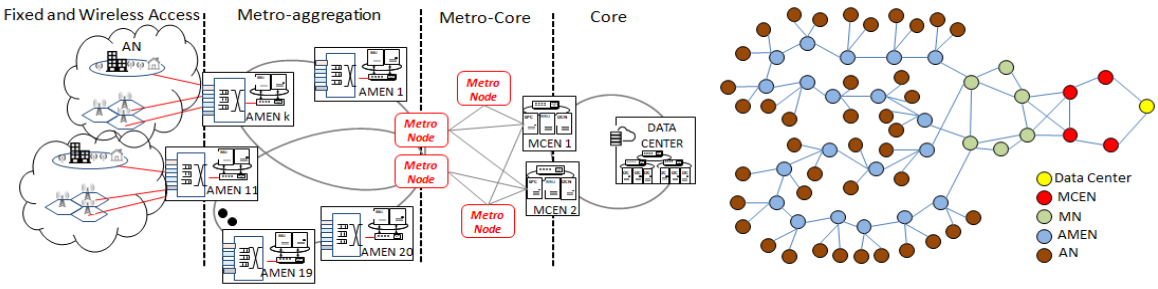

Figure 38: A schematic representation of the network topology showing the 4 network levels (left) and the network topology considered in our study (right).

As for the transponders considered, we consider two coherent transponders with Dual-Polarization (DP) that operate at line rates of 100 and 400 Gbps, with a DP-QPSK modulation format with 37.5 GHz channel bandwidth and DP-16QAM modulation format with a 75 GHz channel bandwidth, respectively. We assume a spectral granularity of 12.5 GHz FSU, and consider 200 FSUs (i.e., a total of 2.5 THz) are available on each fiber, also including 1 FSU guard band between the optical channels.

4.2.3.2 Illustrative Numerical Results

In this section, we present illustrative numerical results, showing the formation of the fiber-trees in the proposed network solutions (i.e., the output of the heuristic approach for filterless network design) and comparing the four proposed network solutions in terms of spectrum utilization.

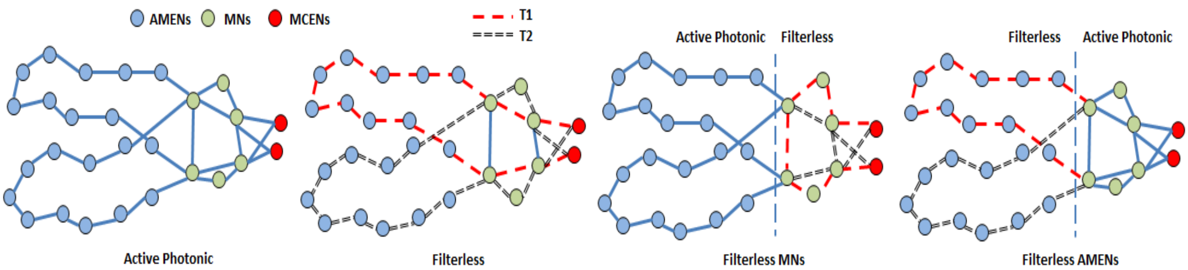

Figure 39 (From left to right) The network topology for the Active Photonic and the fiber tree formation for the Filterless, the Filterless MN and the Filterless AMENs. Note that the ANs are not sketched as they act as leaves and belong to the same fiber.

Fiber Trees Establishment: Figure 39 illustrates the network topology for the *Active Photonic* and filterless solutions. For each of the *Filterless*, *Filterless MNs* and *Filterless AMENs* network deployments we obtained 2 fiber trees. The 2 fiber trees (T1 and T2) in the *Filterless* deployment span all network levels and connect the AMENs to the MCENs. In the *Filterless MNs* deployment, the 2 fiber trees guarantee network connectivity between MNs and MCENs, while in the *Filterless AMENs* deployment the 2 formed fiber trees create a 'horseshoe' topology (to avoid laser-loops), connecting AMENs of the same ring to two different MNs.

Now we compare the spectrum usage of the 4 network deployments in terms of the maximum allocated spectrum slot number (*max FSU*), and the overall network spectrum utilization (*total spec.*), i.e., the total number of FSUs utilized.

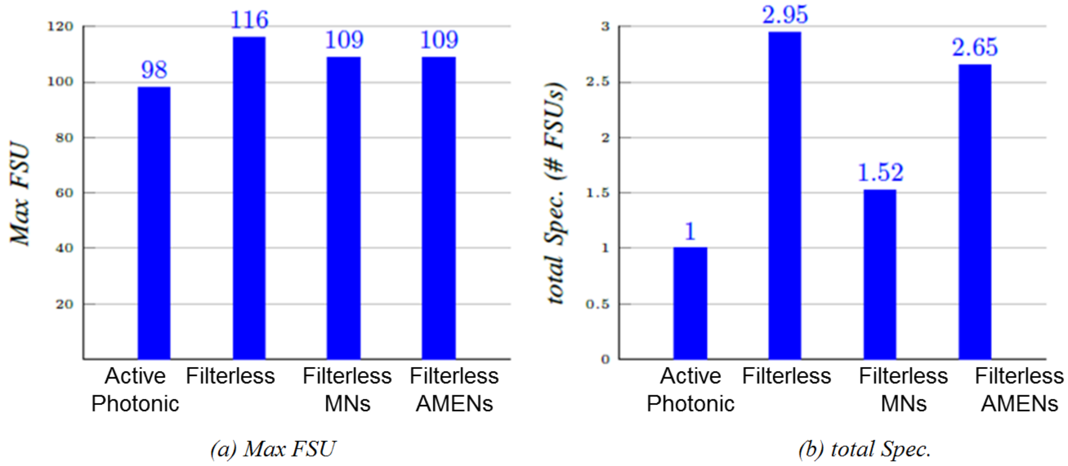

Figure 40: Comparison of the network deployments in terms of (a) Max FSU and (b) total Spec. (normalized to that of Active Photonic).

Spectrum Utilization Comparison: Figure 40 (a) and (b) show *max FSU* and *total spec.* of the four network deployments. On one hand, we see that the *max. FSU* values of the filterless solutions are comparable to that of the active photonic (only around 15% difference), while on the other hand, the *total spec.* values of the filterless and semi-filterless solutions are significantly higher (around 2 to 3 fold) than that of the active photonic solution. This reveals that the *max FSU* values are a result of congestion of the links which are forced to carry most of the traffic from the MCEN to the nearest AMEN of the AMENs ring. In other words, the high max FSU values are not due to widespread congestion among all network links, but only because a few links suffer from high utilization. This also shows the *max FSU* cannot be considered a reliable metric to measure the overall spectrum utilization in hierarchical ring-based metro networks, where specific links are forced to carry most of the traffic. As for the *total spec.,* we see that the fully-filterless solution has the worst performance (highest amount of required FSUs). This is due to the length of the fiber trees which span all the hierarchical levels of the network causing excessive signal broadcasting.

Similarly, filterless AMENs solution shows a high *total spec*. with only a slight improvement concerning the fully-filterless in spite the MNs being equipped with filters. This is because the presence of filters at the MNs network level helps reduce the spectrum utilization by preventing signal broadcasting among the MNs, but then a high spectrum is utilized once the signal propagates into the AMEN rings due to the filterless fiber tree connecting the AMENs. As for the Filterless MNs solution, it shows the best performance in terms of spectrum utilization (lowest *total spec.*) as the AMENs are equipped with filters, and thus no signal propagation of a request occurs beyond the AMEN to which the destination-AN is connected. This also shows that the major part of the wasted spectrum is in the ring interconnecting the AMENs, and thus a strategic deployment of filters is motivated among AMENs more than among MNs.

In summary, the results suggest that filterless AMENs, due to their location in the considered network topology, cause a significant waste of spectrum. An alternative solution is to equip AMENs at strategic locations in the network topology with filters and wavelength blockers, such that the number of signals being broadcast is narrowed. Moreover, because the MNs are meshed-connected and that the traffic between them follows a specific path (e.g., incoming traffic from an MCEN is destined to the nearest AMEN), spectrum waste reduction is possible among filterless MNs through the use of programmable optical switches. These aspects are currently under investigation to discover other possible advantages of filterless architectures in hierarchical metro networks.

4.2.4 Metro-Haul Components

This subsection details some of the physical components, such as wavelength blockers and direct-detection transceivers, as well as the options are chosen amongst the devices developed within Metro-Haul which act as reference implementations in the candidate architectures, and which are therefore included in the techno-economic modeling. Further details and specifications regarding these components may be found in [D3.1].

4.2.4.1 Programmable Sliceable Transceiver Based on Multicarrier Modulation (MCM) - CTTC SBVT-WP3 component

A programmable sliceable transceiver based on multicarrier modulation (MCM) with a receiver configuration adopting direct detection (DD) is designed and implemented within the Metro-Haul project (here a brief description is provided for the sake of clarity, further details can be found in the specific WP3 deliverables). Offline digital signal processing (DSP) is considered. Two bandwidth/bitrate variable transceiver (BVT) modules are enabled to transmit a high capacity single-flow through the network. Additional slices/modules can be included to further enhance system capacity. Cost-effective optoelectronic subsystems and simplified DSP solutions must be adopted to address the stringent cost target of metro network infrastructure (related to MH8 and MH9 KPIs). Thus, a simple front-end, which includes Mach-Zehnder modulators (MZMs) driven by tunable laser sources (TLSs) are included at the sliceable bandwidth/bitrate variable transmitter (S-BVTx). At the receiver side, simple photo-detectors (PINs) with transimpedance amplifiers (TIAs) are used to perform DD. Wavelength selective switches (WSSs) are included at the transceiver to aggregate/distribute the different flows/slices. These elements can also serve as optical filters to perform single side band modulation (SSB), enhancing robustness against fiber impairments such as chromatic dispersion. The architecture of the device is shown in Figure 41.

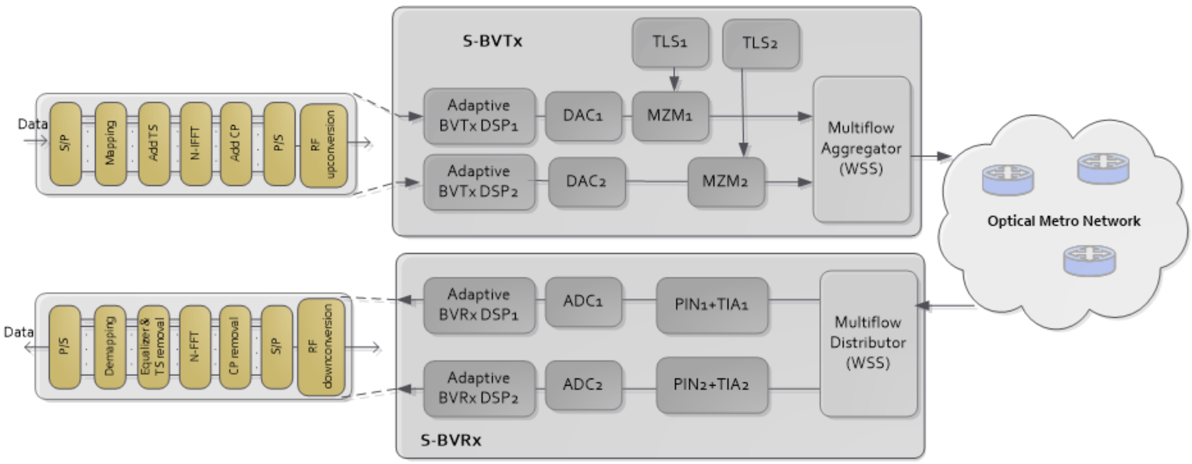

Figure 41. Programmable S-BVT architecture composed of two BVT modules based on MCM and DD.

Table 30 summarizes the S-BVT features and specifications. In particular, it can be seen that each slice can serve a maximum net capacity of 50 Gb/s, corresponding to an aggregated capacity of 100 Gb/s. 40Gb/s per slice can be supported over a 100-km path without any dispersion compensation modules. An extended reach of 200 km can be achieved, but at the expenses of the system data rate. The modular S-BVT architecture provides variable capacity, scalable with the number of enabled active slices (and additional modules following a pay-as-you-grow approach), with adaptive rate/reach performance according to the network needs, thus higher metro transport rates can be achieved with improved efficiency (related to the MH6 KPI).

Different programmable elements and reconfigurable parameters are identified in order to exploit the S-BVT potential to optimally adapt the transmission according to the traffic demand and network configuration (related to the MH1 and MH8 KPIs). A group of digital parameters, such as signal to noise ratio (SNR) and bit error rate (BER), have been identified for monitoring purposes. Thanks to this information, loading algorithms can be implemented to maximize system capacity/performance.

The proposed programmable multi-rate/format/reach/flow transceiver includes advanced functionalities/capabilities, such as sliceability, sub-/super-wavelength granularity, flexibility, reconfigurability and rate/distance and bandwidth adaptability.

Table 30. S-BVT features and specifications.

S-BVT feature	**S-BVT SPECS**
Number of enabled slices	Two (2)-> Extendable adding modules
FEC (Target BER)	HD-FEC (4.62e-3), SD-FEC (2e-3)
Max distances	200 km
Max capacity per slice	50 Gb/S (B2B)
Granularity	Subcarrier spacing (slice BW/subc)
Programmable elements	DSP, TLS, WSS, OA

Reconfigurable parameters	# of enabled active slices
	Bandwidth occupancy (BW)
	FEC type (HD/SD)
	Adaptive DSP mode
	Loading algorithm:
	  # of loaded subcarriers (Nsubc)
	  Modulation format (BL/UL)
	  RA (Gap)/MA (Bit rate)
	MZM bias (~null/quadrature)
	Aggregator transfer function per port
	Equalization type

Reconfigurable parameters within Openconfig model	Laser central wavelength (TLS), Laser output power
Monitoring parameters	<u>Digital:</u> CSI estimation (SNR), Pre-FEC BER
Potential capabilities	Sliceability, Programmability, Reconfigurability, Flexibility, Low cost
Specific sub-KPIs for the S-BVT	BER, Bit rate variable capacity, Sliceability functionality and Programmability

A set of sub-KPIs has also been identified for the modular S-BVT to define performance targets. In particular, a target BER, according to the specified hard-/soft-decision forward error correction (HD-/SD-FEC), has to be achieved. Moreover, bit rate variable capability enhances the system/network efficiency and can be achieved by varying the modulation format per subcarrier according to the channel profile. Hence, bit/power loading algorithms are applied analyzing/evaluating their impact on the system performance. Thanks to the sliceability functionality, multiple flows/slices can be aggregated and distributed towards different network nodes. This capability is demonstrated/assessed by enabling two slices of the S-BVT transceiver. Finally, the transceiver programmability can also be validated by implementing SDN agents to reconfigure/adapt different transmission parameters. Thanks to the adoption of the SDN paradigm, the S-BVT can be reconfigured to adapt the transmission to the network condition/requirements, in order to efficiently use the metro transport resources. Specifically, SDN agents are developed within the project to control, manage and configure the programmable S-BVT, considering the OpenConfig model.

4.2.4.2 CNIT/TEI DD transceiver

Within Metro-Haul, DD transceivers offering high dispersion tolerance at high bit-rates without requiring the installation (or the upgrade) of bulk dispersion compensating modules have been developed. This is achieved by exploiting a combination of alternative modulation formats with silicon photonics integrated modules. Tolerance to chromatic dispersion is enhanced through an appropriate combination of coding and pulse shaping, the “combined amplitude and phase shift” (CAPS) codes. While the receiver maintains the same hard threshold symbol-by-symbol receiver used for OOK, the generation of CAPS codes requires an encoder for the shaping of the signaling pulse and the use of a 2-channels DAC for the in-phase and quadrature components. A simplified implementation scheme that approximates the CAPS codes with simpler analog electronics is depicted in Figure 42(a). Such an architecture, named IQ-duobinary (IQduo) can reach a similar performance as CAPS and is obtained by associating a quadrature component to a duobinary coded signal, using an IQ-MZM. The tolerance to chromatic dispersion concerning simple OOK is strongly increased while maintaining the same receiver as used for OOK, at the expense of a marginally higher cost of the transmitter. At a bit-rate of 50 Gb/s, CAPS and IQduo provide dispersion tolerance for up to 20 km and 15 km of uncompensated G.652 SMF, respectively. To enable further reach extension for C-band high bit-rate transmission without requiring complex DSP, dispersion compensating devices based on integrated photonics are embedded in the optical transceiver. This optical dispersion compensator (ODC) is based on all-pass microring resonators and has been designed and implemented in sub-blocks. Each sub-block is able to compensate the chromatic dispersion of 10-km G.652 SMF, at a wavelength of 1550 nm, and is implemented as a cascade of three all-pass microring resonators that have been appropriately designed. By combining ODC and optical switches, the transponder can be configured for operation over a wide range of fiber lengths (Figure 42 b-c). Table 31 summarizes the main transceiver features and specifications, while Table 32 provides details on the achievable optical reach at 25 Gb/s and 50 Gb/s with different ODC configurations.

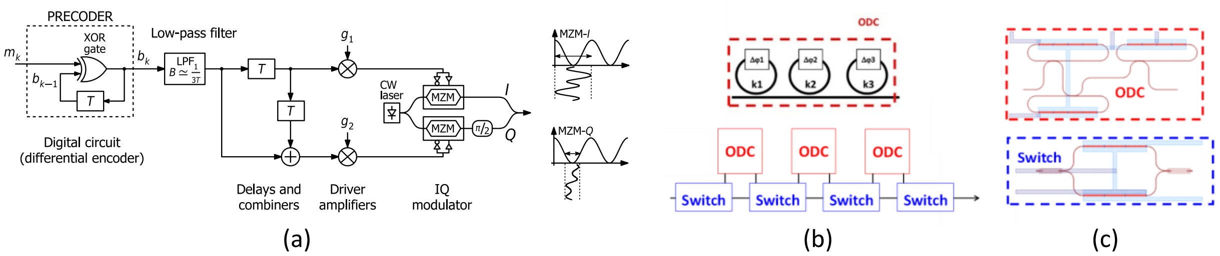

Figure 42 Schematic of the IQ-duobinary transmitter (a); schematic of the three microrings ODC and of the reconfigurable ODC module (b); mask layout of the ODC and the optical switch (c).

The transceiver encompasses several components that can be configured depending on the connection to be established as well as the link condition, e.g. line rate, laser wavelength, number of active ODC modules. In Metro-Haul, a NETCONF/OpenConfig agent module is implemented to enable the transceiver configuration in agreement with the YANG models of the device. A NETCONF server is implemented using the ConfD framework, acting as the northbound interface of the agent for communication with the SDN controller and an OAM (Operation, Administration and Maintenance) handler. Transmission parameters such as laser wavelength, modulation settings as well as the ODC configuration are internally configured via specific software modules that interface with the NETCONF agent via a NE handler implemented in Python. The NE handler maintains the current view of the transceiver and incorporates the different drivers for communication with the sub-blocks device, relying on proprietary APIs for the different components.

Table 31. Transceiver specifications

Tx feature	**Tx SPECS**
Max capacity per lambda	50 Gb/s up to 50 km
Optical reach (G.652 SMF)	50 Gb/s: 50 km 25 Gb/s: 100 km 12.5 Gb/s: >200 km
Transmitter Bandwidth (-3 dB)	18 GHz (50 Gb/s) 9 GHz (25 Gb/s)
Line Rate	26.75 – 53.5 Gb/s
Programmable elements	Central wavelength Laser output power Line rate N° of active ODC modules
Output power range	0 – 10 dBm
Tunability range	1528 nm – 1568 nm 191.300 THz – 196.300 THz
Reference FEC (Target BER)	7% HD-FEC (3.8e-3)
Reconfigurable parameters within OpenConfig model	Laser central wavelength Laser output power N° of active ODC modules
Monitoring parameters	Pre-FEC BER
Potential capabilities	Programmability, Reconfigurability, Low cost
Specific sub-KPIs	Optical reach, bit rate, low cost, reconfigurability, power consumption compliant to QSFP form factor.
Rx feature	**Rx SPECS**
Photoreceiver type	PIN / TIA
Input power range	-10 / +5 dBm
Receiver bandwidth	40 GHz
Minimum required OSNR in BtB	19 dB at 50Gb/s, 7% HD-FEC (BER=$3.8x10^{-3}$)

Table 32. Optical reach for different modulation formats at 25 Gb/s and 50 Gb/s with and without ODCs. Longer reaches can be enabled by different ODC configurations.

Speed	Modulation Format	w/o ODC	1 ODC	2 ODC	3 ODC
25 Gb/s	CAPS	80 km	90 km	100 km	110 km
	IQduo	70 km	80 km	90 km	100 km
50 Gb/s	CAPS	20 km	30 km	40 km	50 km
	IQduo	17 km	27 km	37 km	47 km

4.2.4.3 ADVA Wavelength-Blocker

The basic structure of a wavelength-blocker based ROADM node is shown in Figure 43. It is a 2-degree ROADM. Physically, to enable the two directions, three devices are employed: two "terminals" which include EDFAs and splitters, and one LCoS-based blocker in the middle, which enables the channel-wise adjustment of power levels and blocking of dropped wavelengths. The signal travelling from left to right through the node in the figure is first amplified by an erbium-doped fiber amplifier (EDFA) which also serves as a pre-amplifier for the signals to be dropped in the node. In principle, all signals are available at the output of the splitter following the EDFA. Those, which are dropped at the node and not broadcast to other nodes are blocked by the wavelength blocker in the middle. The internal structure of the wavelength blocker ROADM is shown on the right-hand side of the figure. The LCoS chip is divided into four sections, where two perform the blocking or power adjustment, one for each direction, and the other two sections act as sweeping filters. The output power of these two filters is detected by monitoring photo diodes. Thus a dynamic adjustment of the power levels in all wavelength channels is possible. After blocking the dropped channels, the subsequent splitters allow the addition of new channels onto the blocked wavelengths. The channels passing though the ROADM are then power-equalized. If the span loss is high before entering the fiber, the signal can be amplified by a booster EDFA included in the terminal. For short spans, this EDFA can also be omitted. Passive spitting without a pre-amplifier can also be realized, depending on the network parameters. The optimum splitting ratios are dependent on the number of channels in the whole network, the number of channels to be added and dropped, and the use of EDFAs. They can be optimized to fit most requirements.

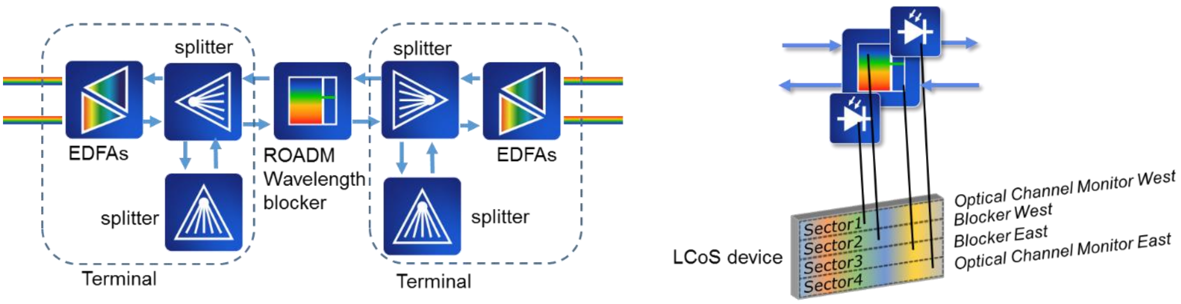

Figure 43: Left: ROADM node based on wavelength blocker, with terminals including amplifiers and splitters. Right: ROADM device with sectioned LCoS chip for blocking/power adjusting and monitoring functionality.

While in principle at the output of the splitters, AWG filters can also be employed to allow the separation of signals with direct detection, in the scope of the project, the use without any additional filters and only coherent detection is envisioned. The wavelength blocker device can be operated with 50-GHz and 100-GHz channel grids as well as with flexgrid channels in the C-band.

In a 40-channel system (e.g. 100-GHz channels grid in the C-band), the cost for a fully functional WSS-based ROADM is roughly twice the cost for a fixed OADM solution. If a wavelength blocker based solution is employed, the additional flexibility and functionality increases the cost by only approximately 30%.

4.3 Economic Impact of Optical Disaggregation Models

The objective of the model and the methodology presented in this section is the evaluation of the economic saving opportunity of equipment and SW disaggregation in the optical layer. According to the proposed approach, the evaluation is made comparing the legacy solution (aggregated) with semi-disaggregated and disaggregated options. The proposed methodology, partially inspired by [BNaud16], is characterized by the following action steps:

1. Assume the generic optical layer equipment models, the digital and analogue domains, and the disaggregated options for the optical layer as they are defined in Section 3.2 of [D3.1].
2. Map the Metro-Haul optical layer equipment solutions as specified in D4.2 to the generic equipment model of point 1.
3. Assign costs to equipment and systems according to the cost model (Section 3.4).
4. Design and take the Bill of Materials (BoM) of one or more metro networks (at least one metro aggregation and one metro core) for all the aggregated and disaggregated options.
5. Apply the cost model to the BoM of point 4 for the alternative aggregated/disaggregated options and make comparisons using a «what if analysis». The analysis is performed defining the potential coefficients of cost variation expected (forecast) when disaggregation is introduced.

In the following subsections the optical layer equipment model, the concept of the domain (analog and digital-to-analog) and the disaggregated options adopted in Metro-Haul and described in [D3.1] are first recalled. Afterwards, a techno-economic evaluation model is described end exemplified using a numerical example. The numerical example does not constitute an output of the evaluation, but is proposed just to show how the model can be applied. The model refinement, the parameters assignment (which is the critical point of the model) and the final results of evaluation require additional work and will be included in deliverable D2.4.

4.3.1 Disaggregation Options for the Optical Layer

The adoption of a disaggregated optical layer has been taken into account in Metro-Haul as one of most relevant elements of innovation to be introduced into metro networks, potentially able to produce significant advantages in terms of flexibility and cost reduction (above all thanks to the overcoming of the vendor lock-in), but also carrying possible issues that must be carefully investigated (mainly related to the unpredictability of the system integration effort).

The concept of "disaggregation" is generally used in a broader context with respect to the optical layer; in particular, regarding packet networks, it is mostly referred to the adoption of general purpose hardware ("commercial off the shelf" or COTS components), acquired separately from the software (commercial or open source), in a variety of different disaggregation levels (classified as "bare metal", "white boxes", "brite boxes"…). The term "boxes" refers to the adoption for the packet network equipment of a "pizza box" form factor, typical of DCs server and storage units, abandoning the traditional telecom assembly of vertical blades to be mounted inside standard ETSI racks, and thus also implying that the evaluation of energy consumption and footprint, together with the cost savings obtainable by the separation of hardware from software, should be taken into account in the assessment of the economic impact of disaggregation.

In the optical layer context, the term "disaggregation" is related to the same form factor change that concerns the packet layer, but implying that different optical functionalities, traditionally performed by different blades integrated in a single item of equipment and interconnected by a back panel, are now performed by different boxes, interconnected by external cables. The main challenge of disaggregation, and potentially also most of the benefits regard the deployment of an optical network using boxes provided by different vendors, where the heaviest implications relate to the requirements in terms of horizontal and vertical interoperability, respectively at the data plane level among different boxes, and at the control plane level towards a common controller and management tool. In this regard, different levels of disaggregation can be introduced, and here the most significant options are briefly analyzed, starting from the classification already made in WP3 (section 3.3 of [D3.1]).

As reported in [D3.1], the optical layer is composed of two domains: the “Digital-to-WDM adaption layer” (DtoWDM), in charge of the adaption of digital client signals to analogue “media channels”, and the “WDM Analogue transport layer” (A-WDM), in charge of “media channels” add-drop, switching, multiplexing, amplification, equalization and transport; the boxes, or optical network elements (O-NEs) that compose a Metro-Haul disaggregated optical network belong to either one of these domains, and can be classified as follows.

O-NEs belonging to the DtoWDM domain:

- Transponders: 1-1 mapping of client interfaces
- Muxponders: N-1 mapping and multiplexing
- Switchponders: N-mapping, switching and multiplexing

O-NEs belonging to the A-WDM domain:

- MD-ROADMs: Multi-Degree Reconfigurable Optical Add-Drop Multiplexer; this O-NE is usually assembled by several boxes, each containing a single degree or add-drop module
- Line Terminals: single line side optical multiplexer
- F-OADMs: Fixed-filtered Optical Add-Drop Multiplexer
- Fls-OADMs: Fixed-filterless Optical Add-Drop Multiplexer
- ILAs: In-Line Amplifiers

Starting from the above classification, 3 main scenarios, concerning disaggregation can be described and analyzed.

4.3.1.1 Fully Aggregated Optical Network

An optical network can be defined as fully aggregated when all the O-NEs, and the controller are provided by the same vendor (Figure 44). As shown in the figure, an open and possibly standard NBI can be provided towards a higher level controller or orchestrator, that would also control for example the packet layer, enabling end-to-end management and service provisioning.

This solution has the advantage for the telecom operator of not being in charge of the system integration activity (neither to be performed internally nor to be outsourced to a third party), as the vendor should solve all the interoperability issues among the network elements and between the O-NEs and the proprietary controller of the optical layer. On the other hand, the cost of the hardware is expected to be penalized by the vendor lock-in.

The fully aggregated option should be considered mainly as a benchmark, to evaluate possible consequences of disaggregation on the overall costs of an optical network. For this purpose, the development of a cost model for this solution should be based on a truly traditional solution, where functionalities belonging to the DtoWDM and the A-WDM domains (with the only exception of ILAs) are integrated into a single case.

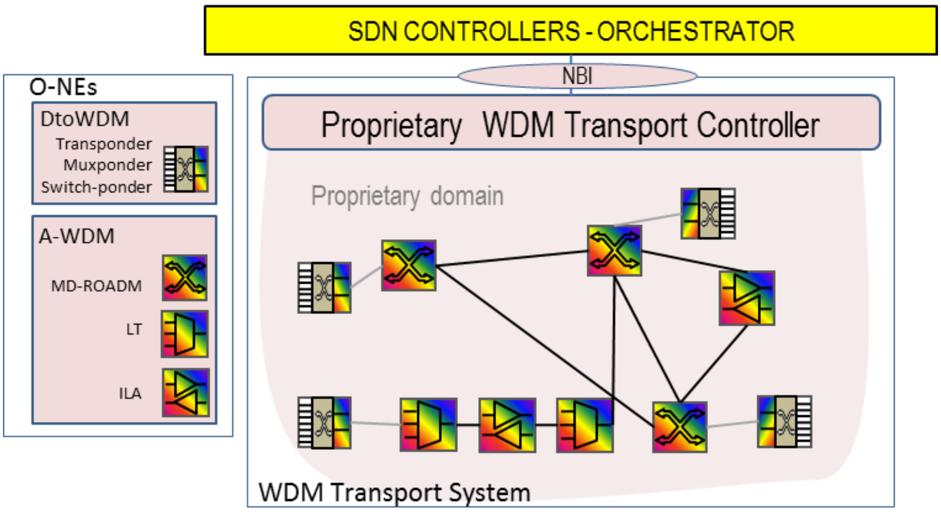

Figure 44. A fully aggregated optical network (figure 13 of D3.1)

4.3.1.2 Partially Disaggregated Optical Network

This solution is represented by the case where all the boxes belonging to the A-WDM domain (that can be together all defined as an "Open Line System", or OLS) and the related controller are provided by the same vendor, and are open to the transport of media channels provided by different vendors O-NEs belonging to the DtoWDM domain. The interoperability between the two domains must be guaranteed by the definition of a "standard" interface at the border, here named the Single Wavelength Interface (SWI). Most of the advantages of the aggregated solution are maintained: in particular the control loops necessary to tune the analog parameters and optimize the optical transmission can still be performed in a proprietary way by the OLS vendor; moreover, if transponders provided by the same vendor are used as a pair at both ends of each connection, many features related to the DtoWDM domain (FEC, DSP at transmitter/receiver, etc.) also don't need to be standard, and can lead to a higher performance. The interoperability issues at the data plane level are limited to the border between the A-WDM and DtoWDM domains, while at the control plane level the integration of the two domains is regarded as being under a single WDM transport controller (Figure 45 (a) and (b) show two possible implementations of this integration).

The effort demanded of the telecom operator related to system integration should then be quite limited for this solution, which represents a good compromise in the short-middle term, when the standardization activities of MSAs like OpenROADM and OpenConfig, regarding both data plane interoperability and open yang models for the control plane, are still being developing, and their implementation by equipment vendors is still at a preliminary stage.

Furthermore, the greater benefit of this solution depends on the fact that the lifecycle of DtoWDM boxes is generally much shorter concerning the OLS, as technology innovation is evolving much faster for the transponders and muxponders then for the optical switches. So the possibility to preserve the infrastructural investment made in the A-WDM domain, while following the technology evolution in the DtoWDM domain, should lead to significant savings for telecom operators during the lifetime of their optical networks.

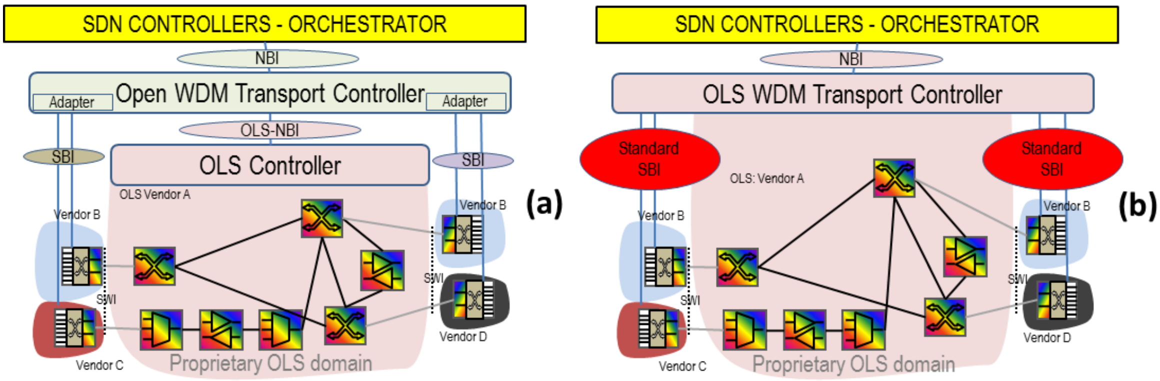

Figure 45. A partially aggregated optical network with (a) proprietary SBI, or (b) standard SBI; (figure 14 of D3.1)

4.3.1.3 Fully Disaggregated Optical Network

The full disaggregation option represents the case where the A-WDM domain can also be composed of boxes provided by different vendors, and is by far much more challenging, in terms of system integration, then the previous scenario (Figure 46).

To ensure data plane interoperability in this scenario, the interface between different boxes belonging to the A-WDM domain, here named Multi Wavelength Interface (MWI) following the OpenROADM terminology, must also be compliant to a standard definition. Moreover, the WDM controller, in charge of managing all the issues related to analog optical transmission in a multi-vendor environment, must be necessarily vendor-agnostic, and the SBI between different vendors O-NEs and the controller must therefore also be compliant to a standard.

For these reasons, the adoption of this model is not likely to become feasible in the near future, and its possible subsequent success will strongly depend on the availability and maturity of standards capable of ensuring horizontal and vertical interoperability.

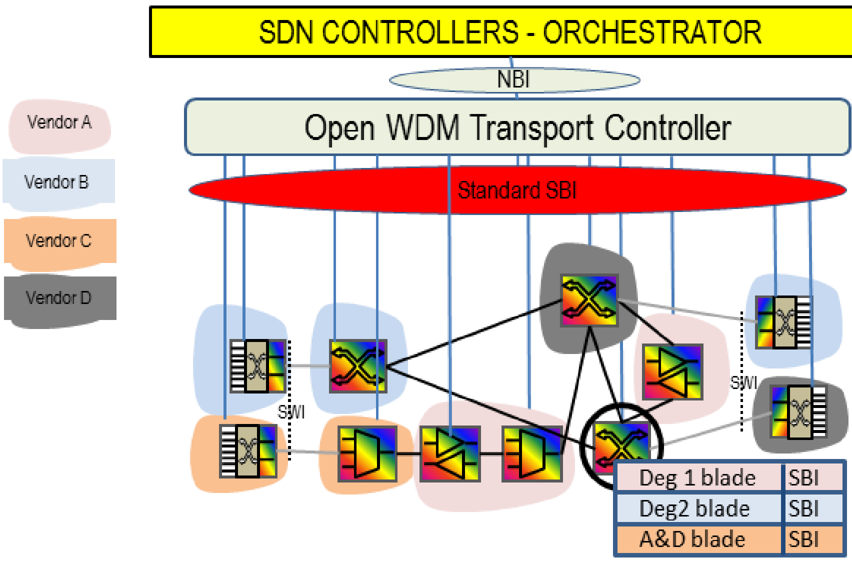

Figure 46. A fully disaggregated optical domain; figure 15 of D3.1)

4.3.2 Model for evaluation of the economic impact of disaggregation in the optical layer

The general assumptions of the techno-economic evaluation model of disaggregated options are listed below.

- All options can deliver the same services in terms of type, the fulfillment of QoS and, in case of a whole network designed and operating, the same volume of carried traffic.
- In the "what if analysis" phase of the approach presented in this section, applied for making comparisons between alternatives, a variation in cost of the disaggregated components (typically a reduction for HW and SW, possibly an increase for integration) is considered when certain conditions are assumed to be foreseeable (e.g., strong competition, high volumes, use of free SW, etc.).
- The condition of a mature market is assumed (all parts for all options, HW and SW, have reached a stable release, i.e., no troubleshooting effort (over and above that required in standard conditions) during systems integration is assumed for all partial and fully disaggregated options). Also, all standards for interoperability are assumed to be available and implemented in commercial products.
- Macro cost components for A-WDM and DtoWDM (both HW and SW) are kept separate in the evaluation and analysis. This is motivated by the following two main reasons:
    1. In the partially disaggregated solution the OLS is only made of A-WDM equipment;
    2. Analogue-optical and digital (optoelectronic) equipment can have different lifecycle and cost reduction margins. This is because optoelectronic equipment has a faster renewal cycle and relies heavily on a large-scale semiconductor industry, which is highly dynamic and subject to more competition than analogue-optical equipment.
- For all solutions the integration cost is broken into three main subparts:
    1. Integration of A-WDM equipment
    2. Integration of DtoWDM equipment
    3. Integration of A-WDM with DtoWDM

As far as the systems modelling is concerned, the components relevant for the cost evaluation are the following:

- HW
    1. A-WDM equipment
    2. DtoWDM equipment
- SW
    1. A-WDM equipment (control SW/agent)
    2. DtoWDM equipment (control SW/agent)
    3. vendor specific WDM transport controller (all WDM equipment)
    4. open WDM transport controller (all WDM equipment, for partially disaggregated option (a) and fully disaggregated option)
    5. OLS WDM controller (for A-WDM equipment only, for partially disaggregated option (a))
    6. OLS WDM controller capable of also controlling the DtoWDM equipment of other vendors (for partially disaggregated option (b))
- Integration costs (includes HW and SW testing, debugging, troubleshooting etc. during the network creation phase.)

1. Integration of A-WDM equipment
2. Integration of DtoWDM equipment
3. Integration of A-WDM with DtoWDM equipment

The system components listed above are generic and are required to be mapped with specific HW and SW elements when the techno-economic evaluation is performed. For instance, splitters/combiners, DD transceivers and their control SW in the evaluation of a specific horseshoe network in the metro aggregation. Some insights are therefore necessary here to clarify the model assumptions. The SW for equipment (both A-WDM and DtoWDM) is determined in different ways depending on the disaggregated option considered.

In the case of the aggregated mono-vendor option, equipment SW is developed within the integrated solution and can benefit from an economy of scale if the provider sells high quantities of equipment to a variety of customers. In this case, the software development and testing costs are very high because they involve an *ad hoc* development, but such costs can be shared between a large number of such items of equipment, and this can lead to a very low cost per unit (i.e., each SW installation on a specific piece of equipment). This is true also for the A-WDM equipment SW of the OLS part of the network, in the case of partially disaggregated options.

In the case of partially or fully disaggregated options (in the first case only for DtoWDM equipment, and in the second case for both A-WDM and DtoWDM equipment) the cost of equipment SW is subject to a different model as compared to the case of the aggregate option. In these cases, the SW installed on equipment can rely on the availability of free SW (for instance on free SW developed within the framework of a standardization body like the Open ROADM Multi-Source Agreement); but, in general, it necessitates additional development for the customization required by the specific network implementation. The cost of SW customization (indeed very difficult to be estimated itself) in this case is shared between a limited number of items of equipment (all the pieces of the same type and bought by the same whitebox HW provider in the network) and this can lead to a high cost per unit, even if the overall cost of development is not as high as in the aggregated option.

Another specification is required relating to the meaning of Integration cost components. In the version of the model presented here, the integration costs don't include the development of the SW to be installed on equipment or for the centralized control plane. These costs are included in the equipment SW and controller cost components.

The components of total integration cost include all the effort spent in integrating the available pieces of the network and assuring that they interwork properly. It includes troubleshooting the network deployment and, in particular, the detection of non-compliance with the specifications of HW and SW of equipment and in the controller, and managing the requests for updates or repairs made to suppliers. In the case of partially and fully disaggregated options, integration is made under the assumption that all equipment items, HW SW and controller, respond to the well-defined specifications, hopefully, issued by standardization bodies.

Another point concerns the meaning of the three components of the integration cost. A-WDM equipment integration cost involves all the costs in integrating the HW and SW of the equipment and assuring compliance with the specifications they must satisfy. This aspect is the responsibility of the HW equipment maker, the equipment SW developer, or of a third party. This integration could be made independently from the specific network to be deployed, and involve only the HW, SW, and the specifications required for the specific types of equipment under consideration. The same holds for the DtoWDM equipment integration cost. Once the integration is performed under given specifications for a given equipment type, no trouble should emerge when such equipment is introduced into the network.

The cost of integration of A-WDM with DtoWDM equipment involves the specific network deployment, and includes all the effort spent in solving the problem of effective interworking of all network pieces. This is the most critical point, especially for disaggregated options which integrate a multitude of different HW and SW items, each of them potentially from different providers.

The methodology proposed for the evaluation is based on the setting of a baseline cost equivalence framework and a subsequent "what if analysis" introduced to assess the quantitative impact on cost variations of disaggregated options in the presence of saving possibilities in some or all of the disaggregated components of the model.

In the baseline cost equivalence framework, the reference case for all disaggregated options is an ideal neutral condition characterized by the statement that making a thing (a piece of HW, a SW application, an integration process) costs the same, independent of the disaggregated option and of the actor involved in the making of that thing. For the reference case, the cost of each piece is simply its standard industrial cost. This basic statement is motivated by the fact that there isn't any *a priori* reason to assume that a big vendor can sustain lower costs than a small supplier or vice versa: it could depend on what item one considers, on the volume of purchased items, and on other conditions that are not predictable without a specific deeper context and trend analysis.

Given a semi-disaggregated or disaggregated option, a cost variation coefficient can be applied to some cost items to take into consideration the fact that such items can take advantage of conditions that are typical for that option.

Examples of dependencies between reductions of some cost components and specific conditions of which disaggregated solutions can benefit are the following:

- Reduction of HW cost of DtoWDM in partially disaggregated scenario due to overcoming vendor lock-in and competition between suppliers;
- Reduction of HW cost of both DtoWDM and A-WDM systems in a fully disaggregated scenario for the same reasons as adduced above;
- Reduction of cost of SW on equipment in the case of disaggregated options (for both A-DWDM or DtoWDM) because the SW can rely on the availability of free SW and the fact that only limited adaptation may be necessary. This point is controversial because the mono-vendor aggregate solution can rely on an economy of scale in developing *ad hoc* equipment SW, as was already discussed above.
- Reduction of the cost of the open WDM transport controller in a fully disaggregated scenario as compared to the cost of a mono-vendor controller, since an open controller can rely on free SW and only require some limited customization/adaptation.

It is important to consider that disaggregated solutions can also lead to an increase in the cost of certain components, typically the cost of integration, which is likely to be more expensive than for the mono-vendor case. This is because the systems to be integrated are very heterogeneous, which requires additional effort in the integration phase.

To illustrate the results that could give an example of the proposed evaluation approach, we see below some breakdowns that are not yet related to real data, but provide an illustrative example: Figure 47, which reports an example of what the cost breakdown could be (here limited to CapEx only) according to three different assumptions as regards the prices of network components that aggregated and disaggregated solutions can assume (diagram (i), (ii) and (iii)). The four network solutions in each diagram are the ones presented above: fully aggregated, partially disaggregated type (a), partially disaggregated type (b) and fully disaggregated.

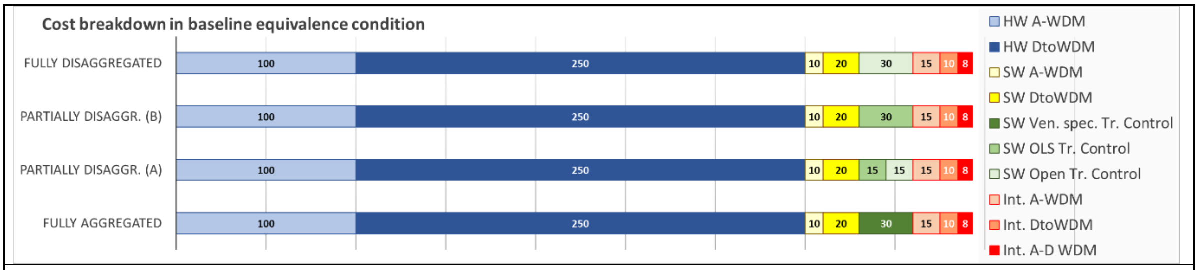

(a) Baseline equivalence

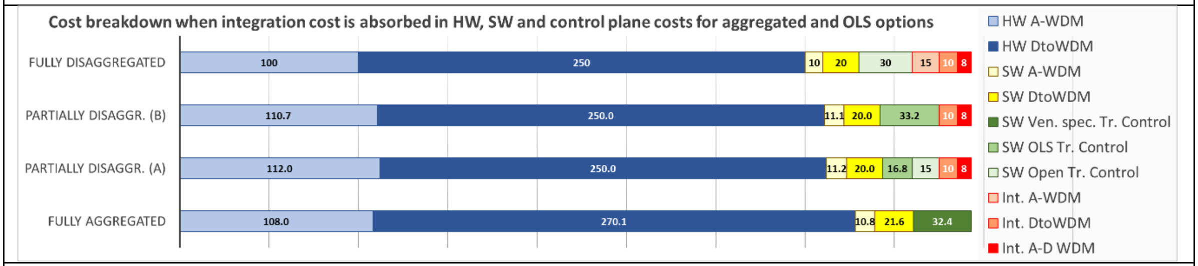

(b) Integration cost absorbed

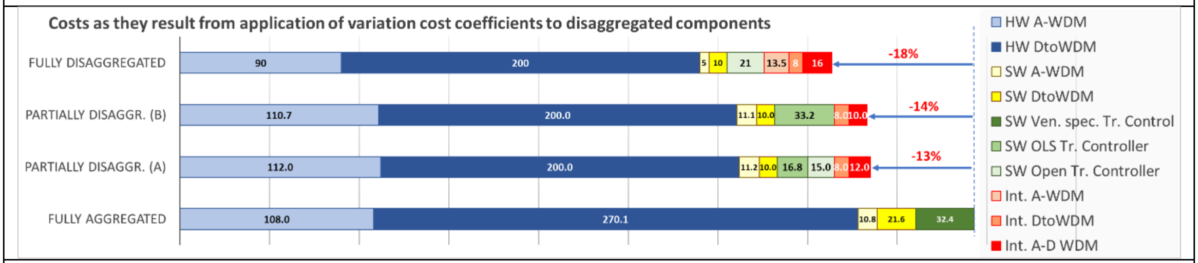

(c) Application of cost variationcoefficients to disaggregated components

Figure 47. Cost breakdown of aggregated and disaggregated options under three different conditions.

The cost breakdown consists of the costs of the following parts:

- Equipment HW, in blue, differentiated for A-WDM (dark blue) and DtoWDM (light blue) equipment;
- Equipment on board SW (i.e., OS and control SW, in yellow, differentiated for A-WDM in dark yellow and for DtoWDM in light yellow;
- Control plane, in green, with three different subtypes: vendor specific control plane for the whole network (dark green), vendor-specific for the OLS part only (intermediate green) and customized open free SW (light green); Please note that for the partially disaggregated type (a), two control plane SW systems are present: one for controlling the OLS and one for coordinating the OLS with the DtoWDM Equipment (transceivers).

Integration, in red, with the specification of its subparts using light red for the integration cost of the analog part only, of the intermediate red integration cost for the digital to analog part, and of dark red for the integration cost of analog with digital to analog parts.Figure 47 (i) the cost breakdown is shown assuming the baseline equivalence condition. Costs are the same for the same function among the different options. Integration costs are made explicit for all options, even though in the aggregated option they are absorbed in the other costs, and in the options involving the OLS, the integration cost of A-WDM part is absorbed in the other HW and SW A-WDM cost components.Figure 47 (ii) a more realistic view of the cost breakdown is displayed for the same network of the case (i). The fully aggregated option (legacy) integration costs are considered as being absorbed in the HW and SW cost and assumed to be proportionally spread to HW, SW and control plane cost. For the OLS options (partially disaggregated (a) and (b)) the A-WDM integration cost is absorbed by and spread amongst the A-WDM cost components, the other integration costs remain explicit. For the partially disaggregated (a) the cost of the control plane components (the OLS and the open ones) is assumed to be 50% each of the total cost of the control.Figure 47 (iii). Some cost components for the partially disaggregated and fully disaggregated options are reduced using coefficients on the base of favorable trends. Due to the criticality of integration in disaggregated options, the integration cost of the pure optical components with the transceivers (Int. A-D WDM) is assumed higher for both partially and fully disaggregated options. Figure 47 (iii) are the ones reported in Figure 47where negative percentages (in green) are the expected cost reductions benefiting from disaggregation, while the positive percentages (in red) are the presumable cost increases due to the higher effort to be sustained in the integration process in the case of disaggregated solutions. Figure 47 (iii), with the cost variation coefficients ofFigure 47, the overall cost of the disaggregated solutions are 13% and 14% lower, respectively, for the partially disaggregated solutions (a) and (b). The overall cost is 18% lower for the fully disaggregated option, when compared to the fully integrated case.

Table 33: Coefficients of cost reduction applied in the cost breakdown diagram of Figure 47(c).

	Cost variation coefficients			
	Fully aggregated	Partially disaggregated (a)	Partially disaggregated (b)	Fully disaggregated
HW A-WDM	0%	0%	0%	-10%
HW DtoWDM	0%	-20%	-20%	-20%
SW A-WDM	0%	0%	0%	-50%
SW DtoWDM	0%	-50%	-50%	-50%
SW Ven. spec. Tr. Controler	0%		-	-
SW OLS Tr. Controller	-	0%	0%	-
SW Open Tr. Controller	-	0%	-	-30%
Int. A-WDM	0%	0%	0%	-10%
Int. DtoWDM	0%	-20%	-20%	-20%
Int. A-D WDM	0%	50%	25%	100%

Continued work on the proposed evaluation model above consists of reviewing and consolidating the model, and applying it to network scenarios for case study comparisons between aggregated and disaggregated options.

An important and critical point is the assignment of the cost variation coefficients of Figure 47. A specific activity will be required for the estimation of the cost variation coefficients, according to the factors that can influence each cost component when disaggregation is introduced. In particular, a critical point to be addressed is the method to estimate the difference in SW cost, which has a different model for the aggregated and disaggregated options, and also to evaluate the factors that influence integration cost components between aggregated/disaggregated options.

The model presented involves the CAPEX only. If one were to also consider the impact on OPEX, an investigation into which factors differentiate the OPEX among the disaggregated options would also be necessary. Within this scope, the work in the following months, to be reported in D2.4, will include refinements to the model and proposed estimates for the identified cost variation intervals.

4.4 Techno-Economic Impact of Edge Computing

This Section reports on various network optimization/dimensioning studies done within the scope of Metro-Haul, that will be harmonized through the input data modelled in Section 3 (topologies, traffic, cost modelling).

4.4.1 Data-Center Dimensioning and Network Cost Analysis based on Latency Classes

The LA-ML-SCA algorithm receives as an input the following parameters: (i) an IP-over-WDM network, (ii) with available IT resources in the COs, and (iii) the service chain requests together with their latency requirements. Its goal is to satisfy the requests by allocating service chains, and creating the needed VNF instances and IP links that are realized via lightpaths over the WDM plant. Latency-aware means that allocations are restricted to satisfy the latency requirements, as well as to avoid IT resource oversubscription and spectrum clashing. LA-ML-SCA first sorts the service chain request according to its injected traffic (higher traffic first), and then sequentially processes each request one by one.

The main core of the algorithm (lines 3-28) first calculates the shortest path $p(r)$ from the origin node to the nearest core node (line 4). It tries the allocation of all the VNFs in the node furthest from the user (closest to the core, and thus more centralized). This potential allocation is evaluated in three approaches: (a) Firstly, lines 7-9 evaluate if the latency requirements are satisfied with $p(r)$ considering OEO and propagation. If so, the algorithm establishes a lightpath per link in $p(r)$, which creates a new logical IP link or increases IP capacity (in case an IP link is present); (b) secondly, lines 11-18 aim to reuse previously established lightpaths (λ_{inter}) evaluating potential lightpaths (λ_{ori}*) in terms of total latency with two OEO conversions; (c) Thirdly, lines 19-23 evaluate if a direct lightpath satisfies the latency requirements.

Finally, LA-ML-SCA iterates the three approaches above, reducing by one hop the shortest path in each iteration (lines 24-25). In case any approach above is satisfied, all VNFs are instantiated in the current destination node of $p(r)$. The output of LA-ML-SCA is used to compute the CO and total network cost.

Algorithm 1 Latency-Aware Multilayer SC Allocation

```
Input: G = (N,E), L_max
Output: VNF placement, IP links and routing, lightpaths
Begin:
1.  Sort service chain requests (SCR) in decreasing traffic
2.  Create empty list of established lighpaths: LP
3.  foreach (r ∈ SCR) do:
4.    Calculate shortest path p(r) in km from oriNode(r) to the
      nearest core node (destNode(r))
5.    while (l_dest(p(r)) > l(r))
6.      Calculate l_dest(p(r))     // latency to destNode(r) with OEO
        conversion penalty
7.      if (l_dest(p(r)) <= l(r))
8.        Establish lightpath: λ ∀ e ∈ p(r) | λ(e) ∈ LP
9.        Instantiate V(r) in destNode(r)
10.     else
11.       foreach λ_inter ∈ LP(ending at destNode(r)) do:
12.         Calculate l_2-hop(p(r)) = l(λ_ori*)+ l(λ_inter)
13.           if (l_2-hop(p(r)) <= l(r))
14.             Establish lightpath to interNode(r): λ_ori
15.             Instantiate V(r) in destNode(r)
16.             break
17.           end if
18.       end foreach
19.       if (l(λ_direct*(r)) <= l(r))
20.         Establish lightpath to destNode(r): λ_direct
21.         Instantiate V(r) in destNode(r)
22.         break
23.       end if
24.       Calculate l_pNode (r)     // Calculate latency to the previous
          node pNode(r) with OEO conversion in p(r) = p(r)-1
25.       Update destNode(r) = pNode(r)
26.     end if
27.   end while
28. end foreach
```

In a preliminary analysis done using realistic data from the Murcia-Alicante region, we considered a network infrastructure with a sufficient pool of IT resources (CPU, RAM, HP) in all COs to place all SCRs without contention. The network size is obtained by executing the proposed LA-ML-SCA algorithm, which places all VNFs over the M-A network. Note that LA-ML-SCA progressively establishes IP links (and lightpaths) to meet the services latency requirements, starting with an empty network.

We investigate the network design cost (in arbitrary units *a.u.*) as a function of the services latency requirements. In particular, we explore a range of L_{max}, latency values (maximum latency from the user to the first VNF) in a range between 1 ms up to 60 ms, which correspond to the most stringent and the most permissive scenarios, respectively. We recall that L_{max} is used to define the maximum latency constraints for each service chain request, which are uniformly picked in the range [0.5, L_{max}].

Figure 48 shows the total network cost as a function of the considered values of L_{max}. Results show that the total network cost that guarantees the services' latency below 1 ms is as high as up to four times the value of the more permissive scenarios $L_{max} \geq 30$ ms. An abrupt cost increase is observed for values below 10 ms. Indeed, low latency values are only attainable by instantiating VNFs close to the origin nodes of the service chains, which decentralizes the IT resources into a larger amount of COs. Conversely, the more permissive scenarios (i.e. $L_{max} \geq 30$ ms) tend to concentrate the IT resources into MCENb, which leverages the economy of scale and reduces the network cost.

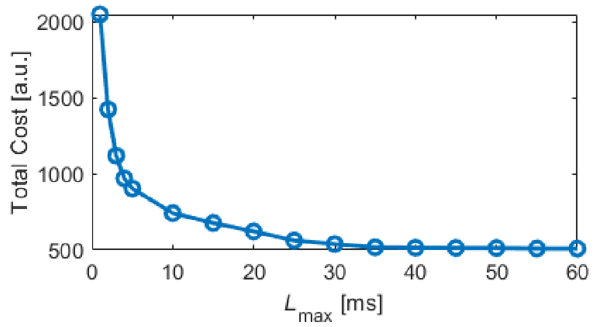

Figure 48: Total network cost as a function of the maximum latency allowed to reach the first VNF.

Figure 49 illustrates the per node cost distribution, when considering L_{max} values of 1 ms, 15 ms and 60 ms. Nodes in the x-axis are sorted in decreasing order according to their cost for L_{max} = 1 ms, denoting MCEN and MCENb as (M) and (Mb), respectively. The MCENb central nodes in the capital cities (Murcia1 and Alicante1) concentrate on the largest amount of resources. For L_{max} = 1 ms, Murcia1 and Alicante1 present a 10-fold value in cost as compared to the remaining nodes. Additionally, the resources not allocated in the MCENb nodes are distributed among various AMEN and a few MCEN nodes. L_{max} values of 15 ms and 60 ms further concentrate the network cost in the two MCENb nodes, up to the situation where no resources are needed in both MCEN and AMEN nodes in the latter case. Relevant to this observation, Figure 53 groups the cost distribution per node type for the three L_{max} values under study. Indeed, as anticipated in Figs. 4 and 5, network cost is concentrated in MCENb nodes for high (i.e. more permissive) L_{max} values whereas low latency values spread the network cost toward nodes not connected to the backbone network.

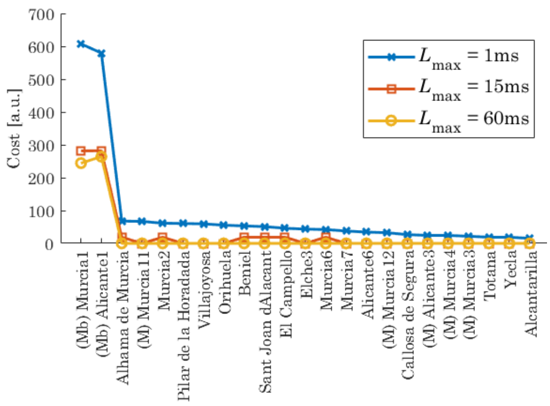

Figure 49: DC costs for maximum latencies allowed of 1, 15 and 60 ms.

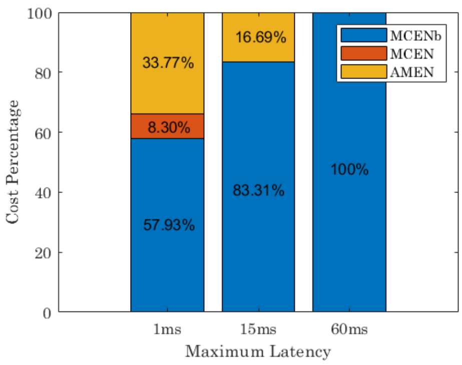

Figure 50: DC cost distribution between MCENb, MCEN and AMEN.

4.4.2 Dynamic Service Chaining

In this section we first describe an algorithm that performs the VNF placement and traffic routing at each time instant, considering the current status of the network. Then we discuss the numerical results obtained by comparing our algorithm with two benchmark algorithms.

4.4.2.1 Heuristic Approach for Dynamic VNF placement (DVNFP)

To summarize the problem, we state: (i) A hierarchical optical aggregation network composed of COs connected to each other using WDM links, is given; (ii) Service Chain (SC) requests are dynamically generated by users in cell sites and terminated in a CO based on the type of the SC; (iii) For each SC request the placement of its VNFs on nodes equipped with processing units (NFV-nodes) need to be decided. Our objective is to minimize the number of active NFV-nodes as well as the blocking probability, taking into account the constraints on maximum link capacity, maximum computational capacity of NFV-nodes and maximum tolerable latency for each SC.

As the first step, our algorithm, builds an auxiliary graph of the network composed of all nodes and WDM links. It takes as input an SC request specified by the properties such as source, destination, number and type of VNFs composing the SC, the number of users requesting it, the latency requirement of the SC, and its holding time. The main steps of the algorithm can be defined as follows:

- *Reusing active VNFs:* Since activating a new instance of a VNF on an NFV-node will impose an additional cost, DVNFP tries to use already-activated NFV-nodes hosting requested VNF instances, as much as possible. Therefore, for each VNF instance belonging to the SC request, DVNFP at first checks whether, in the network, there is an already-activated VNF instance of the same type or not.
- *Selecting among active VNFs:* When more than one active NFV-node hosting the required VNF instance, with enough capacity, exists in the network, DVNFP uses the locality-awareness metric. Here, for each NFV-node, the metric is defined as the length of the shortest path between the source of the SC request and the NFV-node, summed up with the length of the shortest path between the NFV-node and the destination of the SC request. It is worth mentioning that the shortest path is calculated in a way that congested links are not considered. Hence, DVNFP chooses the NFV-nodes with locality awareness value less than a predefined threshold, and they are chosen based on the topological aspect of the network. Among these NFV-nodes, depending on the SC request, a node is chosen; that is, if the SC requested needs large computational resources then the NFV-nodes closer to the Core CO (which is more likely to have the large computational capacity) will be chosen. However, for SCs with stringent latency requirements, DVNFP chooses the NFV-nodes among the COs at a lower level of the metro hierarchy. After choosing the best NFV-node, the VNF is placed at that node by allocating the required resources.
- *Activating new VNF instance:* If there is no already-active VNF instance of the required VNF on the network, DVNFP tries to instantiate a new one. Therefore, at first, it calculates the shortest path between the source and destination of the SC request. After that, the closest NFV-node to the source, along the shortest path, with enough computational capacity is chosen. If the VNF cannot be placed on any of NFV-nodes along the shortest path, the algorithm checks the capacity of all other NFV-nodes on the network and chooses

the one with the higher *betweenness centrality* (defined as the number of shortest paths passing through a node) and better locality-awareness.

Note that, at each step, the source of the SC request is replaced by the NFV-node chosen to host the VNF requested, at the previous step and the above-mentioned procedures will be repeated until all the VNFs of the SC are placed.

After all the VNFs are placed, DVNFP checks if the latency requirement of the SC is satisfied. If it is the case, the SC is provisioned, and when its holding time expires the resources occupied by this SC are released. However, if the latency requirement of SC is not met, the algorithm calculates the latency of all virtual links and finds the one that has highest latency value. Then, the computational resources used on the end points of this virtual links are released and their VNFs are placed at their adjacent virtual nodes (if they have enough capacity). After that, the shortest path between these new endpoints is replaced with the virtual link that has the highest latency. This procedure is repeated until either the latency requirement of the SC is satisfied, or all VNFs of the SC are consolidated onto one NFV-node.

4.4.2.2 Benchmark Algorithms

To evaluate the performance of DVNFP, we compared its performance with two benchmark algorithms:

- *Centralized service chaining:* In this approach, the main objective is to reduce the expenses (CapEx and OpEx) by using just one NFV-node with the highest computational capacity and serving all the SCs using that node.
- *Distributed service chaining:* In this approach, VNF instances are enabled on all the NFV-nodes whenever they are needed. Therefore, even if an already activated instance of a VNF exist in the network, the algorithm activates a new instance on NFV-nodes along the shortest path between the source and destination of the SC request.

4.4.2.3 Numerical Results

In this section, we compare the performance of our algorithm with benchmark algorithms. The performance metrics used for comparison are as follows; the blocking probability, that is calculated as the fraction of SC request served out of total number of SC requests; the average number of active NFV-nodes, which is calculated by considering the number of NFV-nodes that have at least one running VNF instance at each time instant; and the latency violation ratio, which represents the number of SC requests that violated the latency requirement out of the full total of SC requests.

The topology considered consists of 80 nodes, 15 of which are NFV-nodes, while the remaining nodes are forwarding nodes. In this topology, there are 170 WDM links, each supporting 16 wavelengths each with 40 Gbit/s capacity. We performed our simulation experiments using a discrete-event driven simulator that generates SC requests as input traffic according to a Poisson distribution of inter-arrival rates and a negative-exponential distribution for the holding times. All the results are obtained within the 5% confidence interval, guaranteeing 95% statistical confidence. The SCs we considered are depicted in Table 34 while the amount of CPU required for each VNF (percentage of CPU per user) is shown in Table 35. The VNFs are Network Address Translation (NAT), Firewall (FW), Traffic Monitor (TM), WAN Optimizer (WO), Video Optimization Controller (VOC), and Intrusion Detection and Prevention System (IDPS).

Table 34: Service chains with corresponding VNFs, bandwidth and latency requirements

Service Chain	Service Chain VNFs	Bandwidth	Latency
Cloud Gaming	NAT-FW-VOC-WO-IDPS	4 Mbps	80 ms
Augmented Reality	NAT-FW-TM-VOC-IDPS	100 Mbps	1 ms
VoIP	NAT-FW-TM-FW-NAT	64 Kbps	100 ms
Video Streaming	NAT-FW-TM-VOC-IDPS	4 Mbps	100 ms
MIoT	NAT-FW-IDPS	100 Mbps	5 ms
Smart Factory	NAT-FW	100 Mbps	1 ms

Table 35: CPU usage for VNFs

VNF Name	NAT	FW	VOC	TM	WO	IDPS
CPU Core Usage per User	0.00092	0.0009	0.0054	0.0133	0.0054	0.0107

As is shown in Figure 51, the blocking probability has an upward trend for increasing traffic load values. However, the blocking probability for DVNFP always lies between the blocking probabilities for the two benchmark algorithms, i.e. centralized and distributed. In addition, for most cases and especially for higher loads, DVNFP guarantees a blocking probability close to the lower bound of the blocking probability (the one returned by a completely distributed service chaining algorithm).

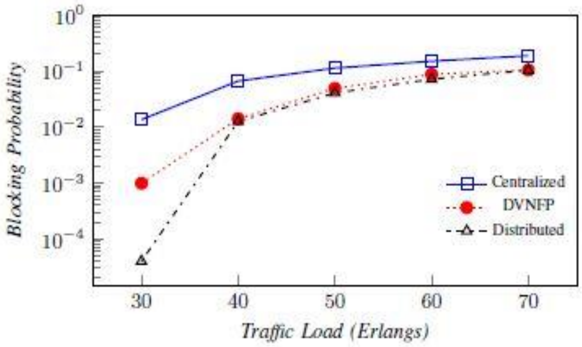

Figure 51: Blocking Probability

Figure 52 depicts the average number of active NFV-nodes. In this case, DVNFP uses up to 50% fewer NFV-nodes in comparison with the distributed service chaining approach. This will help telecom operators to almost halve the SC provisioning costs, since activating NFV-nodes imposes additional costs.

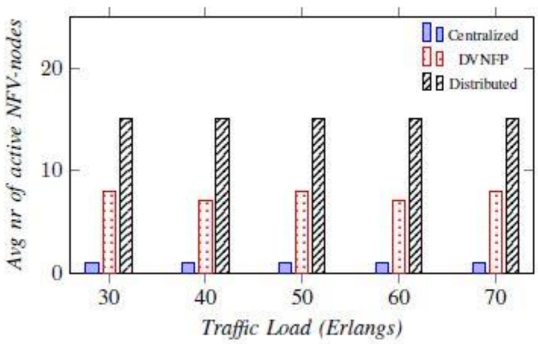

Figure 52: Average number of active NFV-nodes

Finally, in Figure 53 it is interesting to note that, although DVNFP requires the activation of fewer NFV-nodes in comparison to the distributed approach, and it provides a lower violation of the latency requirement. This is because after placing the VNFs, if the latency requirement of the SC is not satisfied, DVNFP performs additional steps to satisfy the latency requirement. Furthermore, in comparison with the centralized approach, DVNFP achieves fewer latency violations, as it is able to choose NFV-nodes based on the requirements of SC (i.e. for latency-sensitive SCs, nodes closer to the source are chosen).

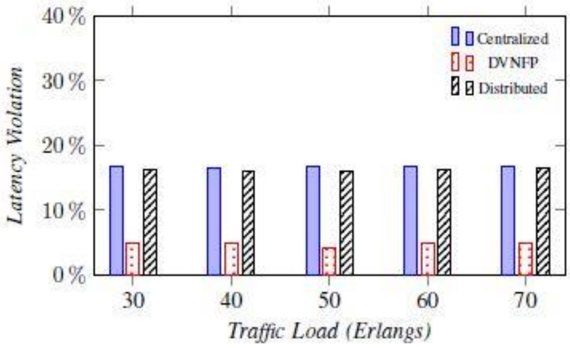

Figure 53: Latency violation

4.4.3 Dynamic and Survivable BBU Placement

In this section, we first describe the network settings we are considering in the evaluation. Then, we present the algorithm for the dynamic placement of Central Units (CUs) considering the functional separation of 5G eNBs in a metro-access network, and finally, we summarize the numerical results.

4.4.3.1 5G Network Settings

A three-layer functional separation of 5G eNBs (often called gNBs) has been identified and agreed in the context of standardization bodies recently [ITU-T-TN5G18]. As shown in Figure 54, these three layers are referred to as: 1) Remote (or Radio) Unit (RU), indicating the RRH at the antenna site; 2) Distribution Unit (DU) as the element including part of the digital signal processing, possibly providing a degree of functions sharing between several RRHs; and 3) Central Unit (CU), including higher layer (e.g., packet-based) processing, typically located at higher layers in a metro network and associated with several DUs. Besides fronthaul traffic exchanged between RUs and DUs, the midhaul traffic must be supported, which is exchanged between the DU and the CU. The mobile traffic is then backhauled towards the core network. In our work, however, we consider a two-layer separation of eNBs, where we assume co-locating the DU and CU within a single element, which we refer to it as CU, as also indicated in Figure 54. Correspondingly, only the fronthaul and backhaul traffic is considered in our evaluation.

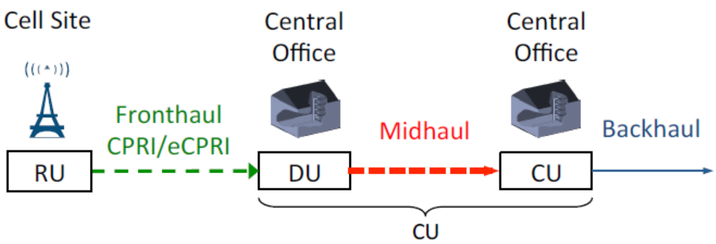

Figure 54: eNB functional separation in 5G networks. In this paper, we assume that the DU is co-located with its corresponding CU

Considering such a scenario, the dynamic placement of CUs is performed to enhance the utilization of processing and transport resources based on the spatial and temporal variation of traffic. For instance, on the one hand, in low-traffic conditions, several virtualized CUs can be placed in a centralized location located in higher network layers, known as a CU pool, to save on power and to enhance coordination. On the other hand, when traffic increases, CU pools can be placed at lower layers, i.e., closer to antenna sites, so as to avoid excessive insertion of fronthaul traffic. Thus, the ability to dynamically reconfigure the CU location allows network operators to achieve the desired balanced between baseband-resource consolidation and network capacity utilization. In our evaluation, we consider a multilayer OTN over WDM network as the underlying transport technology, so our algorithm must perform grooming, routing and wavelength assignment (GRWA) in an OTN over WDM aggregation network, and explore the interaction of the GRWA with CU placement to reach the objective of minimizing the average number of active pools, i.e., nodes hosting CUs, while achieving a satisfactory blocking probability.

Of course, adopting a multilayer OTN over WDM transport architecture to perform fronthaul traffic grooming has an impact on the latency between CUs and RUs and thus it affects CU placement; however, OTN technology is still considered a promising technology for fronthaul/midhaul transport. In fact, network operators are already working on optimizing OTN technology to fit with 5G service requirements, e.g., to reduce mapping latency from 10 ms to around 1 ms or less through the so-called Mobile-optimized OTN [RJing17].

4.4.3.2 Dynamic CU Placement/Handover (DCPH) problem

The Dynamic CU Placement/Handover (DCPH) problem in WDM access-aggregation networks can be stated as follows. **Given** 1) a hierarchical multi-stage access-aggregation network topology, represented by a graph G(N,E), where N is the set of nodes (including COs and Cell Sites (CSs)) and E the set of optical fiber links, and 2) random dynamically-generated backhaul traffic demands originated by CSs and directed to the Core CO, **decide** the placement/handover of CUs and the Grooming, Routing and Wavelength Assignment (GRWA) of backhaul and fronthaul traffic, while minimizing the average number of active pools in the network, i.e., the nodes hosting at least one active CU, **constrained by** 1) network links capacity (i.e. wavelength capacity and number of wavelengths per fiber) and 2) maximum fronthaul latency.

Note that, although only backhaul traffic demands are randomly generated and taken as an input of the DCPH problem, in general, once a CU location is selected for the RU source of the backhaul demand, one fronthaul traffic demand also has to be routed from the RU to the CU together with the backhaul demand between the RU and the Core CO. In this context, for a given backhaul demand originated by a CS, two special cases may arise according to the location selected for the CU, i.e.: 1) in the case where the CU is co-located with the RU, only the backhaul demand needs to be routed; 2) if the CU is located at the Core CO, only the fronthaul demand is routed.

4.4.3.3 DCPH Algorithm

In what follows, we describe the algorithm for dynamic CU placement/handover for CU consolidation. The algorithm consists of three different parts: 1) Optimal CU location identification; 2) CU Placement/Handover; and 3) GRWA.

1- *Identify optimal CU location*

The identification of the optimal CU location for a given demand is performed as follows. A list of candidate nodes is created to search for the optimal CU location for an arriving demand; the different solutions, i.e., the candidate nodes in the list, are sorted according to their cost, while considering a cost function that takes into account the activation of a new pool (i.e., in a node without other active CUs) to host the CU for the demand and the establishment of new lightpaths to provision the demand. Note that also trivial solutions, i.e., locating the CU at the cell site or at the Core CO, are also included in a list of candidate nodes.

2- *CU Placement/handover*

After computing the amount of required fronthaul traffic, which depends on the backhaul traffic, the list of candidate CU locations is scanned, starting from the first node in the list. First, the algorithm checks if a CU is already present in the network for the RU at CS. If such a CU is present, and it is already located at the optimum location (i.e., the first node in the list of candidate locations), the available capacity in the lightpaths already used between the RU and the CU (for fronthaul traffic) and between the CU and the Core CO (for backhaul traffic) is decremented by the amount of fronthaul traffic and backhaul traffic, respectively. In such a case, a trivial GRWA is performed for the demand, and the corresponding bandwidth values are deallocated from the lightpaths at the time of deprovisioning. Note that, if the available capacity in one or more of these lightpaths is not sufficient to provision the bandwidth requested by the demand, the demand is blocked, and the algorithm then considers a subsequent demand. On the other hand, if a CU is already present for the RU at CS, but its location does not coincide with the optimum location, CU handover needs to be performed. In this case, the GRWA for the demand takes place, and it is performed similarly to the case where no CU is already present for the RU at CS.

3- *GRWA*

In the case where a new CU is deployed or a CU handover is performed, the GRWA is performed. Note that, in the case where a CU handover takes place, besides the traffic for the demand into consideration, the traffic of all the existing demands originated by the RU at CS must also be considered at this step. This process, in general, involves the execution of GRWA for both the fronthaul and the backhaul traffic, and is performed on a shortest-path basis, also considering the possibility of using residual capacity of the existing lightpaths in the network, which are used to transport the traffic of other demands.

The first task of the GRWA step is to perform GRWA for the fronthaul traffic, due to the fact that the fronthaul has more stringent requirements in terms of latency and required network capacity. The main cost metric used in our algorithm is the hop count. However, to favour the utilization of the residual capacity in already-provisioned lightpaths, costs are assigned to a given lightpath-edge by considering the number of physical links it traverses, divided by two. Moreover, to discourage unnecessary grooming, we assign to grooming-edges a cost equal to 0.6. The value 0.6 allows us to break the tie in the case when applying the Yen algorithm, equal-cost paths are obtained between a short route where a new lightpath must be established and a longer route re-using existing lightpaths. Note that, when fronthaul traffic for a new demand is routed and there are already existing demands from the same CU, the different fronthaul flows can be transported along parallel lightpaths between the RU-CU pair.

If the fronthaul latency budget is respected for the demand considered and for all the existing fronthaul flows possibly affected, GRWA is performed also for the backhaul traffic. Note that, performing traffic grooming for fronthaul and/or backhaul flows of a demand may affect existing fronthaul flows. Therefore, every time a GRWA is performed for a demand, the fronthaul latency budget is checked, not only for the current fronthaul demand, but also for the other existing fronthaul flows, which may be affected due to the switching latency contribution introduced when performing traffic grooming. Moreover, in the case where the GRWA solution cannot be used due to the violation of a latency constraint, the first solution is removed from list, and the subsequent solution is analyzed. In the case where no solution is found, the current candidate CU location is removed from list of candidate locations and the subsequent candidate CU location is analyzed, i.e., the process is repeated.

If no solution is found for any of the candidate locations, e.g., due to the lack of network capacity and/or the violation of the fronthaul latency constraint, the demand is blocked. Conversely, if a solution is found for the demand, the corresponding backhaul and fronthaul traffic is deprovisioned once the connection duration of the demand passes, and, in the case when the used lightpaths are not used for any other demand, such lightpaths are torn down.

4.4.3.4 Case Study and Simulation Settings

To perform our numerical evaluation, we developed a C++ event-driven simulator, where we randomly generate the arrival of 55,000 demands originated by the RUs. Arrivals are generated according to a truncated-Poisson distribution, used to capture the fact that CSs support a limited backhaul traffic, and are uniformly distributed among RUs in the network. Demands duration is assumed to be exponentially-distributed with mean = 1 second. We consider a 5G HetNet scenario, where 80 nodes, consisting of 50 Macro CSs (MC) and 30 COs also inserting mobile traffic, cover a square region of 200 km^2 and are interconnected via a 4-stage topology as shown in Figure 55. Each MC is interconnected via a lower-layer tree to 10 Small Cells (SCs), not shown in Figure 55 for the sake of figure clarity, via optical fiber links at a maximum distance of 2 km. Each fiber supports 8 wavelengths at 100 Gbit/s each, in line with [InfineraMetro]. This scenario follows the guidelines of a 5G urban mobile aggregation network, as identified in [ComboD3.4].

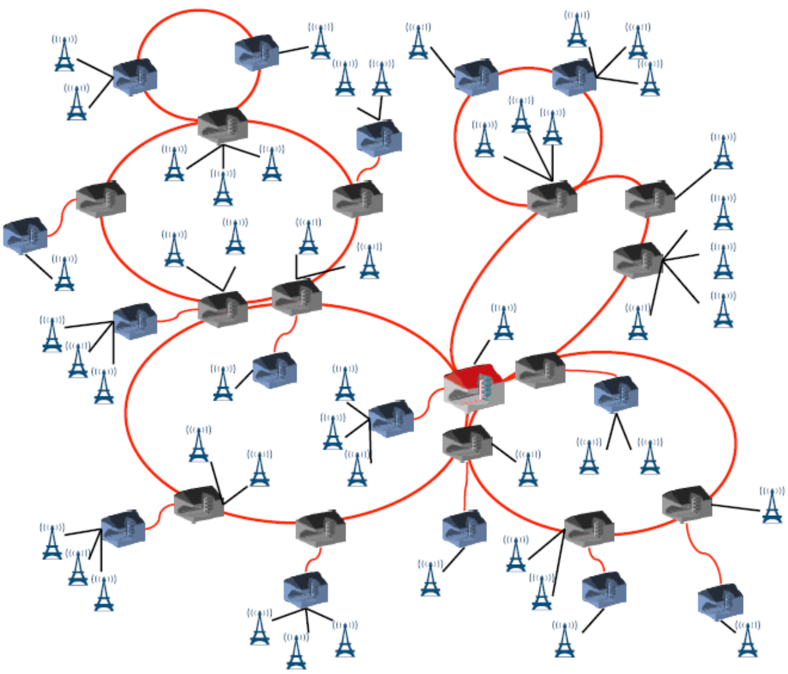

Figure 55 Hierarchical access-aggregation network architecture

MCs are assumed as 3-sectored sites with maximum backhaul traffic of 15 Gbit/s each, corresponding to an antenna configuration with 125 MHz spectrum, 256 QAM and 8x8 MIMO. We consider the same configuration for SCs, though we assume single-sectored sites, thus requiring a maximum of 5

Gbit/s traffic. Each demand requires a fixed bandwidth of 300 Mbit/s backhaul and, as we assume RAN split option II, the corresponding fronthaul traffic is 1.2 Gbit/s leading to a maximum fronthaul of 60 Gbit/s and 20 Gbit/s per MCs and SCs, respectively. The maximum tolerated latency for the considered RAN split is set to 100 microseconds. The choice of the RAN split is motivated by the fact that, among the eCPRI splits with fronthaul traffic proportional to backhaul, eCPRI split II enables the highest degree of functions centralization. Note that, considering a RAN split with backhaul-proportional fronthaul traffic allows us to evaluate the importance of traffic grooming when solving the DCPH problem.

4.4.3.5 Numerical Results

To evaluate the effectiveness of the proposed approach, which we refer to as MaxC-H (Maximum Centralization with CU handover) and which is characterized by a mode of operation where CUs can be moved during their activity, i.e., performing CU handover, to an algorithm, which we refer to as Adaptive, where a CU location could not be modified during operation (e.g., if it is receiving traffic from an RU).

As a preliminary result, we evaluate the performance of the developed algorithm by considering the average number of active pools metric, P_{av}, as shown in Figure 56. To better capture the performance difference of the two algorithms, we show the *Normalized* P_{av}, i.e., we normalize the average number of active pools with respect to the number of provisioned demands by the two algorithms.

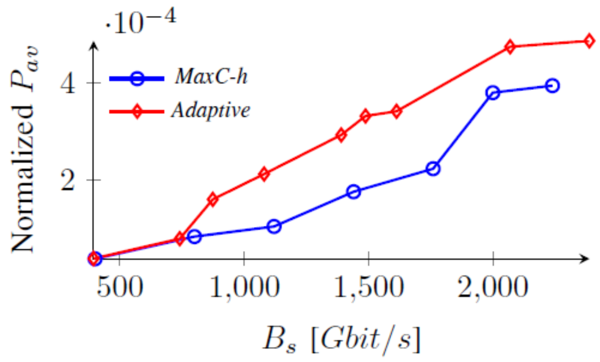

Figure 56 5G HetNet topology used for the numerical evaluation

As shown in Figure 56, MaxC-h always provides a lower number of active pools per demand, mainly due to the possibility of performing CU handover in the case where it is convenient to improve CU consolidation.

MaxC-h and Adaptive have comparable performance in terms of normalized P_{av}, only for lower served traffic, confirming that MaxC-h is better able to adapt to the dynamic changes of network traffic behaviour. In other words, this demonstrates that the MaxC-h algorithm is able not only to reduce the number of active pools, but also supports more users traffic thanks to the opportunity of moving CUs and consequently reduce the amount of fronthaul traffic which might lead to network congestion. As a matter of fact, for the considered arrival rates, no demands are blocked in the MaxC-h case. Conversely, the Adaptive algorithm exhibits a higher blocking behaviour, i.e., of the order of 20% higher, even for medium traffic (e.g., 20 Gbit/s per RU).

4.4.4 CDN Dimensioning for Cache/Bandwidth Balancing

In this section we provide a discussion of the cache deployment in a hierarchical metro-area network. As mentioned in Section 4.2 of [D2.1], in the vision of the Metro-Haul network architecture, the CDN operator caches popular video contents at AMENs and MCENs for several objectives, such as offloading traffic from the metro network and meeting end-user requirements. In this context, we focus on optimally deploying and dimensioning caches located at AMENs and MCENs such that the overall network bandwidth consumption due to video-on-demand content delivery is minimized.

To provide a quantitative estimation of the bandwidth consumption due to video delivery, a *cache deployment,* for a given network topology, needs to be adopted, in addition to a case study, characterizing a geotype scenario (i.e., number of users and average video delivery bit-rate) and a video-content catalogue (i.e., number of video contents, average video content size and popularity distribution). While the geotype scenario and the video-content catalogue remain fixed for a given case study, the *cache deployment* is variable and depends on the budget invested by the operator. With *cache deployment,* we refer to the deployment of caches, in terms of the number of caches, their location and dimension (i.e., storage capacity). Indeed, an optimal cache deployment, for a given investment, is a deployment that minimizes the overall network resource occupation, i.e., a deployment which minimizes the overall bandwidth requirements for video-delivery.

In the following, we first present a discussion on the optimal cache deployment for a given investment in terms of storage capacity for the case study described in [D2.2], and then, based on the cache deployment considered, we quantitatively estimate the bandwidth requirements to deliver the required video-content from AMENs and MCENs.

4.4.4.1 Optimal Cache Deployment

Given a budget-constrained investment represented by a maximum overall amount of storage capacity to be deployed, the network topology (i.e., number of AMENs, AMENs rings and MCENs), potential location of caches, and characteristics of the content catalogue (catalogue size, popularity distribution), we find the optimal cache deployment such that the overall average Resource Occupation (RO) in the network is minimized. Similar to previous work (e.g. [HSyed14]), we assume the average hop-count as the main metric to estimate the overall RO, where the RO is assumed to be the product of the average hop-count and the average bit-rate. To solve this problem, we develop and use an event-based dynamic simulator for VoD content caching and distribution and compare different cache deployments by varying the storage capacity distribution among caches at different network levels (at AMENs and MCENs) to find the optimal cache deployment. The simulator generates video requests according to the VoD content catalogue popularity model and provisions them according to the implemented cache deployment strategy.

As a case study, we consider a content catalogue characteristic similar to the one described in Section 4.3 of [D2.1], where the available storage capacity to be deployed = 160,000 GB and the network topology consists of 32 AMEN caches uniformly-distributed over 4 AMEN-rings (8 AMENs per ring) and 1 MCEN cache. We perform different simulations for different cache deployments. In the simulations, we vary the amount of content stored in caches deployed at AMENs from 0 (the case where all the storage capacity is utilized in the cache located at the MCEN), to the maximum possible amount of content to store in each of the AMEN caches, i.e., the case where all the storage capacity is utilized in the caches located at the AMENs. In each simulation, we simulate the arrival of 400,000 VoD requests, assumed as Poisson-distributed, at an arrival rate guaranteeing negligible blocking probability, to provide a fair comparative analysis between the different cache deployments. We denote by $RO_{avg/req}$ the average resource occupation of a video request per second under a given storage capacity distribution, and which is represented by the product of the average number of hops and the average bit-rate of all video requests. Note that, since the number of caching nodes at different network layers varies (the network topology consists of more AMENs than MCENs), storing content in the metro-aggregation level (i.e., at AMENs) utilizes more storage capacity with respect to the metro-core level, i.e., at MCEN (where only one copy of the content needs to be stored).

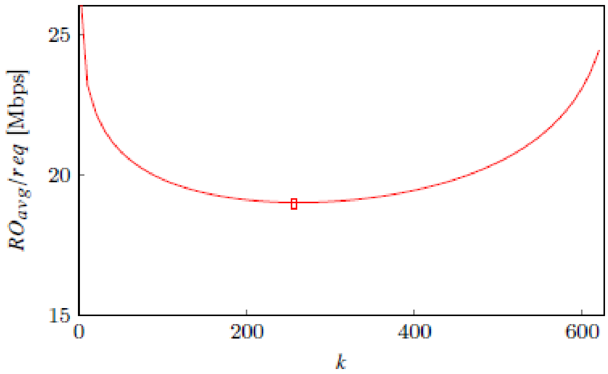

Figure 57: ROavg/req concerning the number of contents stored in the AMEN caches for the considered case study

Figure 57 shows $RO_{avg/req}$ (i.e., the average resource occupation per VoD request) as a function of the number of contents stored in the caches hosted by AMENs, denoted by k. Results show that $RO_{avg/req}$ initially decreases as k increases (as more contents stored in the caches located at AMENs allow to serve more requests from locations near end-users) until a certain value of k, after which *RO* increases again (as it becomes less-advantageous to deploy more storage capacity in the caches of the AMENs and more-advantageous to deploy the storage capacity in the MCEN cache). Why is the optimal solution not to deploy storage capacity at AMENs? This is due to the fact that, when the storage capacity is limited, it becomes more-advantageous not to store duplicates of several popular contents at AMENs, but rather to store one copy of a larger set of content, thus pulling more content from the origin server into the network.

Table 36: Values of K*, storage capacity and hit-ratio of AMENs and MCEN caches

*k**	AMENs Cache (GB)	AMENs Cache hit-ratio	MCEN Cache (GB)	MCEN Cache hit-ratio
267	2136	0.45	91648	0.46

In Table 36 we show the value of *k**, i.e., the number of contents that, if stored in the AMEN caches, guarantees an optimized cache deployment, the storage capacity and the hit-ratio of the caches located at AMENs and MCEN for the resulting cache deployment. This shows that for a given investment in terms of cache budget, and a given case study, in terms of network and services characteristics, there exist an optimal deployment and dimensioning of caches which yields a minimal resource occupation. In the following paragraph, we use this optimal cache deployment, considering the hit-ratio of the caches deployed at the AMENs and the MCENs, and approximate the bandwidth requirements for the CDN use case.

4.5 Evaluation of Autonomous Networking and E2E Orchestration

This Section overviews the Metro-Haul control-plane architecture in greater detail, tying it to the KPI evaluations and how they can be modelled in the techno-economic frameworks. It should be highlighted that, by its very nature, it is not straightforward to include such issues directly within the cost/capacity/power consumption scope that is the main objective of this deliverable. Nevertheless, some of the proposed evaluations, specifically the network-level studies in section 4.4, can be augmented with specific considerations of the Metro-Haul COM and its subsystems. For this reason, the following descriptions are provided.

4.5.1 Control-Plane Architecture, Orchestration and Slicing

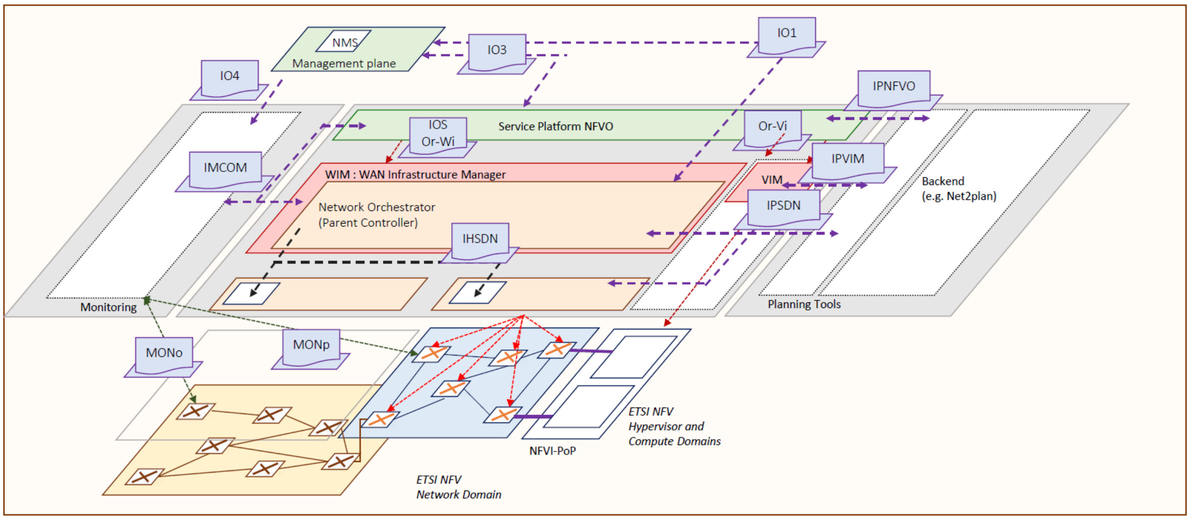

Figure 58: Metro-Haul unified service platform

The Metro-Haul COM system (See Figure 58) is responsible for the dynamic provisioning of services. The COM is defined spanning the following layers: i) the **(Network) Control Layer** addressing, mainly, the SDN control of the network infrastructure (hierarchical for several network domains/layers) and the control and dynamic provisioning of Computing and Storage resources (in that sense, specific elements of ETSI NFV architecture such as the VIM component are considered as part of the control layer; ii) the **MANO (Management and Orchestration) Layer** refers to specific components of the ETSI NFV architecture that enable the deployment of VNFs and VNF-FGS across the Metro-Haul infrastructure. This terminology is inherited from the ETSI normative references. Finally, iii) the **Slicing Layer** is a layer that enables the deployment of Slices (self-contained logical networks that are service-tailored and that can be controlled independently) over the MANO layer.

Regarding the network control, the COM follows the IETF ACTN (Abstraction and Control of Traffic Engineered Networks) framework, to control the infrastructure that spans across several network domains that include the optical layer, the packet layer and PON access networks. It relies on the functional elements defined therein, and the macroscopic architecture can be seen in Figure 59.

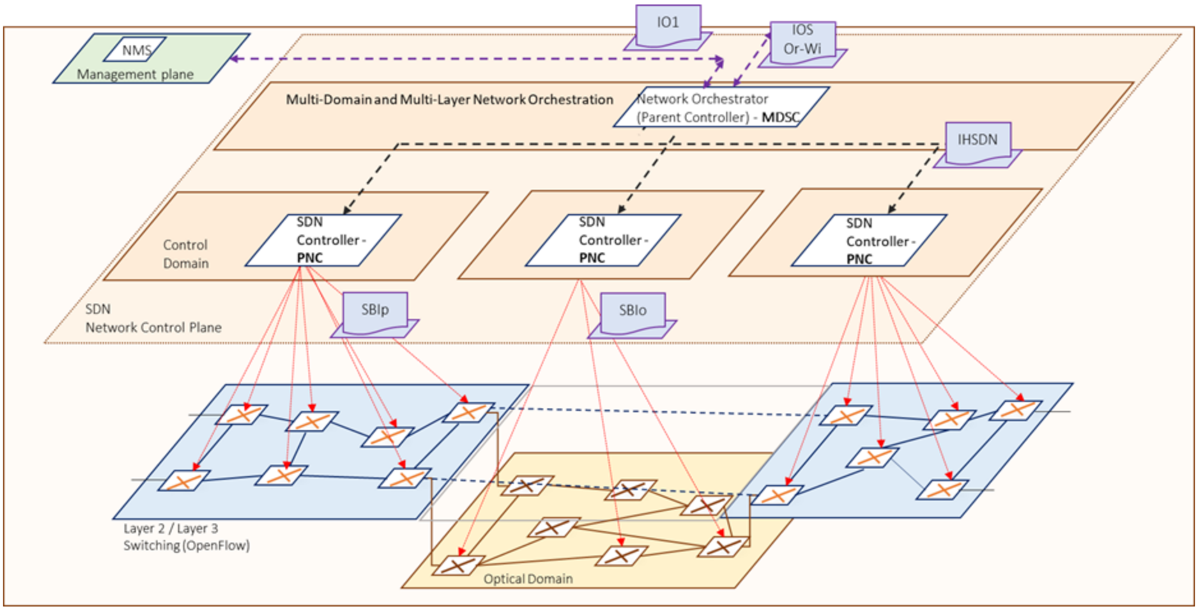

Figure 59: METRO-HAUL Network Control layer

4.5.1.1 Network Layer

4.5.1.1.1 Optical Domain

The Optical domain SDN Controller is a single component, controlling the establishment of optical channels across distant O-NEs (Optical Network Elements) and which exposes a Northbound Interface (NBI) to the parent SDN controller to instantiate connectivity services. The local node controller agent implements a NETCONF server that is accessed by the Optical domain SDN Network Controller to configure the individual components, as depicted in the following figure for the case of the partially disaggregated optical network.

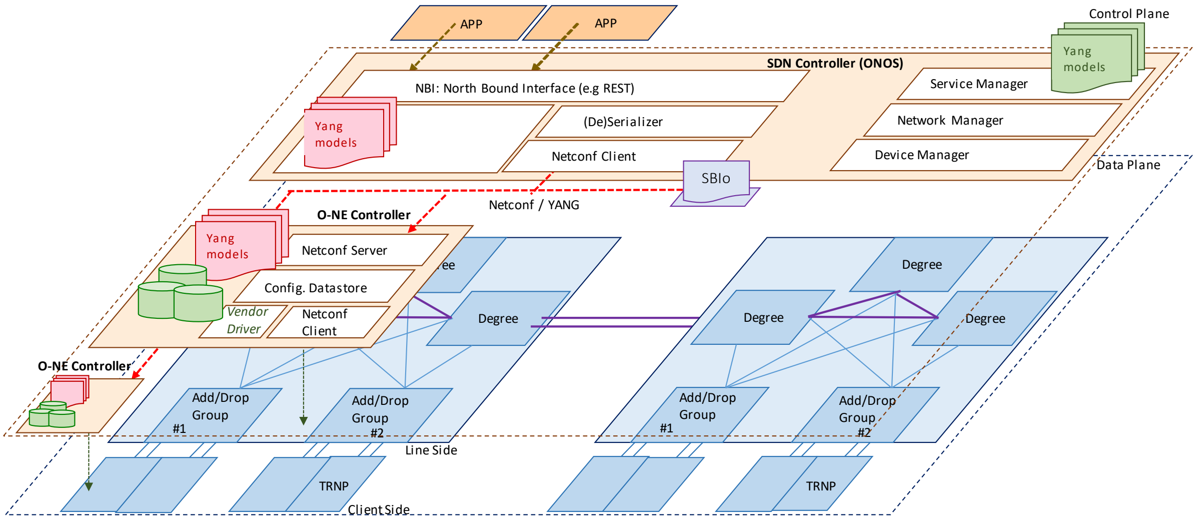

Figure 60: NETCONF / YANG control of the partially disaggregated optical network

4.5.1.1.2 Packet Domain

A packet domain is deployed to connect the compute and storage infrastructure present in the Metro-Haul nodes to the optical domain, to get connectivity services to remote nodes, and to provide local connectivity services between local compute and storage nodes.

4.5.1.2 Orchestration Layer

This is designed to coordinate the joint orchestration of the cloud as well as network resources for the implementation of NFVs. PoPs are interconnected using a network infrastructure, while the orchestrator is responsible for instantiating VNFs.

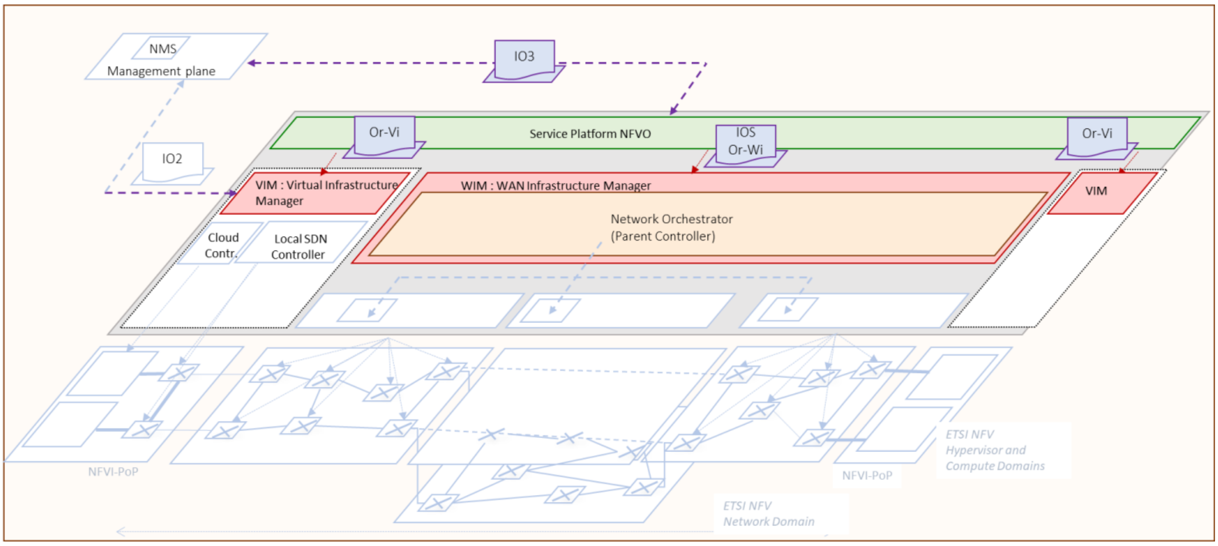

Figure 61: Metro-Haul Orchestration layer based on ETSI/NFV MANO

Figure 61 shows the Metro-Haul orchestration layer that is based on the design of ETSI NFV Management and Orchestration (MANO). In this case, the orchestrator is referred to as the NFV Orchestrator (NFVO) and coordinates with various components to orchestrate an end-to-end network service. The NFVO orchestrates the network resources interconnecting the AMENs and the MCENs using the WAN Infrastructure Manager (WIM). The WIM is responsible for provisioning connectivity paths between the VNFs in a Wide Area Network domain. It relies on / or includes the parent SDN controller which controls multiple heterogeneous network domains. Note that the Placement, Planning, and Reconfiguration Subsystem (namely, the Network Planner) enables the optimization of resource allocation in the optical metro network to effectively provision VNFs in specific computing nodes considering heterogeneous requirements.

4.5.1.3 Slicing Layer

Network slicing is related to the concept of network virtualization, including the partitioning (slicing) of a single (commonly physical) infrastructure to construct multiple (logical) infrastructures; but there are differences that are worth highlighting: emphasis is given to the actual network functions and how they are arranged and configured, forming a complete logical construct or network, tailored, customized and optimized for a given service or service set, or to support a given actor or customer (e.g. vertical industry). They combine both data and control plane functions and functional elements, which are an inherent part of the slice.

Metro-Haul relies on the ETSI NFV framework as a starting point for a concrete implementation of generic slicing architecture, in which network slice instances are NFV Network Services (NS), encompassing NS endpoints and one or more VNFs interconnected by logical links, forming VNF Forwarding Graphs (VNFFGs). Logical links are thus mapped to supporting network connectivity services which may, in turn, span multiple network segments. This is shown in Figure 62, where multiple NFVO (green, blue), potentially managed by different users or operators can have shared access to a common NFVI managed by their respective VIM/WIMs, and each NFVO instantiated network service (with its corresponding VNFs) is a slice instance.

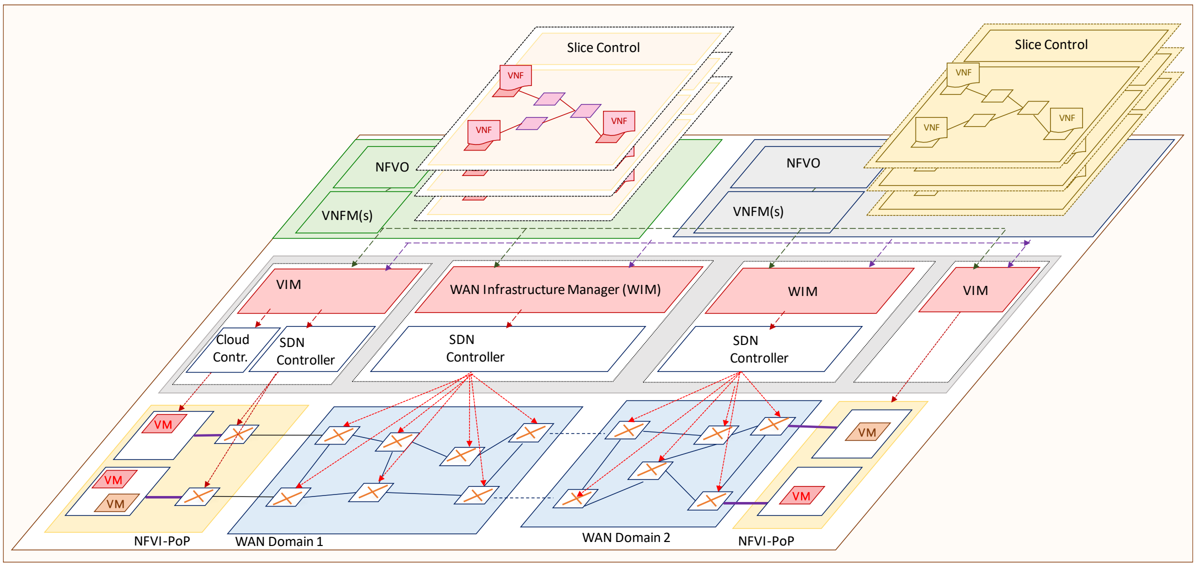

Figure 62: Network Slicing using the integrated SDN/NFV framework.

4.5.2 Evaluation of Metro-Haul COM KPIs

This section complements Section 2.1 regarding the Metro-Haul KPIs that are specific to the COM subsystem. The evaluation of these KPIs is carried out in the scope and context of WP4 activities and will be reported in D4.2 and subsequent deliverables. In the next sections, we detail for each of the KPIs, the methodology to be used, along with its implications. A significant effort is used in characterizing the service setup time, which is a key indicator to demonstrate the OPEX savings associated to automation, especially when compared to long-lasting quasi-manual processes that characterize service provisioning in current networks.

4.5.2.1 MH1. Optical PtP connection set-up time

This KPI is by definition the latency perceived by a user when requesting a lightpath or optical connection, from a transceiver client port, to another transceiver client port. Since it involves the optical hardware as well as SDN control components, this is a WP3-WP4 joint KPI. The process to measure the KPI can be automated, and is defined as the time passed between the reception of a request in the SDN controller NBI and the moment the lightpath is considered active.

The NBI can be in terms of a TAPI connectivity request specifying the involved Service Interface Points (SIPs) and includes:

- The parsing and processing of the request, by the SDN controller, mapping the TAPI Service Interface Points to a given transceiver client port, identified by the SDN controller as a Connect Point.
- The creation of an optical connectivity intent between the two connect points, along with any applicable constraint.
- The computation, by the SDN controller, of the end-to-end lightpath taking into account resource constraints and the capabilities of the transceivers and the network elements, including tunability constraints.
- The deduction of configurations to be established at the different network elements. Such configurations are sent to the network elements via the Netconf sessions that are assumed existing and stable.
- The report of the operation to the user via the NBI.

There are aspects that may affect the KPI:

- The transmission latency and bandwidth between the SDN controller and the element that is controlled by it. Although in lab trials this latency is of the order of a few milliseconds, real deployments need to account for the latency associated with, e.g., a dedicated VPN to access the network elements.
- A lightpath may need subsequent validation (along with additional processes related to e.g. power management) before it can be considered "active". For the purposes of the Metro-Haul target scenarios, a lightpath is considered active upon reception of the acknowledgement from each of the network elements along its path.

The key components involved are:

- The SDN controller of the optical Layer 0 / Layer 1 network.
- The agents of the network elements, transceivers, open line systems or ROADMs.

4.5.2.2 MH2. Metro-Haul E2E PtP connection set-up time

This KPI is strongly related to MH1 and extends the connection setup time to account for the fact that end-to-end services involve a multi-layer network (packet switching over circuit switching) and that there is a hierarchy of controllers involved in the service deployment.

Similar to the previous case, the latency in MH2 is defined as the time between the reception of a request from the parent controller north bound interface and the moment the connection is provisioned. In particular:

- The parent controller receives a request in its NBI. Although for the optical domain TAPI is the selected interface, in the case of the parent controller there also exists the possibility of using a Layer 2 Service Model (L2SM) to request packet connections with constraints in terms of, e.g. VLANs between endpoints.
- The parent controller performs a basic computation taking into account the domains involved (packet and optical) and may query the underlying controller. The parent controller may reuse an existing optical connection for subsequent packet connections as long as the different capacity constraints are respected.
- The request, from the parent to each of the involved child controllers. This can be done using TAPI (e.g. see MH1) or the controller native interface. Similar considerations to the case in MH1 apply at each domain.
- Once each element is configured, the different controllers report the status or failure of the operation to the parent controller.

As in the previous KPIs, there are similar aspects that may have a direct impact on the KPI values. It is important to qualify the scenario in which the performance is evaluated to get an insight into the applicability and order of magnitude/relevance of the measure.

A potential optimization of this KPI is defined by the fact that in an SDN context, a given SDN parent controller can send the requests to the underlying child controllers in parallel, thus benefiting in a reduction in setup time. This is at the expense of more complex logic that needs to account for potential failures and be able to rollback in such a case.

For MH2 KPI, the involved elements and components are:

- Those of MH1. It is worth mentioning that the establishment of an end-to-end connection of interconnecting VNFs relies on optical connectivity across the Metro-Haul segments.
- The parent controller, which coordinates the operations.

4.5.2.3 MH3. Set-up time of network service slice across Metro-Haul

MH3 captures a complex service as defined within the scope of WP4 COM. We consider a slice as an ETSI NFV network service (a set of interconnected virtual functions), and the KPI targets the measurement of the service deployment.

For the purposes of the KPI the following hypotheses are made:

- We assume that the network service and VNF descriptors are previously configured at the level of the NFV-O.
- The slice shall contain at least 2 VNFs, each of which has to be allocated in different NFVI PoPs that correspond to different Metro-Haul nodes.
- The slice is considered active when all the VNFs supporting the business logic are active and their interconnection supported by one or more connections

Consequently, the KPI is defined as the time it takes between a user or operator to request the instantiation of a network service/slice and it becomes active. It is clear from the previous descriptions that this KPI is a superset of the previous two ones. It is expected that at least one end-to-end connection needs to be established for this purpose.

For the evaluation of this KPI, the following components are involved:

- The net2plan network planner, which takes a request from the user or operator as specified in the descriptors of the service and computes the best placement. In this sense, the running time of the algorithm contributes to the service setup delay.
- The OSM NFVO, which receives the request and instantiates the different VNFs across the Metro-Haul infrastructure. It decides which logical links need to be supported, so that it requests such connectivity to the parent controller (see MH2). The setup delay is then affected by the time it takes to instantiate the VMs and the connectivity.

A target for this service is of the order of minutes, but it may vary depending on the number of VNFs, their size and complexity, etc.

4.5.2.4 MH4. Capacity of Metro-Haul controller

The adoption of a centralized control plane and SDN principles for the control of the disaggregated optical network raises questions regarding the scalability of the solution. Assuming a single centralized controller, MH4 KPI aims at evaluating the following aspects:

- How many Netconf devices can be controlled from a single instance of the controller (assuming ONOS SDN controller running or medium server hardware). It is a design target of the Metro-Haul solution to support the control of 10 – 100 nodes (AMENs/MCENs, i.e. basically Open Disaggregated ROADMs)
- What is the control plane overhead that this imposes?

For the former, the KPI will be measured by progressively increasing the number of Netconf devices (including the use of device emulation) and measuring the latency of the controller in performing basic operations.

For the latter, network analyzers will be used to estimate the overhead (e.g. in terms of packets per second or required bandwidth) between the controller and the devices, obtaining basic guidelines and recommendations for its usage.

4.5.2.5 MH5. Fault/degradation detection time

In contrast with other time-related KPIs defined above, the concept of fault/degradation detection time requires some previous definitions. Specifically, we focus on optical connection failures defined as events that negatively affect (e.g., degrade or interrupt) established connection/s. We can distinguish between:

- *Hard failures:* unexpected events that suddenly interrupt the connection/s.
- *Soft failures*: events that progressively degrade the quality (e.g., quality of transmission) of the connection/s.

Because of the definition, hard failures can be only detected after they have happened, whereas soft-failures can be anticipated. Therefore, our focus is on detecting soft-failures/degradations before they become hard-failures and cause the connection/s to be disrupted. Specifically:

MH5 focuses on detecting connection degradation while leaving enough time to implement counter-measures (e.g., tuning device parameters or re-route connections) to avoid such a degradation becoming a hard-failure.

Note that the duration of the degradation (from the time it starts until it becomes a hard-failure) can be very variable: from few milliseconds (real-time), e.g., in the case of transmission-related effects to several days/weeks/months, e.g., in the case of device ageing. In this regard, the following operational scenarios are defined for the quantification of this KPI:

1. *Real-time* (instant degradation): In this scenario, we are targeting soft-failures that occur in a short time (millisecond) scale. As an example, fluctuations of physical parameters, such as the state of polarization (SOP) rotation can noticeably impact pre-FEC BER. Then, by the anticipated detection of SOP fluctuations, an algorithm could configure the local *soft-decision* FEC module to ensure robust optical transmission. Another example could be the detection of a condition that cannot be solved locally but needs a mitigating action in a remote device, like changing the modulation format. Note that communication between two distant devices (through the control plane) is needed in this case, which would entail longer times.
2. *No Real-time* (gradual degradation). Typical examples include laser and filter failures. The detection of those failures can trigger e.g., device parameter tuning, network re-optimization to re-route the affected connections, as well as scheduling preventive maintenance.

It is clear that for being able to detect soft-failures, the monitoring/telemetry interval, i.e., the time to collect performance measurements must be on the same time scale as the failure that is being targeted. In this regard, we assume that the monitoring/telemetry interval is, at least, such that it provides a minimum number of samples (e.g., 10) within the target KPI time. Table 37 describes the quantification of the KPI for the identified operational scenarios and the required monitoring/telemetry interval:

Table 37: Monitoring/Telemetry intervals per operational scenario.

Scenario	***Counter-measure***	***Target anticipation time***	***Monitoring / telemetry Interval***
Real-Time (instant degradation)	Local node controller parameter tuning	*> 100 msec*	10 msec
	Remote node controller parameter tuning	> 10 sec	1 sec
No Real-time (gradual degradation)	SDN controller connection re-routing	It might vary from minutes in case of parameter tuning to 1 hour in the case of network re-optimization + connections set-up time	1 min

For the evaluation of this KPI, the following components are involved:

- The node/device controller is involved in general, for the collection of measurements, as well for the detection of real-time degradations.
- The SDN controller might be involved in both cases of real-time and non-real-time degradations.
- The monitoring and data analytics (MDA) system is the main component for the gradual degradation scenario, where monitoring samples are collected from the network nodes and analyzed and, in the case that degradation is detected, a recommended action is issued to the SDN controller.
- The network planning tool can be involved in the case of network re-optimization.

5 Conclusion

This deliverable D2.3 has provided a detailed overview of the candidate architectures and methodologies to be employed in the techno-economic dimensioning of Metro-Haul solutions, and subsequent validation according to the project's KPIs. The structure of the deliverable mirrors the organization of task T2.3 within the project, which is responsible for defining both the network scenarios of interest, as well as the network architectures intended to address each scenario. The deliverable also covers the segmentation of the E2E Metro-Haul architecture into mostly self-contained domains that can be more easily investigated, providing building blocks to create Metro-Haul solutions suited to various reference deployment scenarios. Finally, special emphasis is also given to methodology aspects such as network dimensioning frameworks and algorithms, which are used in the scope of the various working streams to evaluate network performance across the most relevant KPIs: capacity, cost and power consumption.

The network performance evaluation relies on the creation of dimensioning workflows and reference network scenarios for analysis, built upon service use-cases defined in the previous deliverables. In this document, four support tasks are detailed to provide the environment for E2E performance evaluations. The first one covers the dimensioning of AMEN/MCEN nodes according to their geographic coverage setting, defining rural, suburban, urban and dense urban geotypes modelled according to data from the operators in the consortium. These geotypes define the coverage of legacy central office nodes repurposed as integration of metro aggregation sites and edge computing platforms. Each geotype defines a range of physical areas and covered households. It further defines the profile of the services reaching the node (e.g., urban/dense urban settings are associated with a higher share of enterprise traffic). This input therefore provides the means to do local node dimensioning, based on the traffic aggregation requirements from the access side. Additionally, this sub-task also defines the structure of the reference networks to consider. The chosen methodology splits the analysis into two aspects: for network-level studies involving traffic optimization aspects and placement of distributed VNFs across the entire metro infrastructure, realistic reference topologies provided by consortium partners are utilized. For device/subsystem-level architectural studies, such as evaluating different express node architectures to be used, a more abstracted approach is taken, consisting of subdividing a Metro-Haul domain into a metro aggregation segment comprising a chain/horseshoe of AMENs bounded by two MCEN nodes, and a metro-core domain featuring a mesh topology interconnecting MCEN nodes and interfacing with the photonic core network. This split enables physical architecture solutions to be studied in a parametric approach, defining target ranges for metro aggregation and metro core solutions (e.g. linked with the defined geotypes), in terms of link/path lengths, the required number of optical express nodes etc.

The second support task pertains to the modelling of traffic from the macro level (traffic growth assumptions over months/years), to the micro level (breakdown of traffic into the different services, daily/tidal fluctuations around peak values, etc.). The framework presented here describes an end-to-end methodology for traffic generation, as the input for various types of evaluations. On a high-level, traffic is subdivided into “mass market” background traffic and 5G vertical use-cases. The former encompasses mobile/fixed access traffic, from residential or enterprise sources. Based on monitored data from network operators, coupled with the geotype modelling assumed, this data provides the peak traffic volume subdivided by high-level traffic types. These types are basically defined by the type of connectivity they require, most notably by whether they require local processing (e.g. based on their latency strictness) at the AMEN/MCEN itself, if they can be forwarded towards another data-center within the same Metro-Haul domain, or if their destination is in the core network (i.e, outside the Metro-Haul domain). On top of this background traffic, specific vertical use cases, covered in detail in deliverables [D2.1] and [D2.2] are specifically modelled over the Metro-Haul network to evaluate the responsiveness of the architecture to their requirements.

Based on the high-level traffic dimensioning and segmentation, two main steps are used to produce dynamic traffic traces at the intended time granularities. The first one applies known daily/weekly/monthly traffic variation profiles to the high-level peak-value based volumes. This is particularly useful in the scope of Metro-Haul, as one of the main advantages of the control-plane driven, highly flexible and distributed approach envisioned within the project are the efficiency gains achieved from adapting the infrastructure to real-time requirements, rather than dimensioning it for peak values. The second step involves modelling the effect of a group of users for different services on the aggregate dynamic traffic requirements, at almost arbitrarily low granularities. This framework, based on the CURSA-SQ methodology, enables the generation of synthetic traffic traces per service and user profiles, bridging the gap between local user level behaviour and the aggregate traffic variation at a node’s input/output. The traces produced by this methodology can be directly used, at the desired time granularity, by network optimization and resource allocation algorithms addressing dimensioning studies.

Following this workflow, the following component in the ecosystem is the availability of a common method to model and use input data, as well as of implementing and cross-checking the validity and performance of planning/provisioning algorithms. This role is taken by the Net2Plan open source planning tool, which in Metro-Haul serves a dual role as the backend planning/placement subsystem interfacing with the orchestration layer, but also as an offline tool for network dimensioning by providing a basic layer of data structures and optimization capabilities for algorithm developers to build upon. This latter role takes special importance in the scope of this task, as it enables multiple partners to collaborate on the same platform, facilitating dissemination and validation of results. Additionally, given the open source nature of the project, it also enables external entities to build upon evaluation scenarios developed for Metro-Haul, to replicate scientific findings, and to accelerate development. In the concrete scope of task T2.3, the Net2Plan tool was augmented with the NIW library, providing base functions to easily model VNF placement problems over optical transport networks, and more generally enabling network dimensioning with joint optimization of optical/IT resources. The interfaces were developed to cover the specific needs of the project, such as native modelling of VNF chains over optical networks, network service descriptors, time-varying traffic, etc. These modelling capabilities, along with native algorithm libraries, can be extended to specific optimization use-cases in the scope of the architectural analysis, mainly for the network-level performance analysis related to edge computing and E2E orchestration use-cases.

Finally, as the basis for the techno-economic evaluations in Metro-Haul, it has become necessary to build a model for the cost and power consumption of the devices/subsystems used to deliver an E2E solution. Given that: 1) some of the components used are prototype-stage devices, and 2) existing optical component cost/power models are reasonably outdated, it was decided to insource a cost/power model within the consortium for the optical layer components required. This model is complemented with a data-center and packet model for other components, where the figures are taken from a mix of insourced estimates provided by partners with actual field experience, and available models in the literature covering these types of components.

The support tasks are used to create an overall framework that enables multiple types of architectural analysis. These analysis are subdivided into various working streams, focused on specific aspects of the Metro-Haul architecture. The first one pertains to the architecture of the central offices, and how it evolves to support a higher heterogeneity of traffic sources with differentiated requirements. A set of six different scenarios were identified with respect to how a central office can be deployed to interconnect access traffic to both the optical transport and the data-center resources co-localized at each AMEN/MCEN. These scenarios mostly vary according to: 1) the amount of layer 3 processing that can be offloaded from physical switches/routers to VNFs in the DC node, and 2) the interconnection topology between the L2/L3 devices within the switching fabric (e.g., how access traffic, the DC node and the optical transport interfaces are interconnected). Based on the cost and traffic modelling work described in this deliverable, this analysis is expected to pinpoint target architectures for COs based on traffic requirements and hardware/software price points. Additionally, it will enable us to verify the suitability of components developed specifically within the project, such as FPGA-based programmable switches or compute nodes with native integration to packet/optical switching fabrics.

The second architectural aspect deals with the physical architecture of the optical nodes, covering both the line system (express and add/drop node architectures, amplifiers, etc.) and the transmission devices. In this evaluation, the analysis was subdivided into the two network tiers within Metro-Haul, assuming that differentiated solutions (in terms of cost, flexibility and capacity) will be needed for the metro aggregation and metro-core tiers. To this effect, this deliverable details the considered architectures in both cases. For the metro aggregation segment, nodes based on filterless, semi-filterless and ROADM-based technology were evaluated against reference metro aggregation implementations based on fixed-frequency filters. These different architectures offer trade-offs between deployment cost, throughput, power consumption and scalability that will likely define optimal target use-cases (in terms of fiber coverage and traffic requirements) for each of them. On the transceiver side, options based on direct-detection and coherent transmission were explored, again offering a trade-off between cost/capacity that must be bounded through dimensioning for specific network use-cases. In the metro-core case, a similar approach is taken, this time with a more restricted set of potential technologies, given a more demanding (and less variable) range of aggregated service requirements. Here, the focus lies mainly on the evaluation of flexibility in the add/drop layer and the use of PIC to drive down costs and power consumption of optical express nodes. The specific metro-haul devices targeted in this evaluation range from direct-detection and coherent transponders developed within the project, to PIC-based ROADM degrees with embedded amplification capabilities, as well as passive devices for filterless/semi-filterless node implementations that lower the footprint of network deployments.

In parallel with the node architecture and technology used in specific devices, Metro-Haul is also committed to supporting open and disaggregated systems that foster innovation and interoperability. However, the degree of disaggregation, particularly in the world of analog optics, encompasses some complex trade-offs between the potential deployment and operational advantages of disaggregating the optical transport, versus the integration costs that must be insourced to the infrastructure owner or some other third party, in order to maintain the same level of performance and reliability. To this end, the analysis in WP3 already defined a subset of potential disaggregation scenarios to be investigated: full integration/aggregation, partial disaggregation and full disaggregation. The evaluation in WP2 takes these subsystems as a starting point and implements an economic modelling that is independent of the technology, that is, it is assumed that the technical solution is the same (e.g., given by the analysis of the physical node architectures) in terms of devices, and the focus is on evaluating the economic aspect of deploying the same basic solution, if choosing between fully aggregated/disaggregated or a partial disaggregation solution. The modelling utilizes a parametric approach, defining integration/disaggregation costs and discounts to different deployment aspects of the optical node: hardware (line system and transceivers) and software (running in the devices and in the centralized controllers). The parameters for overhead/discount factors in each option, such as HW/SW integration costs, volume discounts, depreciation factors etc., are being insourced from within the consortium based on the experience from the different industry players (operators, system vendors and SMEs). The expectation is that this analysis indicates, with some degree of confidence, that conditions for network deployment that favour a particular disaggregation option (e.g., in terms of hardware volume required, traffic growth conditions, line system vs. transceiver ratio, etc.).

One of the broadest and most relevant work streams identified within the Metro-Haul scope is the evaluation of the techno-economic effect of supporting 5G-enabled edge computing platforms from the perspective of the optical transport network. While this goal roughly translates to being able to model IT and optical resources jointly in the same optimization frameworks, the wide array of service use-cases identified in [D2.1] and [D2.2] impose a very different set of requirements from both the storage/compute platforms and the optical transport. As such, the evaluation of specific vertical use-cases (as well as that of the general "background" traffic evaluation for capacity purposes) necessitates targeted optimization frameworks capable of modelling the intricacies of each scenario. The type of analysis to be performed in this work stream is exemplified by the frameworks present in this deliverable: a network-level analysis of DC resource requirements based on the latency budget of the overall traffic, in realistic metropolitan network use-cases; an optimization framework for dynamic assignment of BBU pools at different metro nodes according to shifting traffic requirements; a CDN bandwidth optimization use-case, outlining the trade-offs between storage/bandwidth across the optical network. Overall, these and other use-case specific optimization frameworks are targeted to be adapted and reused within the modelling scope outlined in this deliverable. Specifically, this involves considering the reference topologies of Metro-Haul, the traffic modelling framework, and, wherever possible, ensuring that the optimization algorithms are implemented and available in the open-source environment of the Net2Plan ecosystem.

The last work stream addresses, in a more indirect way, how the Metro-Haul COM affects the business case of infrastructure owners providing access to network/IT resources, as well as that of verticals looking for the best way to quickly deploy their services across a commoditized infrastructure. This evaluation looks at issues such as the controller and orchestrator scalability, the responsiveness of backend components such as the planning tool for resource optimization and allocation procedures, or the MDA subsystem for early detection of soft failures and respective preventive action, allowing the network to operate at a higher reliability threshold without overbuilding resources.

References

[3GPP-TS-22.261] 3GPP TR 38.801, V14.0.0, "Study on new radio access technology: Radio access architecture and interfaces", Mar. 2017

[5G-PPP-KPI] The 5G Infrastructure Public Private Partnership KPIs. Online: https://5g-ppp.eu/kpis/

[AEira18] A. Eira et al. "On the capacity and scalability of metro transport architectures for ubiquitous service delivery," *in Proc. International Conference on Transparent Optical Networks (ICTON)*, July 2018.

[AGCOM18] "Communication Markets Monitoring System", AGCOM (Italian Communication Authority) report, Q4 2018- [online]: https://www.agcom.it/documents/10179/3864782/Allegato+21-1-2019/dcdf6338-05f3-4589-a443-2a566f94b193?version=1.0

[AMozo18] A. Mozo et al. "Forecasting short-term data center network traffic load with convolutional neural networks", PLOS ONE, Feb 2018.

[ARao11] A. Rao et al, "Network characteristics of video streaming traffic," in Proc. CoNEXT, 2011.

[BNaud16] Naudt et al. "How can a mobile service provider reduce costs with software-defined networking?," Network Management, vol.26, no.1, January 2016.

[CiscoGCI16] "Cisco Global Cloud Index: Forecast and Methodology, 2016–2021", [online]: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/global-cloud-index-gci/white-paper-c11-738085.html

[CiscoVNI17] "Cisco Visual Networking Index: Forecast and Trends 2017–2022", [online]: https://www.cisco.com/c/en/us/solutions/collateral/service-provider/visual-networking-index-vni/white-paper-c11-741490.html

[ComboD3.3] ICT Project COMBO, Deliverable 3.3. [online]: http://www.ict-combo.eu/data/uploads/deliverables/combo_d3.3_pu.pdf

[ComboD3.4] ICT Project COMBO, Deliverable 3.4. [online]: http://www.ict-combo.eu/data/uploads/deliverables/combo_d3.4_wp3_17june2016_v1.0.pdf

[D2.1] D2.1: Metro-Haul Definition of Use Cases, Service Requirements and KPIs

[D2.2] D2.2: Metro-Haul Functional Architecture Specifications and Functional Definition

[D3.1] D3.1: Metro-Haul Selection of metro node architectures and optical technology options

[D4.1] D4.1: Metro-Haul Control and Management Requirements and Framework.

[EricssonMobRep18] "Ericsson Mobility Report", November 2018, [online]: https://www.ericsson.com/assets/local/mobility-report/documents/2018/ericsson-mobility-report-november-2018.pdf

[FGuio17] F. Guiomar et al. "Comparing different options for flexible networking: probabilistic shaping vs. hybrid subcarrier modulation," *in Proc. European Conference and Exhibition on Optical Communications (ECOC)*, September 2017.

[FJMor18] F. J. Moreno-Muro et al., "Joint optimal SC allocation, VNF instantiation and network resource management demonstration," *in Proc. Optical Fiber Conference (OFC)*, paper TU3D.10, March 2018.

[FJMor18-2] F. J. Moreno-Muro et al, "Latency-aware optimization of service chain allocation with joint VNF instantiation and SDN network control," *in Proc. European Conference and Exhibition on Optical Communications (ECOC)*, September 2018.

[FRamb13] F. Rambach et al. "A multilayer cost model for metro/core networks,", IEEE/OSA Journal of Optical Communications and Networking (JOCN), vol.5, no.3, pp.210-225, 2013

[FXu17] F. Xu et al. "Understanding Mobile Traffic Patterns of Large Scale Cellular Towers in Urban Environment", IEEE/ACM Trans. Networking 25(2), April 2017

[HAzwa14] H. Azwar et al. “H.265 video delivery using dynamic adaptive streaming over HTTP (DASH) on LAN network,” in Proc. International Conference on Telecommunication Systems Services and Applications (TSSA), 2014.

[HSyed14] Syed et al. “Trade-offs in optimizing the cache deployment of CDNs,” *in Proc. IEEE IEEE Conference on Computer Communications (INFOCOM)*, 2014.

[InfineraMetro] “Cloud scale metro networks”, [online]: https://www.infinera.com/gometro/

[ITU-T-TN5G18] “Transport Network Support of IMT-2020/5G,” ITU-T GSTR-TN5G Technical Report, February 2018

[JAHer19] J. A. Hernandez et al. “Meeting the traffic requirements of residential users in the next decade with current FTTH standards: how much? how long?”, IEEE Communication Magazine (accepted), 2019

[JGome08] J. V. P. Gomes et al. "Analysis of Peer-to-Peer Traffic Using a Behavioural Method Based on Entropy," 2008 IEEE International Performance, Computing and Communications Conference, Austin, Texas, 2008, pp. 201-208.

[JMFab16] J.M. Fabrega et al. “On the filter narrowing issues in elastic optical networks,” IEEE/OSA Journal of Optical Communications and Networking (JOCN), vol.8, no.7, pp.A23-A33, 2016.

[JPedr18] J. Pedro et al. “Metro transport architectures for reliable and ubiquitous service provisioning,” *in Proc. Asia Communications and Photonics Conference (ACP)*, October 2018

[LHuan17] L. Huang et al. “Analysis of User Behavior in a Large-Scale VoD System,” in Proc. NOSSDAV, 2017.

[MBala18] M. Balanici et al. “Hybrid electro-optical intra-data center networks tailored for different traffic classes”, in IEEE J. Opt. Commun. Netw. 10(11), pp. 889-901, 2018

[McKinseyRoadto5G18] “The road to 5G: The inevitable growth of infrastructure cost”, McKinsey report, Feb. 2018, [online]: https://www.mckinsey.com/industries/telecommunications/our-insights/the-road-to-5g-the-inevitable-growth-of-infrastructure-cost

[MGarr19] M. Garrich et al., “The Net2Plan-OpenStack project: IT resource manager for metropolitan SDN/NFV networks,” *in Proc. Optical Fiber Conference (OFC)*, paper M3Z.16, March 2019.

[MRuiz16] M. Ruiz et al., "Big Data-backed Video Distribution in the Telecom Cloud," Elsevier Computer Communications, vol. 84, pp. 1-11, 2016.

[MRuiz18] M. Ruiz et al. "CURSA-SQ: A Methodology for Service-Centric Traffic Flow Analysis," IEEE/OSA Journal of Optical Communications and Networking (JOCN), vol. 10, pp. 773-784, 2018.

[NCala17] N. Calabretta et al. “Photonic integrated WDM cross-connects for disaggregated optical metro-access networks,” *in Proc. Asia Communications and Photonics Conference (ACP)*, October 2017

[Net2Plan] Net2Plan: The open-source network planner [online]: http://www.net2plan.com/

[NGMNWPLTE11] “Guidelines for LTE Backhaul Traffic Estimation”, NGMN Alliance white paper, 2011. [online]: https://www.ngmn.org/fileadmin/user_upload/NGMN_Whitepaper_Guideline_for_LTE_Backhaul_Traffic_Estimation.pdf

[NIW] Excel file template of the NIW library. [online] : https://github.com/girtel/MetroHaulPlanner/blob/master/MetroNetwork_Template.xlsx

[PVela16] P. Velan, J. Medková, T. Jirsík, P. Celeda: “Network Traffic Characterisation Using Flow-Based Statistics” in IEEE/IFIP NOMS Conference, 2016.

[RBoll08] R. Bolla et al. "Characterizing the network behavior of P2P traffic," 2008 4th International Telecommunication Networking Workshop on QoS in Multiservice IP Networks, Venice, 2008, pp. 14-19.

[RJing17] R. Jing et al. “China Telecom’s requirements on 5G transport,” *in ITU-T Workshop on the Evolution of Transport Networks to Support IMT-2020/5G,* October 2017.

[SAlba17] S. Albandea et al. "Valley times in the Spanish Academic Network", in IEEE/IFIP IM Conference, May 2017

[VLope16] Elastic Optical Networks: Architectures, Technologies and Control, ed: V. López and L. Velasco, Springer, 2016.

[WHedd12] W. Van Heddeghem et al: "Power consumption modeling in optical multilayer networks", Photonic Netw. Commun 24, pp. 86-102, 2012

List of acronyms

5G-PPP	5G Infrastructure Public Private Partnership
AMEN	Access-Metro Edge Node
BER	Bit Error Rate
BGP	Border Gateway Protocol
CAGR	Compound Annual Growth Rate
CapEx	Capital Expenditure
CAPS	Combined Amplitude and Phase Shift
CO	Central Office
CORD	Central Office Re-Architected as a Data-Center
CPRI	Common Public Radio Interface
D&W	Drop & Waste
DCI	Data-Center Interconnect
DHCP	Dynamic Host Configuration Protocol
DL	Downlink
DP-QPSK	Dual-Polarization Quadrature Phase Shift Keying
DSP	Digital Signal Processing
DuFiNet	Dual-Bus Filterless Network
E2E	End to End
EDFA	Erbium-doped Fiber Amplifier
ETSI	European Telecommunications Standards Institute
FPGA	Field-Programmable Gate Array
ILP	Integer Linear Programming
IoT	Internet of Things
IP	Internet Protocol
KPI	Key Performance Indicator
LCoS	Liquid Crystal on Silicon
LTE	Long Term Evolution
MANO	Management and Orchestration
MCEN	Metro-Core Edge Node
MCM	Multi-Carrier Modulation
MRSA	Modulation Format, Routing and Spectrum Assignment
NFVI	Network Functions Virtualization Infrastructure
NIW	NFV-over-IP-over-WDM
NR	New Radio
NRZ	Non-return-to-zero
OLT	Optical Line Termination
ONF	Open Networking Foundation
P2P	Point to Point
PIC	Photonic Integrated Circuit
PNF	Physical Network Function
PON	Passive Optical Network
PoP	Point of Presence
QoE	Quality of Experience
QoS	Quality of Service
RAT	Radio Access Technology

ROADM	Reconfigurable Optical Add/Drop Multiplexer
RRU	Remote Radio Unit
SDN	Software Defined Networking
SNR	Signal-to-Noise Ratio
SSB	Single Side Band
TIA	Trans-impedance Amplifier
UL	Uplink
UPF	User Plane Function
VLAN	Virtual Local Area Network
VNF	Virtual Network Function
VRF	Virtual Routing Functions

[end of document]