Autonomic NFV Network Services on Top of Disaggregated Optical Metro Networks

P. R. Esmenats[1], R. Casellas[2], Ll. Gifre[3], A. P. Vela[1], M. Ruiz[1], R. Martinez[2], and L. Velasco[1*]

(1) Universitat Politècnica de Catalunya (UPC), Barcelona, Spain
(2) Centre Tecnològic de Telecomunicacions de Catalunya (CTTC), Barcelona, Spain
(3) Universidad Autónoma de Madrid (UAM), Madrid, Spain
e-mail:lvelasco@ac.upc.edu

Abstract: Control and orchestration of metro network to support dynamic provisioning and reconfiguration of services based on Virtual Network Functions are demonstrated through a CDN example. Interaction among CASTOR MDA, OSM and ONOS is exhibited.

OCIS codes: (060.4250) Networks; (060.4256) Networks, network optimization

1. Introduction

In view of the requirements associated to 5G networks, the control, orchestration and management (COM) system of metro networks needs to support dynamic service provisioning involving IT and networking resources. A common trend is to design such service and resource orchestration systems by adopting, extending and building on top of frameworks that follow Software Defined Networking (SDN) and Network Function Virtualization (NFV) principles. While SDN encourages control plane and data plane separation bringing network programmability and enabling an application layer, NFV focuses on replacing dedicated network appliances with software implementations running on common shared hardware, becoming Virtualized Network Functions (VNF). The ETSI NFV architecture defines the NFV Infrastructure (NFVI) deployed across multiple points of presence (NFVI-PoP) supporting the instantiation of Virtual Machines (VM), along with the Management and Orchestration (MANO) subsystem dealing with the orchestration of VNFs and their deployment as components of the so-called NFV *Network Services* (NS).

Nonetheless, to achieve true *autonomic* NFV NSs, SDN control and NFV management need to be augmented with data-driven decision-making, using advanced monitoring and machine learning (ML) tools, feeding control and orchestration planes. The architecture is developed in the framework of the METRO-HAUL EU project and includes CASTOR [1], a monitoring and data analytics (MDA) framework that collates monitoring data records from network and NFVI-PoPs and contains ML algorithms to enable making ML-driven decisions to trigger actions, essentially connecting data-driven automation with policy-based orchestration and management planes.

This demonstration, based on our previous work in [2], focuses on an illustrative autonomic NFV network service use case where the application manager runs on top of the control and management plane; a virtualized Content Delivery Network (CDN) service is able to autonomously adapt to the load by requesting the instantiation of new VMs in selected leaf cache nodes, as well as by incrementing the capacity of the network connecting users with the caches. CDN autonomous decisions made using optimization techniques, are enabled by data analytics techniques applied on monitoring data collected from the MDA controller.

2. Control, Management and Orchestration System

The COM architecture (Fig. 1a) consists of four main building blocks for *i*) the control and orchestration of the network, *ii*) managing the computing and storage infrastructure, *iii*) NFV, and *iv*) monitoring and data analytics.

Since transport networks are increasingly segmented in domains, coordinated control of packet and optical layers is needed to provision end-to-end services and steering traffic coming from aggregation networks. Regarding the optical domain, partially disaggregated optical networks are considered. Here, we adopt open interfaces exporting programmability along with unified and systematic information and data modelling.

A variable number of computing, storage, and virtualization servers are available on every Central Office (CO), which constitute part of the NFVI. The set of servers in each CO are managed by a dedicated Virtual Infrastructure Manager (VIM), within a unique administrative domain.

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 with hierarchical transport network control planes, which are abstracted as WAN infrastructure managers (WIM). The ETSI framework is used as a starting point to provide NFV NSs, encompassing NS endpoints and one or more VNFs interconnected by logical links, forming VNF forwarding graphs (VNFFGs). Logical links are mapped to supporting network connectivity services spanning multiple network segments.

Finally, autonomic networking entails the capability to perform data plane measurements and generating data records that are collected and analyzed to discover patterns. In our architecture, a centralized *MDA controller*

The research leading to these results has received funding from the EC through the METRO-HAUL (G.A. nº 761727) and the Spanish MINECO TWINS (AEI/FEDER TEC2017-90097-R) projects, and the Catalan Institution for Research and Advanced Studies (ICREA).

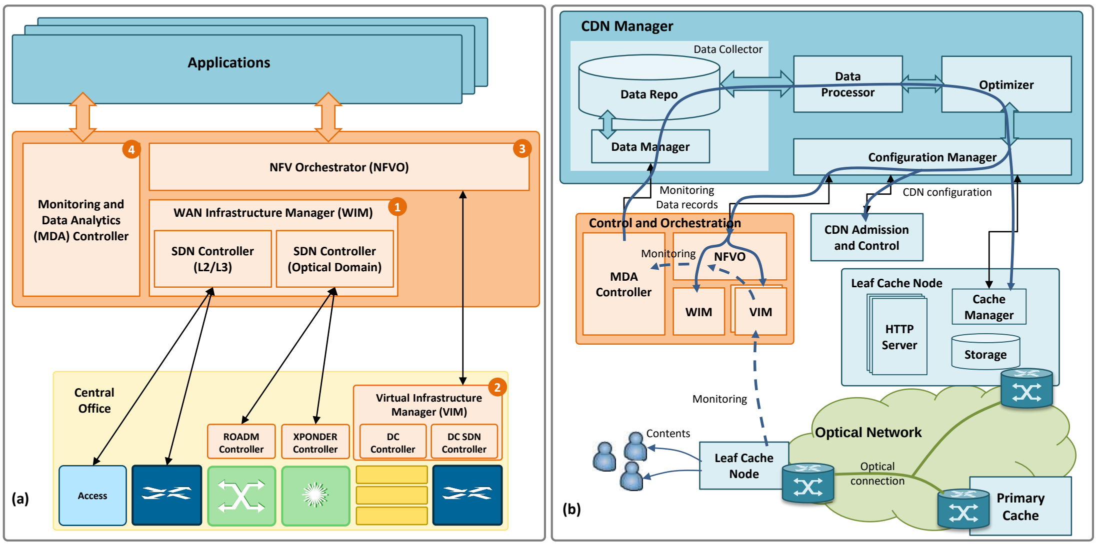

Fig. 1. Control, Orchestration and Management architecture (a) and demonstrated control loop for CDN reconfiguration (b).

that contains a big data repository and data analytics capabilities is used to collate monitoring data records from COs. Data records and notifications are stored and processed in correlation with operational databases.

3. Virtualized Content Delivery Network

To assess the applications support based on VNFFGs of the COM architecture, the demonstration targets the adaptation of a typical operator-owned CDN, where a CDN manager is in charge of the virtualized CDN function for video on demand (VoD) distribution [3]. The virtualized CDN collects monitoring data from the caches and dynamically reconfigures them to fulfil current and future demand. By virtualizing the CDN function, CDN costs can be minimized by dynamically adapting computing and network resources to current and future users' needs while ensuring the highest quality.

The CDN uses the MPEG Dynamic Adaptive Streaming over HTTP (MPEG-DASH) technique to distribute contents. MPEG-DASH enables media content delivering over HTTP protocol using standard HTTP web servers' infrastructures to users' devices; it divides contents into a sequence of small file *segments*, each containing a short interval, e.g., 2 seconds, of the content and provides mechanisms to request the segmented contents. Distribution of segments can be load balanced, among others, to mitigate disruptions in the video distribution service during the reconfigurations and spikes in demand.

Fig. 1b presents the proposed architecture and the control loops that allow dynamically adapting the CDN. A virtualized hierarchical CDN can be deployed on a metro infrastructure with some few primary caches that are the entrance point for new multimedia contents and a number of *Leaf Cache Nodes* running in COs, close to end users: a centralized *CDN Admission and Control* module implements CDN user access policies and redirects users' requests, e.g., based on their geographical location, to the leaf cache node that will serve them. Leaf cache nodes distribute VoD contents that are stored locally based on its popularity.

A virtualized leaf cache node would consist of *HTTP servers* and *Cache Managers* running as software inside VMs deployed in the same CO; the *HTTP server* component serves end users' segment requests, whereas the *Cache Manager* is the entry point of the cache node; it receives users' requests, performs appropriate security checks, identifies which contents will be locally stored, and redirects users' requests to the appropriate HTTP server. Each component usually consists in a pool of resources for load balancing and redundancy purposes. The amount of resources in every resource pool can be dynamically adapted in response to spikes in demand, e.g. during a sports event. Additionally, the Cache Manager exposes a north-bound interface (NBI) to the CDN manager to allow its remote reconfiguration.

Finally, a *CDN Manager* is responsible for adapting the CDN to the current and future load. To that end, monitoring data collected by the MDA controller is retrieved and analyzed; specifically, the logs from the cache manager that contain useful information regarding user activity and contents access together with the activity of the packagers and HTTP servers are analyzed. Specifically, contents' popularity can be computed not only at the leaf cache level, but also at the CDN manager level. This enables reassigning users among leaf caches without affecting the structure of the cached contents. The analyzed performance and load of the CDN is used to elastically adapt its resources to current and near future service needs.

A configuration manager is in charge of interfacing the NFVO to request and release IT and networking resources and of properly configuring every component.

4. Demonstration

The setup proposed for the demonstration is presented in Fig. 2. The WIM is implemented extending the ONOS controller framework. The WIM is the entry point for other entities to operate the transport network that interconnects NFVI-PoPs in the COs; it is a dual layer network based on ONOS that encompasses a packet switched, OpenFlow controlled layer and an optical layer implemented as a disaggregated optical network and using YANG/ NETCONF interfaces.

For the demonstration, a point to point link between two OpenFlow packet switches relies on a provisioned Optical Connectivity Service, involving 3 emulated NETCONF/ YANG devices: two transceivers and one ROADM; the former two use OpenConfig data models while the latter uses OpenROADM. From the point of view of the WIM NBI, a service YANG data model is registered into the ONOS YANG subsystem, so a RESTCONF based interface can automatically be used to trigger service provisioning. At this stage, the service provisioning is based on a simple YANG data model to request connectivity. Upon request through its NBI, the WIM proceeds to configure the optical elements through *edit-config* NETCONF messages.

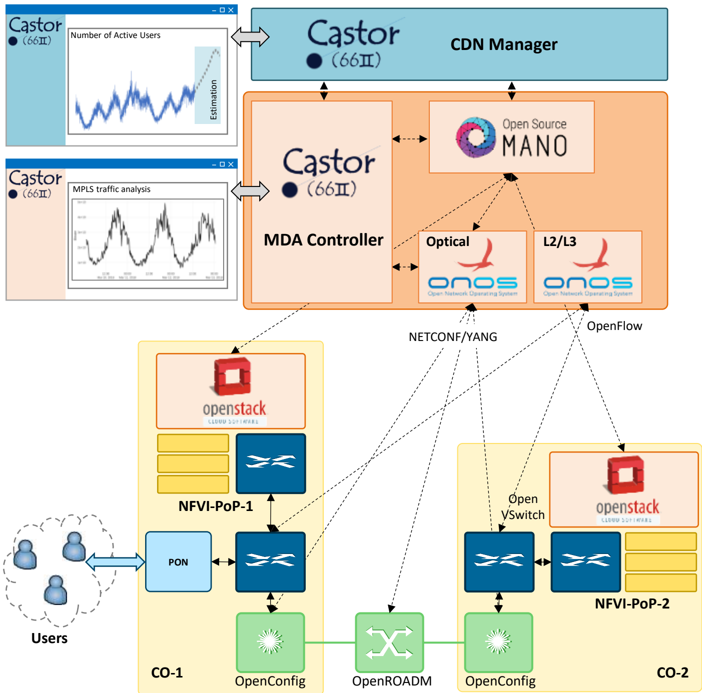

Fig. 2. Experimental set-up.

OpenStack and Open Source MANO (OSM) have been selected as VIM and as NFVO, respectively, whereas the MDA system is based on CASTOR MDA platform.

The Ceilometer module in OpenStack periodically collects measurements such as CPU, disk, and network statistics that are collected by OSM that stores them in its local database. In addition, applications' logs are exported toward a Logstash instance. The MDA controller periodically collects monitoring data from both OSM and ONOS, becoming a global repository and potential intelligence system for the whole infrastructure.

The CDN manager has been developed on a different instance of CASTOR, enabling data collection from the MDA controller through a REST API. ML algorithms are used to model CDN load evolution, which is used to periodically adapt the CDN resources to the expected near future load. An heuristic algorithm to compute near optimal resource allocation for the next period is executed using the expected near future load as input data. Finally, cache managers and CDN admission and control modules have been developed in Python.

The demonstration first shows the relation among the COM modules and from them to the network and IT infrastructure. In particular, the control of OpenConfig and OpenROADM devices, as well as the collection of monitoring data. Such data will be visualized through the MDA controller GUI. As for the CDN self-adaptation, it will be demonstrated the event of an increment/decrement in the number of users being served (see Fig. 1b). The evolution in the number of users is being monitored by the CDN manager by analyzing the logs from the HTTP servers in every leaf cache. A predictive model allows to estimate the number of users for the next period (e.g., one hour). In the event of an increment in the number of users, load balancing among the different leaf caches in a given metro area allows to serve the expected number of users. To this end, an optimization problem that uses the expected number of users from the data analytics module, as well as the current resources allocated in every leaf cache (i.e., active HTTP servers, VLANs, etc.) is solved by the local planning tool and the optimal result is used to reconfigure the CDN. Note that this autonomic NFV service operation requires from the support of the proposed COM architecture.

5. Innovation and Relevance to OFC Conference

The demonstration aims at bringing cognitive capabilities to NFV services with the support of a holistic control and orchestration platform specifically designed for metro networks. In the demonstration, integration among the different elements in COM system, as well as with a disaggregated multilayer data plane, is exhibited. To this end, extensions have been developed in ONOS and OSM frameworks, including YANG data models to control disaggregated optical nodes, and REST APIs based east-west interfaces in ONOS and OSM to connect with CASTOR MDA controller. Finally, the autonomic CDN manager connects to OSM and the MDA controller through their exposed NBIs.

The demonstration exhibits timely innovations for autonomic NFV services supported by a holistic control, management and orchestration architecture, relevant for future 5G networks. Hence, the demonstration will attract the community's attention and it will be of interest to a broad OFC audience.

References

[1] L. Velasco *et al*., "An Architecture to Support Autonomic Slice Networking [Invited]," IEEE/OSA JLT, 2018.
[2] R. Casellas *et al*., "Metro-Haul: Supporting Autonomic NFV Services over Disaggregated Optical Networks," EuCNC, 2018.
[3] M. Ruiz *et al*., "Big Data-backed Video Distribution in the Telecom Cloud," Elsevier Comp. Comm., 2016.