

Metro-Haul: Supporting Autonomic NFV Services over Disaggregated Optical Networks

R. Casellas¹, S. Llana², Ll. Gifre³, R. Vilalta¹, J-L. Izquierdo-Zaragoza², R. Muñoz¹, M. Ruiz², R. Martínez¹, and L. Velasco^{2*}

¹ CTTC/CERCA, Castelldefels, Spain. ² UPC, Barcelona, Spain. ³ UAM, Madrid, Spain.

Abstract—The Metro-Haul Control, Orchestration and Management (COM) architecture is presented from a high-level point of view; A key, novel component of the COM architecture is the Monitoring and Data Analytics (MDA) system, able to collect monitoring data from the network, datacenters and applications which outputs can be used to proactively reconfigure resources thus adapting to future conditions, like load or degradations. To illustrate the architecture, a use case of a Content Delivery Network (CDN) network service taking advantage of the MDA is experimentally demonstrated.

Index Terms—Control, Orchestration and Management. Disaggregated optical networks. Autonomic networking.

I. INTRODUCTION

THE overall objective of the Metro-Haul project [1] is to architect and design cost-effective, energy-efficient, agile and programmable metro networks that are scalable for 5G access, encompassing the design of all-optical metro nodes (including full compute and storage capabilities), which interface with both 5G access and multi-Tb/s elastic core networks. A key component is the control, orchestration and management (COM) system.

A. Metro-Haul Infrastructure

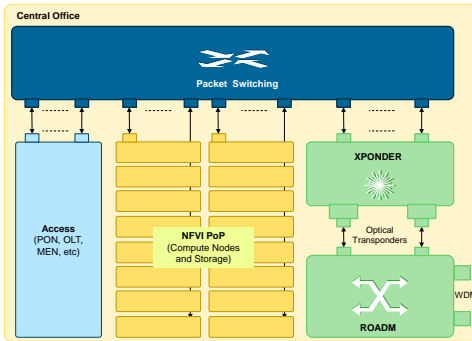


Fig. 1. Metro-Haul Central Office Infrastructure

The infrastructure includes a set of Central Offices (COs) and their connectivity (Fig. 1) with: 1) The *Optical disaggregated transport network*, which provides high bandwidth, low latency connections; 2) *Computing and Storage Infrastructure*, available at each CO; 3) A Layer 2/3 packet switched network at each CO that aggregates traffic and provides connectivity to DC; 4) Active Monitoring probes at 10 and 100 Gb/s.

B. Control, Management and Orchestration System

The COM architecture consists of building blocks: 1) Network Control and Orchestration, 2) Compute and Storage Integration via SDN/NFV, integrating the NFV MANO functional elements (with one or more orchestrators or NFVO) with hierarchical control planes which are abstracted as WAN infrastructure managers (WIM), as shown in Fig.2, 3) MANO / Slicing Integration and 4) Monitoring and Data Analytics.

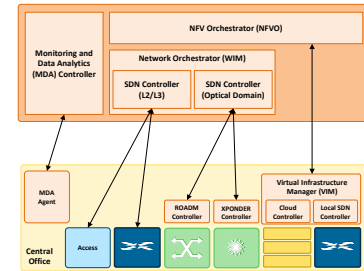


Fig. 2. Metro-Haul Control Orchestration and Management

II. USE CASE: CONTENT DELIVERY NETWORK (CDN)

A virtualized hierarchical CDN can be deployed on the Metro-Haul 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 and live-TV: VoD contents are stored in leaf caches based on its popularity [2] and live-TV content is locally prepared in leaf cache nodes [3]. A virtualized leaf cache node would consist of the following components running as software inside VMs deployed in the same CO. A *CDN Manager* is responsible for adapting the CDN to the current and future load. To that end, monitoring data is collected from the MDA controller and analyzed. Fig. 3 presents the proposed architecture and the control loops that allow to dynamically adapt the CDN. WF1 is related to CDN self-adaptation scaling up/down in the event of an increment/decrement in the amount of contents being served from a given leaf cache node, whereas WF2 focuses on pushing new contents from primary caches toward the leaf cache nodes and it is supported by the dynamic set-up of optical connections. Fig. 4 details the proposed workflows.

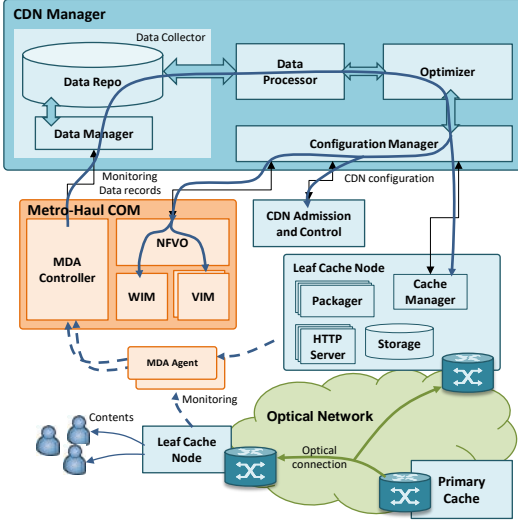


Fig. 3. Autonomous network service management

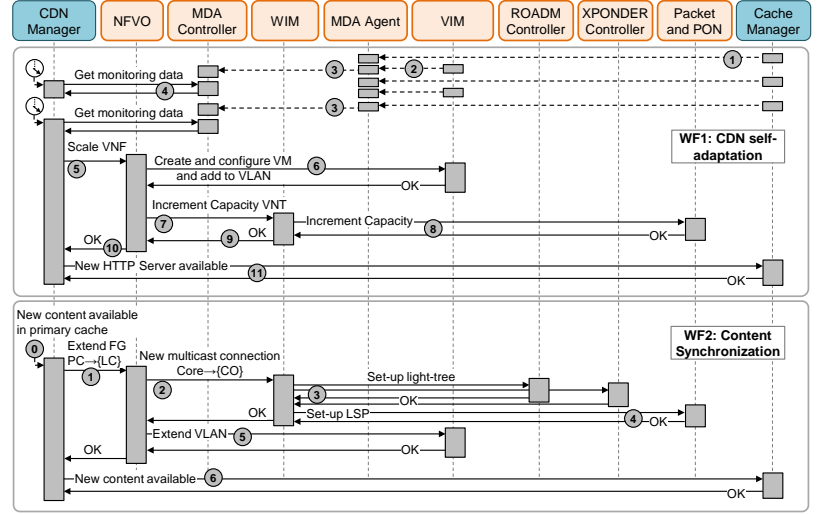


Fig. 4. CDN-driven workflows supported by the Metro-Haul COM system.

A. Experimental Demonstration

We demonstrate workflow WF1, carried on a distributed test-bed. The WIM is implemented extending the ONOS controller framework. A point to point link between two OpenFlow packet switches relies on an optical connection supported over 3 devices emulating an Open Line System (OLS) with a transceiver, a WSS and a receiver. A service YANG model is registered in the ONOS YANG subsystem. The MDA system is based on UPC's CASTOR platform [4].

SrcIP	DstIP	Protocol	Info
CacheManager	MDAAgent	HTTP	POST /receiver/logs HTTP/1.1
MDAAgent	CacheManager	HTTP	HTTP/1.1 200 OK (application/json)
VIM	MDAAgent	HTTP	POST /receiver/samples?batch=False HTTP/1.1
MDAAgent	VIM	HTTP	HTTP/1.1 200 OK (application/json)
MDAAgent	MDAController	CFLOW	IPFIX flow (668 bytes) Obs-Domain-ID= 0
CDNManager	MDAController	HTTP	GET /getMonitoring HTTP/1.1
MDAController	CDNManager	HTTP	HTTP/1.1 200 OK (application/json)
CDNManager	NFVO	HTTP	POST /scaleVNF HTTP/1.1
NFVO	VIM	HTTP	POST /identity/v3/auth/tokens HTTP/1.1
VIM	NFVO	HTTP	HTTP/1.1 200 OK (application/json)
NFVO	VIM	HTTP	POST /compute/v2.1/servers HTTP/1.1
VIM	NFVO	HTTP	HTTP/1.1 200 OK (application/json)
NFVO	WIM	HTTP	POST /onos/restconf/data HTTP/1.1
WIM	TxTPCtrl	NetConf	NetConf
TxTPCtrl	WIM	NetConf	NetConf
WIM	RxTPCtrl	NetConf	NetConf
RxTPCtrl	WIM	NetConf	NetConf
WIM	NFVO	HTTP	HTTP/1.1 200 OK (application/json)
NFVO	CDNManager	HTTP	HTTP/1.1 200 OK (application/json)
CDNManager	CacheManager	HTTP	POST /publishHTTPServer HTTP/1.1
CacheManager	CDNManager	HTTP	HTTP/1.1 200 OK (application/json)

Fig. 5. Message list for CDN self-adaptation (WF1).

Fig. 5 presents the capture of messages exchanged for WF1; the same message numbering as in Fig. 4 has been used for the sake of clarity. A REST API has been used for the Cache Manager and the VIM to send monitoring data periodically to the local MDA agent (messages 1 and 2), whereas IPFIX [5] is used between MDA agents and the MDA controller (3). The CDN manager uses a REST API to retrieve monitoring data for a set of OPs from the MDA controller (4) and uses such measurements to decide scaling leaf caches and connectivity (5). In the demonstration, one HTTP server is added to one of the leaf caches and the connectivity capacity between users

and that cache is increased (6-9). Fig.4 shows the details of the Tx device after the handshake and other initial exchanges (port discovery) have been completed. Finally, once the VNF has been scaled, the new HTTP servers are announced to the Cache Manager (11) by means of an ad-hoc REST API, so the server enters in the round-robin selection as soon as the VM starts. The scale VNF request (messages 5-10) was completed in less than 1.5 seconds, where the REST API POST VM creation (messages 6) took 0.43 seconds (plus VM stat-up time) and the RESTCONF POST operation for link provisioning (messages 8) was completed in 0.89 seconds (optical hardware configuration delay was not considered).

III. CONCLUDING REMARKS

The Metro-Haul COM Architecture has been presented, a use case has been used to illustrate how the proposed architecture supports self-adaptive virtualized services; specifically a CDN. The experimental demonstration has been carried out on a distributed test-bed connecting CTTC and UPC premises¹.

REFERENCES

- [1] METRO High bandwidth, 5G Application-aware optical network, with edge storage, compute and low Latency (Metro-Haul). [On-line] <https://metro-haul.eu>.
- [2] D. de Vleeschauwer and D. C. Robinson, "Optimum caching strategies for a telco CDN," Bell Labs Technical Journal, vol. 16, 2011.
- [3] A. Asensio *et al.*, "Scalability of Telecom Cloud Architectures for Live-TV Distribution," in Proc. OFC 2015.
- [4] L. Velasco *et al.*, "An Architecture to Support Autonomic Slice Networking [Invited]," IEEE/OSA Journal of Lightwave Technology (JLT), 2018.
- [5] B. Claise *et al.*, "Specification of the IP Flow Information Export (IPFIX) Protocol for the Exchange of Flow Information," IETF RFC 7011, 2013.

¹ This research has received funding from the European Community's through the METRO-HAUL project (G.A. n° 761727), from the AEI/FEDER TWINS project (TEC2017-90097-R), from the MINECO project DESTELLO

(TEC2015-69256-R) and from the Catalan Institution for Research and Advanced Studies (ICREA).

