Metro-Haul: SDN Control and Orchestration of Disaggregated Optical Networks with Model-Driven Development

Ramon Casellas, Ricardo Martínez, Ricard Vilalta, Raül Muñoz
Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA), Castelldefels, Spain
e-mail: ramon.casellas@cttc.es

ABSTRACT
SDN solutions for optical transport networks are often associated to single-vendor optical domains, managed as single entities, in a deployment model that is commonly referred to as fully aggregated. Such controllers do export and expose interfaces for the limited control of abstracted resources and operations via north-bound interface (NBI) to operations/business support systems (OSS/BSS), but such APIs are often vendor specific and internal control aspects related to provisioning, monitoring and resource management remain proprietary and not disclosed.

Disaggregation of optical networks refers to a deployment model of optical systems, by composing and assembling open, available components, devices and sub-systems. This disaggregation is driven by multiple factor (the mismatch between the needs of operators and the ability to deliver adapted solutions by vendors or the increase in hardware commoditization) and disaggregated networks are an excellent use case for open and standard interfaces, showing the benefits of a unified, model-driven development.

In this paper, we address the SDN control of a disaggregated optical networks and its role in a wider Control Management and Orchestration (COM) architecture to offer ETSI NFV Network Services in a metropolitan infrastructure, characterized by multiple geographical locations and NFVI. We highlight the main use cases and the implemented extensions to the ONOS Platform for the aforementioned purpose. We detail the design of the controller (core systems and applications); the approach taken for model driven-development, including the YANG modelling language for the different optical devices and the NETCONF protocol to remotely configure the devices and, finally, the experimental validation of the approach and implementation with selected scenarios.

Keywords: SDN Control of Optical Networks, Disaggregated Optical Networks, Model-Driven Development, YANG modelling language, NETCONF protocol.

1. INTRODUCTION

Current SDN solutions for optical transport networks are often associated to single-vendor optical domains, managed as single entities, in a deployment model that is commonly referred to as fully aggregated. Although resources are abstracted by the SDN controller at the networking level and exposed via a north-bound interface (NBI) to operations/business support systems (OSS/BSS), internal resource management, monitoring, and control remain proprietary. Disaggregation of optical networks refers to a deployment model of optical systems, by composing and assembling open, available components, devices and sub-systems. This disaggregation can be partial or total (down to each of the optical components) and is driven by multiple factors, notably, the mismatch between the needs of operators and the ability to deliver adapted solutions by vendors; the increase in hardware commoditization; the different rate of innovation for different components; the promised acceleration on the deployment of services and the consequent reduction in operational and capacity expenses. Disaggregation aims at providing a new degree of flexibility, allowing component migration and upgrades without vendor lock-in, a trade-off in terms of current and potential performance, vendor support and cost are an excellent use case for open and standard interfaces, showing the benefits of a unified, model-driven development. There is a need to have better configuration management, a clear separation of configuration and operational data, while enabling high level constructs more adapted to operators' workflows supporting network-wide transactions, rollback capabilities and transactional semantics. However, optical networks are particularly challenging to model due to the lack of agreed-upon hardware models, and this is critical for the development of an interoperable ecosystem.

2. THE METRO-HAUL PROJECT

This section presents the Metro-Haul project [1], along with its data and control architecture. The goal of the project is to architect and design cost-effective, energy-efficient, agile and programmable metro networks that are scalable for heterogeneous 5G access and future requirements, address the anticipated capacity increase and its specific characteristics e.g. mobility, low latency, low jitter; and support a wide variety of services and use cases with special emphasis on services from various industries vertical to the ICT. This macroscopic goal encompasses the design of optical metro-nodes (including full compute and storage capabilities), which interface effectively with both 5G access and multi-Tbit/s elastic core networks, combine heterogeneous resources (processing, storage and networking) in variable sized-pools as well as the design of a control, orchestration and management subsystem that relies on existing SDN [2] and NFV [3] Frameworks and on unified information and data modelling across devices, infrastructures and services.

2.1 The Metro-Haul Infrastructure supporting 5G services

The Metro-Haul infrastructure (spanning multiple geographic locations) relies on macroscopic nodes, combining networking, processing and storage resources. Such modular devices are composed of different components operating at different layers and technologies, and of different vendors, realizing hardware and software disaggregation inside the node. In particular, it implements layer 0-1 (optical domain), layer 2 transmission and switching and computing capabilities, provided by local pool of servers to instantiate virtual network functions (VNFs) with configurable amount of processing, memory and storage. Concrete specializations of the generic architecture are at the Access Metro Edge nodes (AMEN nodes) to interface with heterogeneous access technologies and at the Metro Core Edge nodes (MCEN nodes). A main challenge and novelty of the system is to offer 5G-oriented services in a context characterized by having multiple, geographically dispersed sites or locations, interconnected by a multi-layer transport network, heavily relying on the optical technology to support the increasing bandwidth and latency requirements. The Metro-Haul infrastructure (see Fig.1.a) includes, notably, the following main parts: 1) The *Optical disaggregated transport network*, which provides high bandwidth, low latency connections between remote locations and constitutes the infrastructure core part. The optical network is obtained as a combination of Optical- Network Elements (O-NEs), where one or more O-NEs are physically located in each location. Current scope is ROADMs and Transponders as main O-NE. 2) *Packet switched networks*. Each location includes a Layer 2/Layer 3 packet switched network that aggregates traffic coming from access and aggregation networks and which provides connectivity to functions and applications running locally. 3) *Computing and Storage Infrastructure*. A variable number of computing, storage, and virtualization servers, attached to the packet switched networks, available at every location, part of the NFVI of the Metro-Haul network.

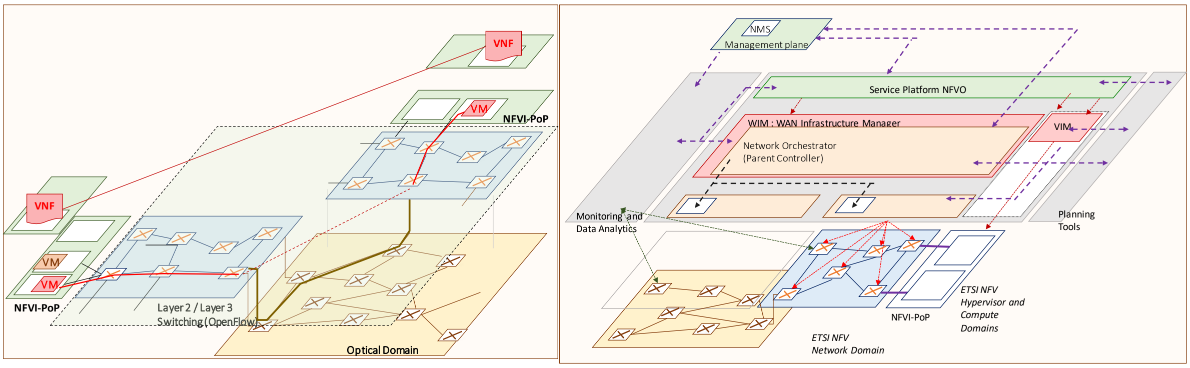

Fig. 1. (a) Metro-Haul Infrastructure, showing the main components: the disaggregated transport network interconnecting Metro-Haul locations (AMEN/MCEN nodes), the packet switched networks and the NFVI PoP encompassing computing, storage resources. (b) Macroscopic view of the Metro-Haul Control, Orchestration and Management (COM) and Service Platform, highlighting the NFV MANO part, the hierarchical SDN control, the Monitoring and data analytics (MDA) and Network Planning (NP) components

2.2 Control, Orchestration and Management (COM) System

The SDN control plane and NFV Management and Orchestration (MANO) system of Metro-Haul is referred to as the Control Orchestration and Management (COM) system (see Fig.1.b). Its architecture consists of four blocks:

1 *Network Control and Orchestration* relies on an over-arching control adopting hierarchical control architectures with a parent SDN controller abstracting the underlying complexity. We adopt open interfaces exporting programmability along with unified and systematic information and data modelling. Finally, the coordinated control of the packet layer and the optical layer is needed to provision end-to-end services and steering traffic coming from aggregation networks.

2 *Compute and Storage Integration via SDN/NFV:* Joint IT/Cloud and Network Orchestration is used to refer to the coordination of resources to deploy services and applications that require storage, computing and networking resources. The approach involves integrating the NFV MANO functional elements [4] over multiple VIMs, with hierarchical control planes which are abstracted as WAN infrastructure managers (WIM)

3 *ETSI MANO / Slicing Integration:* the ETSI NFV framework can be used as a starting point for 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).

4 *Monitoring and Data Analytics* Autonomic networking entails the capability to do measurements on the data plane and generating data records that are collected and analyzed to discover patterns (knowledge) from data. Such knowledge can be used to issue re-configuration/re-optimization recommendations toward COM modules, such as an SDN controller or orchestrator.

3. SDN CONTROL OF DISAGGREGATED NETWORKS

The core of this work is the SDN control of the optical transport network. The approach is based on having a SDN controller controlling one or mode devices, which are characterized by their data model. In general, a device *Information Model* macroscopically describes the device capabilities, in terms of operations and configurable parameters, using high level abstractions without specific details on aspects such as a particular syntax or encoding. A *Data Model* determines the structure, syntax and semantics of the data that is externally visible. YANG [5] is a data modelling language, where a model includes a header, imports and include statements, type definitions, configurations and operational data declarations as well as actions (RPC) and notifications. The language is expressive enough to structure data into data trees within the so called *datastores*, by means of encapsulation of containers and lists, and to define constrained data types (e.g. following a given textual pattern); to condition the presence of specific data to the support of optional features and to allow the refinement of models by extending and constraining existing models (by inheritance/augmentation), resulting in a hierarchy of models.

3.1 Common Yang models and configuration protocols for Optical Devices

In our work, we are considering mainly two sets of data models for abstracting optical hardware devices. OpenConfig [8], for terminal optical devices within a DWDM system and the OpenROADM multi-source agreement [7], which focuses on functional disaggregation and covers pluggable optics, transponders and ROADMs. A *protocol* offers primitives to view and manipulate the data, providing a suitable encoding as defined by the data-model. For YANG, the NETCONF protocol [6], enables remote access to a device, and provides the set of rules by which multiple clients may access and modify a *datastore* within a NETCONF server (e.g., device). It is based on the exchange of XML-encoded RPC messages over a secure connection. The layering mode relies on having configuration or notification data (Content Layer) that is exchanged between a client and a server, with a set of well-defined operations encapsulated in RPC messages or notifications (Message Layer) and using a Secure Transport. The data is arranged into one or multiple configuration datastores, a complete set of configuration information that is required to get a device from its initial state into a desired operational state.

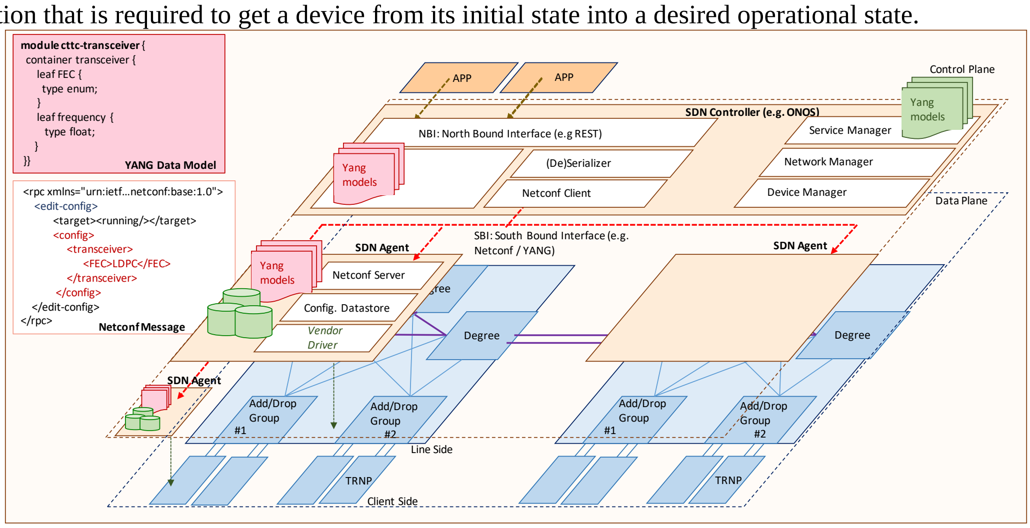

Fig.3. SDN Control of a disaggregated optical network (transceiver and/or optical line system)

3.2 Metro-Haul SDN Controller within the ONOS Framework

The Metro-Haul SDN controller is being developed within the ONOS framework (Fig.3), relying on ONOS support for model driven development, which applies both at the NBI and at the SBI levels. The use of YANG is possible to describe devices as well as services or other control and management constructs (topologies, inventory). A common design goal is to hide and encapsulate the details of the YANG modelling language from the actual application development, so application logic can be programmed using languages and constructs the developers are familiar with. Common Yang development kits have a set of basic functions: i) to be able to tokenize, parse and compile (a set of) model files that describe e.g., a device's capabilities, configuration and operational data and notifications; ii) to have a central repository to manage registered YANG models along with the capabilities to query the repository; iii) to generate common (often referred to as boilerplate) code that can be used "as is" by applications in order to parse and construct configuration objects and messages to be serialized / de-serialized and encoded (e.g. XML as the main payload of NETCONF messages). More advanced frameworks may offer validation functions and integrity checks to make sure that a given configuration is consistent. The ONOS Yang Tools subproject is one of such frameworks, although a development design is to make it useful standalone. We have used the ONOS Yang tools consistently to register within ONOS service models that are consumed by applications and device models for the actual ONOS devices.

4. EXPERIMENTAL VALIDATION

The SDN control has been validated by considering a point to point link between two endpoints relying on a set of devices emulating an Open Line System (OLS) with a transceiver, a ROADM and a receiver. From the point of view of the NBI, a service YANG model is registered into the ONOS YANG subsystem, so a RESTCONF based interface can automatically be used to trigger service provisioning. Upon request through its NBI, the SDN controller proceeds to configure the optical elements. The Tx and Rx models enable to select parameters like the central frequency, and the ROADM model is based on the OpenROADM 2.2 device model including a 2- degree, 2-SRG ROADM model. A connection is requested using TAPI 2.0 [9] connectivity service model, which is then mapped to the configuration of the devices. From the Fig.4 we can see the exchange of NETCONF messages triggered by a TAPI NBI request. The whole procedure is completed after ~40ms, but this is without taking into account hardware delay (emulated).

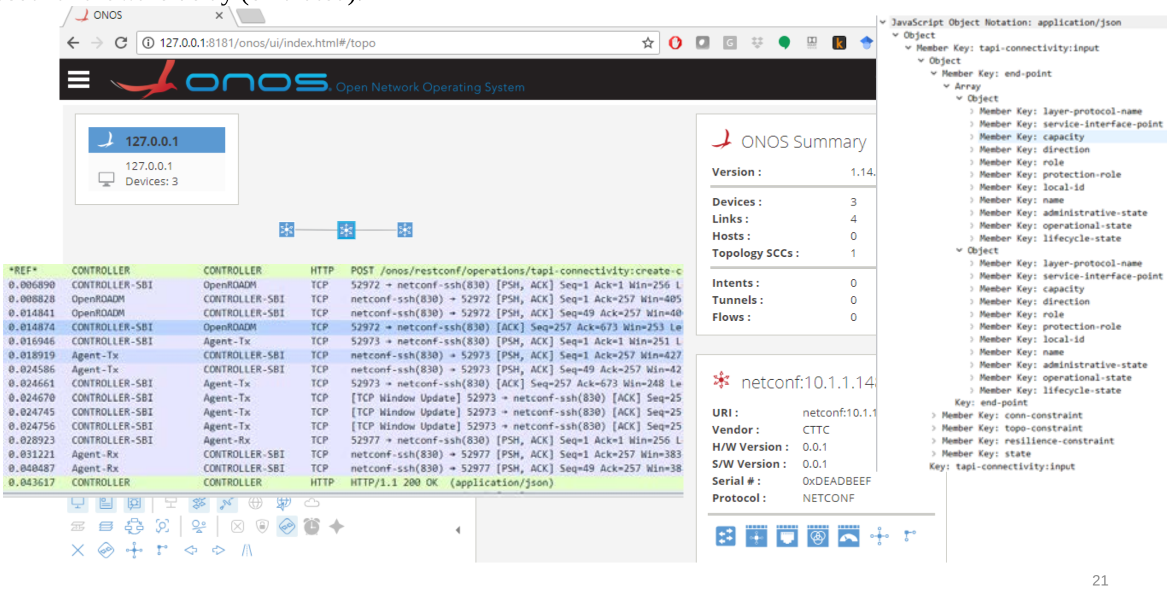

Fig.4. Experimental Validation of NETCONF/YANG device configuration after a TAPI NBI Request.

5. CONCLUSIONS

An optical transport network is a key infrastructure to support heterogeneous and 5G services with stringent constraints of bandwidth and latency. The disaggregation of such networks presents an opportunity for network operators (in terms of potential cost savings, but also piecewise upgrades, different innovation rates and maximum flexibility in composing open line systems and deploying networks) and, in particular, is a major use case for the use of open interfaces and model driven development.

ACKNOWLEDGEMENTS

The work presented in this paper has been partially funded by Spanish MINECO DESTELLO (TEC2015-69256-R) project, and project H2020-ICT2016-2 METRO-HAUL (Proposal number: 761727).

REFERENCES

[1] METRO High bandwidth, 5G Application-aware optical network, with edge storage, compUte and low Latency, the Metro-Haul project [online] https://metro-haul.eu/

[2] ONF ONF TR-521, "SDN Architecture, Issue 1.1, 2016.

[3] ETSI GS NFV 002 (V1.2.1): "Network Functions Virtualisation (NFV); Architectural Framework", December 2014.

[4] ETSI "Network Functions Virtualisation (NFV); Management and Orchestration", ETSI GS NFV-MAN 001 V1.1.1 (2014-12), available: http://www.etsi.org/deliver/etsi_gs/NFV-MAN/001_099/001/01.01.01_60/gs_NFV-MAN001v010101p.pdf

[5] M. Bjorklund, Ed., "The YANG 1.1 Data Modeling Language", IETF RFC 7950, August 2016.

[6] R. Enns, ed. "Network Configuration Protocol NETCONF", IETF RFC 6241, June 2011.

[7] The Open ROADM Multi-Source Agreement (MSA), online: http://www.openroadm.org

[8] The OpenConfig project, online: http://openconfig.net/

[9] ONF Transport API [online] https://www.opennetworking.org/wp-content/uploads/2017/08/TAPI-2.0-Updates-Overview.pdf and https://github.com/OpenNetworkingFoundation/TAPI