5G Management and Orchestration – From Cloud-Native to 5G-Ready Applications

Roberto Bruschi, Franco Davoli, Jane Frances Pajo

Abstract In the complex scenario of highly virtualized "softwarized" networking that has been progressively unveiled by the emergence of 5G in its enhanced integration of wired and wireless resources, network management and control strategies have become increasingly relevant. Their role entails orchestration of the lifecycle of all needed functionalities, supervision and optimization in the allocation of resources, satisfaction of Key Performance Indicators and constraints, creation and lifecycle maintenance of network slices, among others, all under the dynamic evolution of user-generated traffic, multiple tenants, service and infrastructure providers. This section tries to highlight the evolution of the orchestration functionalities from the world of cloud-native applications to the Telecom Operators' (telcos) world and their mutual interaction through the mediation of the telcos' Operations Support Systems.

1 Introduction

Modern cloud technologies and architectures are largely recognized as the foundations of the upcoming 5G ecosystem [1, 2]. These technologies are expected to not only provide the needed means to allow the "softwarization" revolution in telecom-

Roberto Bruschi
CNIT National Laboratory of Smart and Secure Networks (S2N), 16145 Genoa, Italy e-mail: roberto.bruschi@cnit.it
DITEN - University of Genoa, 16145 Genoa, Italy

Franco Davoli
CNIT National Laboratory of Smart and Secure Networks (S2N), 16145 Genoa, Italy e-mail: franco.davoli@cnit.it
DITEN - University of Genoa, 16145 Genoa, Italy

Jane Frances Pajo
CNIT National Laboratory of Smart and Secure Networks (S2N), 16145 Genoa, Italy e-mail: jane.pajo@tnt-lab.unige.it
DITEN - University of Genoa, 16145 Genoa, Italy

munication infrastructures – mainly through the Network Functions Virtualization (NFV) framework [3] – but also to act as key enablers for new (more pervasive and more network-integrated) computing paradigms like, for instance, Fog and Mobile Edge Computing (MEC) [4, 5, 6].

One of the most significant added-value aspects that contribute to this ever-increasing success can be found in the cloud architecture itself, which allows a clear and effective splitting of roles among the main actors involved. This splitting, along with the disruptive "as-a-Service" paradigm and the agility of modern computing/networking virtualization technologies, has contributed to make the cloud enter the mass-market and to become an irrefutable means for any vertical market.

The anatomy of today's cloud ecosystems can be organized along three main layers [7], namely *Infrastructure-as-a-Service*, *Platform-as-a-Service*, and *Software-as-a-Service*, which clearly define the type and the boundaries of the offered service per actor. Actors offering IaaS services are providing their computing and/or networking infrastructures to third-party platform or software providers, usually referred to as *tenants*, through Virtual Infrastructure Managers (VIM), also referred to as Cloud Management Software (CMS) [8].

Through such VIM interfaces, tenants are allowed to monitor and to manage the entire lifecycle of their applications and services. In detail, the VIM manages the computing, storage, and network physical infrastructure in a datacentre, and it serves as a sort of conduit for control-path interaction between multiple virtualized (isolated) infrastructures, each one associated to a tenant, and the physical level. Broadly speaking, the VIM provides tenants with inventories, provisioning and de-provisioning operations, and the management of virtual compute, storage and networking resources, while also communicating with the underlying physical resources (e.g., hypervisors, network switches, etc.). The VIM is also responsible for operational aspects such as logs, metrics, alerts, etc.

Given the rising complexity of such services and the challenging performance requirements associated with them, a large part of these operations is often delegated and automated by a "Service Orchestrator" [9], which constitutes along with VIMs the backbone of any advanced and modern cloud system. Furthermore, it can be noted that also the definition of the ETSI NFV Working Group is perfectly compliant with the Orchestration/VIM layering infrastructure. Any actor (PaaS or SaaS providers) playing on top of the VIM layer shall have its own Orchestrator and, in case of PaaS providers, multiple Orchestrator modules might act in cascade.

On the other hand, although the NFV framework can be considered as an application of standard cloud computing technologies, as demonstrated in detail by its own designers [10], it is a common opinion that the rising of 5G technologies will significantly affect the cloud evolution. In this respect, Fog and MEC are two clear preliminary signs of this trend [6, 11]. Moreover, since the programmable resources will be an integral part of the 5G softwarized infrastructure, datacentres supporting 5G functions and vertical applications are supposed to be owned and maintained by telecom infrastructure providers, and to offer "private" (and in some case "hybrid") services [10]. Thus, public cloud platforms (like AWS, GoGrid, Microsoft Azure, Google Compute Engine, IBM Smart Cloud, etc.) are expected to have a limited relevance to the 5G/NFV ecosystem, while, during the latest few years, the NFV community selected OpenStack [12] as the reference VIM for 5G/NFV environments.

In this complex integrated scenario, it has become of paramount importance to clearly separate orchestration concerns between the application world, which is closer to the cloud-native approach, and the telco NFV platforms. The latter, though stemming from a similar paradigm, are oriented to providing network services to the vertical applications, in order to offer flexible and dynamic resource allocation that should be tailored to the applications' needs, while at the same time allowing Network Service Providers (NSPs) and Infrastructure Providers (InPs) to play their respective roles in full autonomy to compete in the telco marketplace to offer the best possible services and to try to maximize their revenue by efficient use of resources. In this framework, the concept of slicing has emerged as a powerful architectural tool [13, 14, 15].

The aim of this section is to review some of the most relevant platforms in the two environments, their separation and their mutual interaction. In doing so, we will also touch upon the philosophy and approach developed within the framework of the H2020 5G-PPP European Project MATILDA (*A Holistic, Innovative Framework for Design, Development and Orchestration of 5G-ready Applications and Network Services over Sliced Programmable Infrastructure*) [16, 17]. The section is organized as follows: we describe the Cloud/Application Orchestration in Subsection 2 and the NFV Orchestration in Subsection 3, respectively; Subsection 4 summarizes the vision, architectural components and approach of the MATILDA framework, where the Operations Support System (OSS) provides the interaction between the two orchestrators (along with the MEC Orchestrator) through the concept of slice intent; Subsection 5 contains the conclusions.

2 Cloud and Vertical Application Orchestration

Among open source VIM implementations, the most relevant ones are Eucalyptus [18], OpenNebula [19], CloudStack [20], and the well-known OpenStack [12, 21, 22]. As already mentioned, a large part of operations is often delegated and automated by a Service Orchestrator, which generally includes different mechanisms and sophisticated algorithms to cope with the following main functionalities:

- automatically instantiating services and their components (in terms of execution environments – e.g., virtual machines);
- monitoring the service and any of its components;
- providing automatic service upgrade procedures;
- acquiring or releasing compute/network/storage resources from VIM(s); and
- scaling the service to meet service level agreements and incoming workloads.

Relevant examples of well-known service orchestration platforms and related tools include Cloudify [23], Canonical's Juju [24], OpenStack Heat [25], etc.

In the NFV scenario, a number of projects and platforms, often competing among themselves in a still complex and very fragmented landscape, have been arising during the latest years both to cope with NFV specifications and to enable further opportunities for telco providers [26, 27]. However, many of these projects can be seen as added-value extensions/evolutions of state-of-the-art cloud tools and means. In some cases (e.g., Cloudify), some cloud-native projects also released specific telecom editions to meet the challenging requirements of the 5G/NFV framework.

As noted in a Linux Foundation whitepaper [28], the aforementioned complex landscape in telecommunications and networking has created the need for an umbrella architecture that harmonizes the multitude of standards and open source projects. Harmonization encompasses a number of aspects that affect both standardization and open source: i) ease of integration through well-defined information models, APIs, and interfaces; ii) common development environment to ease the integration and testing of components in a highly automated manner; and, iii) close coordination among the activities, to align functional requirements, schedules, etc., based on use cases.

The Linux Foundation, maintaining numerous open source projects relevant to 5G, has forged a unified Open Architecture for Networking and Orchestration (OS-N&O), to position and harmonize the many OS-N&O projects and standards (see Fig. 1, taken from [29]). In detail, the unified OS-N&O relies on three layers, which correspond to the reference architectures defining high-level functionality. As introduced in more detail in the remainder of this sections, these layers roughly correspond to the VIM, the NFV service and the vertical application orchestration, respectively.

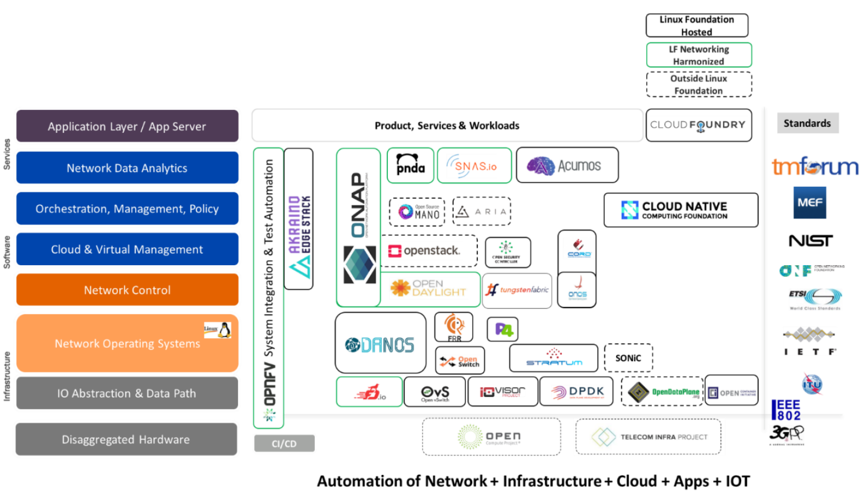

Fig. 1 Open Source Networking Communities, Projects and Initiatives [29].

Orchestrators are the component in the cloud ecosystem where the rising of new projects and products — often with diverse visions and technological basis — largely concentrated. In the public cloud, almost every cloud provider offers proprietary service orchestration tools to their customers, while other platforms provide multi-cloud compatibility. While referring to [9] for the most interesting and market-prominent orchestrators for cloud applications, we cite here explicitly the ARCADIA H2020 Project Orchestrator [30], as it has been the basis for the further 5G Vertical Applications Orchestrator (VAO) developed within the MATILDA project. The ARCADIA Orchestrator and Smart Controller supports the dynamic setup and management of highly distributed applications over a programmable infrastructure. Any type of application/service can be denoted in the form of a service graph, including the parts related with the setup and support of network functionalities. A set of functionalities, like horizontal scaling capabilities, multi-IaaS network connectivity establishment, firewall setup and operation, are supported and activated — if requested – based on the requirements imposed on behalf of a services provider or even the application developer.

3 NFV Orchestration

Owing to the heavy coupling of the current 5G ecosystem design with NFV concepts and paradigms, NFV orchestration platforms are foreseen to play crucial roles in the 5G landscape with respect to the maintenance and lifecycle management of mobile access and core network functionalities. Akin to the "war" among Software-Defined Networking (SDN) controllers (e.g., ONOS, OpenDayLight, FloodLight, etc.), the new generation of NFV-related open source projects are competing under the umbrella of organizations/consortia like ETSI and the Linux Foundation. Although ETSI has made efforts to integrate these open source projects into its community and promote a concerted effort on standards, they can clearly take on a life of their own. A brief overview on the most interesting and well-known platforms is provided in the following.

Open Source MANO (OSM) [31] is an ETSI-hosted project to develop a production quality open source NFV Management and Orchestration (MANO) software stack aligned with ETSI NFV and released under Apache 2 License. OSM was seeded by the OpenMANO efforts from Telefónica, and in combination with RIFT.io's orchestration and Canonical's Juju acting as a Virtual Network Function Manager (VNFM), provides for a more complete and standard-compliant MANO solution. OSM also includes OpenVIM [32]. OSM has now garnered about 60-member organizations worldwide that are actively working on expanding the capabilities of the platform.

The key aspects behind OSM are the usage of a strong data model based on the YANG language and the flexible architecture resulting from the usage of a plugin system. By leveraging the high-level API creation to automatic tools that parse the YANG models, the development of the OSM can be accelerated. Moreover,

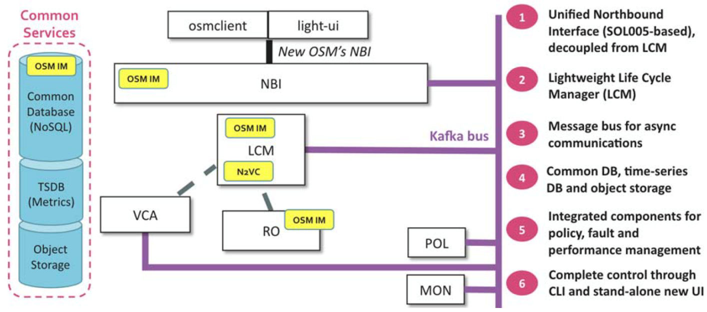

Fig. 2 OSM Release 5 architecture. (IM: Information Model; MON: Monitoring Module; N_2VC: NS-to-VNF Communication; POL: Policy Module; RO: Resource Orchestrator; TSDB: Time Series Database; VCA: VNF Configuration and Abstraction)

consistency can be assured in API calls using automatic validation. By using a plugin system, the Resource Orchestrator can be easily extended to support a wide range of commercially available SDN-enabled networks and VIMs. Additionally, OSM is designed to support brown field developments and interoperability, easing the adoption of OSM in pre-existent infrastructure.

Currently, the focus of the OSM development is providing a production-ready orchestration mechanism that supports service assurance, security (via authentication and rule based authorization), full scalability of VNFs and NSs, nested NSs and service chaining, as well as 5G network slicing.

Born by the integration between the Open-O and ECOMP projects, Linux Foundation's Open Network Automation Platform (ONAP) [33] is a comprehensive platform for real-time, policy-driven orchestration and automation of physical and virtual network functions that will enable software, network, IT and cloud providers and developers to rapidly create new services. ONAP consists of a number of software subsystems, which are part of two major architectural frameworks: i) a design-time environment to design, define and program the platform; and, ii) a run-time environment to execute the logic programmed in the design phase.

Tacker [34] is an official OpenStack project building a Generic Virtual Network Function Manager (VNFM) and an NFV Orchestrator (NFVO) to deploy and operate NSs and VNFs on an NFV infrastructure platform like OpenStack. It is based on the ETSI MANO Architectural Framework and provides a functional stack to orchestrate NSs end-to-end using VNFs. Tacker uses TOSCA (Topology and Orchestration Specification for Cloud Applications) [35] for VNF meta-data definition – specifically, the NFV profile schema. It does not internally support Service Function Chaining, but it rather exposes north-bound APIs for this purpose.

Fraunhofer's Open Baton [36] is an extensible and customizable framework capable of orchestrating NSs across heterogeneous NFV Infrastructures. It includes : i) an NFVO completely designed and implemented following the ETSI MANO specification; ii) a Generic VNFM and a Generic Element Management System (EMS) able to manage the lifecycle of VNFs based on their descriptors; iii) a Juju VNFM

Adapter in order to deploy Juju Charms or Open Baton VNF Packages using the Juju VNFM; iv) a driver mechanism supporting different types of VIMs without having to re-write anything in the orchestration logic; v) a powerful event engine based on a pub/sub mechanism for the dispatching of the lifecycle events execution; vi) an Autoscaling Engine which can be used for automatic runtime management of the scaling operation operations of VNFs; vii) a Fault Management System; viii) a network slicing engine; ix) a monitoring plugin; x) a Marketplace useful for downloading VNFs compatible with the Open Baton NFVO and VNFMs; and, xi) a set of libraries (in Java, Go, and Python).

Cloudify [23] also provides integrated end-to-end NFVM with full lifecycle automation, management and control of the lifecycle of disparate VNFs, by bringing them up or down as needed, or moving VNFs to different locations. The Cloudify platform is capable of automating deployment, driving configuration while managing auto scaling, software updates, and upgrades. It is TOSCA-based, ONAP-aligned and interfaces with other key middleware and tools. It also provides open edge orchestration and automation solution. [1]

NTT's Gohan [37] is a general-purpose API Gateway Server that enables the creation of RESTful services by orchestrating microservices. Due to its flexible architecture, Gohan can be applied to several use cases, including MANO for both VIM and legacy network devices. Similarly to OSM, modelling plays an important role in Gohan, with the information model being defined in rich YAML schemas.

The T-NOVA project [38] has designed and implemented a management/orchestration platform named TeNOR for the automated provision, configuration, monitoring and optimization of Network Functions-as-a-Service (NFaaS) over virtualised Network/IT infrastructures. In other words, T-NOVA combines IT/cloud virtualisation and Network-as-a-Service concepts to offer a complete end-to-end Cloud Network service. For TeNOR, a microservice based architecture was selected, to ensure a lean and modular implementation and operation of the system. Microservices are organized in two groups: one dedicated to NSs, which provides services to the upper layers and requests services from the second group, which is dedicated to VNFs related operations.

4 Bridging Cloud-native Applications and 5G Networking Capabilities – The MATILDA Vision and Approach

MATILDA, an H2020 5G-PPP Innovation Action (IA) coordinated by CNIT and comprising 18 partners from 10 European countries, aims to fill the integration gap between the digital systems that enable enhanced cloud-native services and the network layer, by providing the tools to foster and speed up the extension/evolution of the cloud paradigm into the 5G ecosystem, intrinsically bridging the vertical application and the NS domains. To do so, it adopts a novel and holistic approach

[1] cloudify.co/solutions

for tackling the overall lifecycle of applications' design, development, deployment and orchestration in a 5G environment.

A set of novel concepts is introduced, including the design and development of 5G-ready applications – based on cloud-native/microservice development principles – the separation of concerns among the orchestration of the developed applications and the required NSs that support them, as well as the specification and management of network slices that are application-aware and can lead to optimal application execution.

While the instantiation and management of the application-aware network slice (including the necessary set of network functions) is realised by the Network and Computing Slice Deployment Platform (NCSDP, in charge of the telecommunications infrastructure provider), the deployment and runtime management of an application is realised by the MATILDA VAO (in charge of the service provider), following a service-mesh-oriented approach. The NCSDP includes an OSS/BSS system, a NFVO and a resource manager handling the set of deployed Wide Area and Virtual Infrastructure Managers (WIMs and VIMs). Based on the interpretation of the provided slice intent, the required network management mechanisms are activated and dynamically orchestrated.

This service and functional separation concept has been clearly reflected into the final architecture, whose main elements, shown in Fig. 3, have been completely developed. An original solution has been adopted for the integration of the 5G vApps into the 5G ecosystem at the VIM-level, by keeping — also in the perspective of the Mobile Edge Orchestrator (MEO) — the tenant spaces of each vApp and NFV/Mobile Edge services in each datacentre, so that each orchestrator has its own isolated resources, quotas, external networks, etc., as illustrated in Fig. 4 [39].

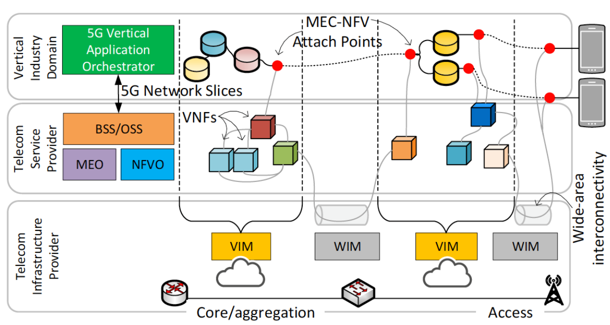

Fig. 3 Example deployment of a vApp into a 5G infrastructure, main involved stakeholders and related architectural key building blocks. Also shown is the deployment of vApp components and VNFs into multiple VIMs, and their attachment to realize the interconnectivity among VIMs and towards UEs in the mobile network.

The MATILDA Telecom Platform has been realized through the composition of the five main architectural building blocks represented in Fig. 5 [40]:

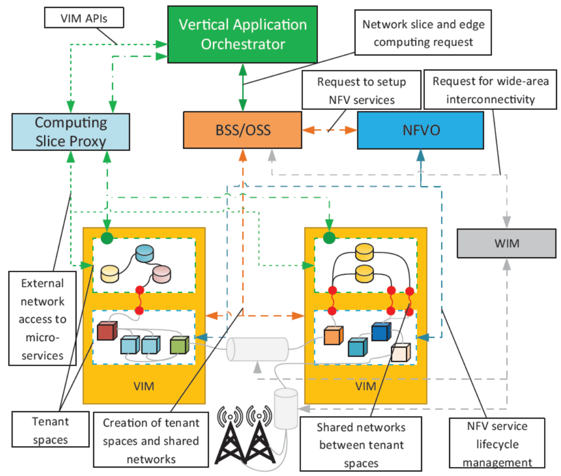

Fig. 4 VIM-level integration [39].

- The **Operations Support System** (OSS), in charge of receiving the slice intents from vertical applications (and more precisely from the VAO), and of coordinating the work of all the other building blocks in the Telecom layer platform to set up and to properly configure base 4/5G NSs, network slices, and edge computing resources. The OSS also acts as the main configuration/interfacing point for Telecom Platform Providers.
- The **NFV Orchestrator** (NFVO), in charge of managing the lifecycle of the NSs composing the base 4/5G services, and of the ones provided to slices in a shared or isolated fashion. The NFVO is also in charge of Day-2 operations for Physical Network Functions – PNFs (e.g., g/eNodeBs). In the MATILDA demonstration pilot, OSM has been selected as reference NFVO.
- The **Wide-area Infrastructure Manager** (WIM), devoted to (i) manage and monitor the wide-area communication resources, (ii) create network overlays to be used in a shared or isolated fashion by vertical applications and base telecommunication services, as well as (iii) provide information on which resources (e.g., VIMs, PNFs, etc.) can be selected in the distributed 5G infrastructure to create slices/services in order to satisfy vertical application performance requirements (e.g., end-to-end latency, bandwidth, etc.). On its southbound, the WIM can be interconnected to networking devices either in a direct fashion or through SDN controllers. In the MATILDA demonstration pilot, the WIM has been realized through an extended version of the Ericsson Network Manager [41] (forked from the official product mainstream in order to support a number of MATILDA solutions).
- The **Virtual Infrastructure Manager** (VIM – one instance per each distributed computing facility) is mainly devoted to abstract and expose computing, storage, and networking capabilities of datacentres within the 5G infrastructures. It has the

key role of isolating the various tenant domains (i.e., NFV domains and Vertical Applications' ones), as well as of creating shared resources to properly "attach" these domains [39]. In the MATILDA demonstration pilot, the VIMs have been realized through the OpenStack software suite (mainly with the Rocky and Queen official releases).

- The **Wide-Area SDN Controller** (WSC), in charge of interconnecting the control agents of the SDN devices in the wide-area network for monitoring and configuration purposes. It exposes a northbound interface mainly towards the WIM and Telecom layer monitoring frameworks. Within the MATILDA WP4 pilot, the OpenDaylight project [42] has been chosen as reference WSC.

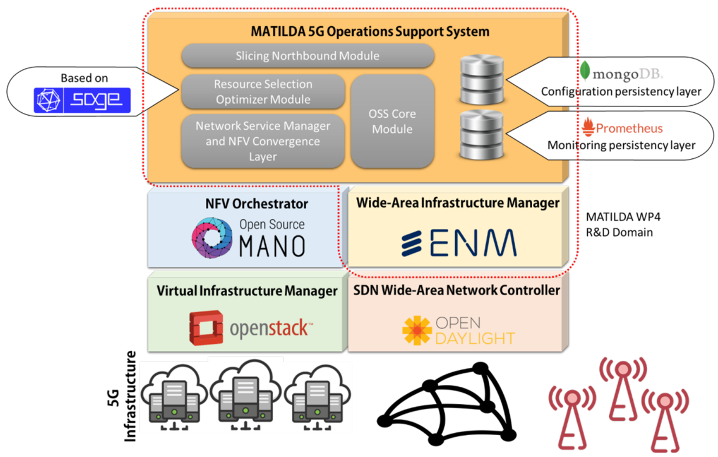

Fig. 5 Main building blocks composing the MATILDA 5G Telecom and Infrastructure Platforms [40]

.

At the same time, the design of a set of intelligent vertical application orchestration mechanisms, in order to realize the proper placement and orchestration of 5Gready applications over the created application-aware network slices, has been completed and developed. The main components constituting the MATILDA VAO are: (i) the deployment and execution manager that supports the production of optimal deployment plans, as well as the management of the overall execution of the application, (ii) a set of data monitoring mechanisms which collect feeds from network and application-level metrics, (iii) a data fusion, real-time profiling and analytics toolkit, which produces advanced insights through machine learning mechanisms and provides real-time profiling of the deployed components, application graphs and VNFs, (iv) service discovery mechanisms for supporting registration and consumption of application-oriented services following a service mesh approach, (v) a context awareness engine providing inference over the acquired data and support of runtime policies' enforcement, and (vi) mechanisms supporting interaction among the VAO and the 5G programmable infrastructure management tools.

A summary of the overall lifecycle of an application created with the MATILDA framework is represented in Fig. 6, highlighting the interaction among the different stakeholders and the usage of metamodels.

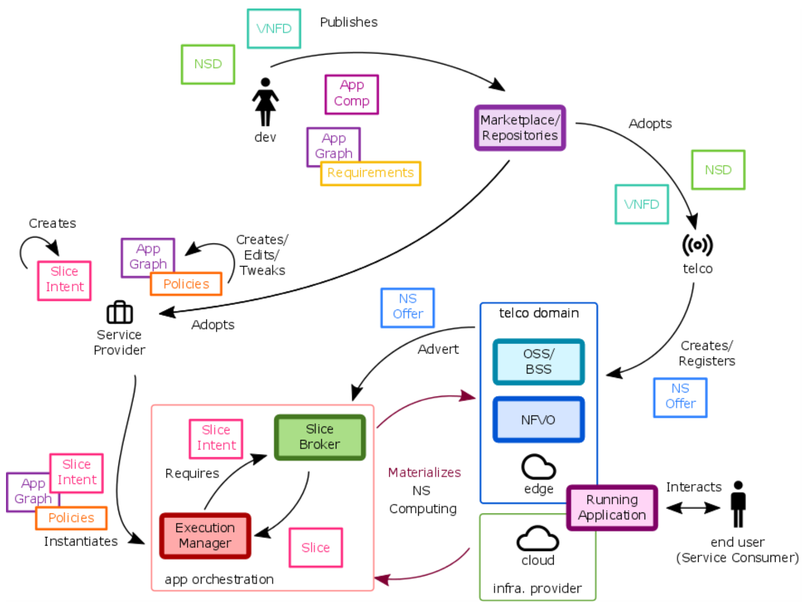

Fig. 6 MATILDA workflow highlighting the different stakeholders and metamodels.

Five vertical applications in different domains will be demonstrated on top of the MATILDA architectural layers and components:

- ***High Resolution Media on Demand Vertical with Smart Retail Venues' Integration***, combining the functionalities of two systems, to provide 5G Personal Assistance in Crowded Events (5GPACE). The new framework can offer end-users Immersive Media Services combined with Machine Learning-based personal retail recommendations.
- ***Testing 4.0 - Distributed System Testing***, based on FastWAN, an experimental communication technology that was developed as a solution for the enablement of geographically separated real-time industrial test benches.
- ***5G Emergency Infrastructure with SLA Enforcement (5GPPRD)***, a 5G system for Public Protection and Disaster Relief (PPDR). It extends the capabilities of a real time intervention monitoring and critical infrastructure protection product suite (iMON), combined with a suite for performance monitoring engines and advanced Operation, Administration and Management (OAM) functionalities to support SLAs (qMON).

- ***Industry 4.0 Smart Factory – Distributed Logistics-Production & Maintenance Application***, providing a logistic and a production scenario: the former allows 5G-based tracking of transport vehicles and monitoring of transported goods; the latter focuses on real-time pattern detection for quality assurance and distance calculation in a human-robot collaborative (HRC) production environment for collision avoidance.
- ***Smart City Intelligent Lighting System***, aimed to provide an easy replicable solution with fast time to market, automated maintenance and a modular approach enabled by 5G application graphs that will assure better monetization of intelligent city lighting solutions.

The five demonstrators will be mapped over three different testbeds:

- The University of Bristol testbed, integrating an extensive environment of LTE radio, WiFi and mmWave devices, interconnected by fibre backhaul, and providing OpenStack on High Performance Computing nodes in Bristol, UK.
- The CNIT-S2N (Smart and Secure Networks National Lab) testbed in Genoa, Italy, based on a cloud infrastructure able to control computing resources at bare-metal level, autonomously instantiate virtual infrastructure managers or software components in automated unattended fashion, and connect them to software-defined networking and radio devices and user equipment (with equivalent testbeds being deployed by Ubitech and Cosmote in their premises in Greece).
- The Orange Romania Smart Lighting testbed deployed in Bucharest, Romania, integrating LTE/5G Lighting Sensors, radio access and VNFs hosted in the Orange Regional Datacentre, along with a Cloud middleware IoT platform.

5 Conclusions

We have briefly examined the ongoing development of architectural and functional solutions for orchestration and management in the cloud and in virtualized telco platforms, along with their perspectives in the framework of 5G. The separation between VAO and NFVO has been highlighted, along with their interaction through the concept of slice intent and the OSS. We have briefly described some of the application and NFV orchestration platforms and mechanisms, and illustrated in some detail the philosophy adopted in this respect within the framework of the MATILDA project.

As a final remark, it is to be noted that, within the challenges posed by the Future Internet in general, and particularly by the strong wireless/wired integration of the 5G environment, four broad topics, among others, can be seen as interacting and mutually influencing: i) flexibility, programmability and virtualization of network functions and services, ii) performance requirements (in terms of users' Quality of Experience – QoE – and its mapping onto Quality of Service – QoS – in the network), iii) energy efficiency, and iv) network management and control.

The first item stems from the evolution of the network towards a multi-purpose "softwarized" service-aware platform upon a heterogeneous infrastructure, to deal with the diverse and integrating service paradigms of 5G. Performance issues have to deal with the very strong requirements imposed by 5G Key Performance Indicators (KPIs), and energy-awareness cannot be neglected in view of sustainability, environmental concerns, and operational costs. In this scenario, network management and control strategies are essential to orchestrate all needed functionalities, supervise and optimize the allocation of resources, to ensure that KPIs are met for network slices under the dynamic evolution of user-generated traffic, multiple tenants, service and infrastructure providers. Indeed, though a general reduction in Operational Expenditures (OpEx) is expected [43] (besides the reduction in Capital Expenditures – CapEx – entailed by the use of general-purpose hardware) from the upcoming revolution in networking paradigms brought forth by SDN and NFV, this reduction will not come without the adoption of specific management and control solutions. Some issues in this perspective have been examined in [44, 45], among others.

References

1. D. Szabó, F. Németh, B. Sonkoly, A. Gulyás, and F. H. P. Fitzek, "Towards the 5G Revolution: A Software Defined Network Architecture Exploiting Network Coding as a Service," in *Proc. ACM Conference of the Special Interest Group on Data Communication (SIGCOMM)*, Aug. 2015, pp. 105–106.
2. R. Vilalta *et al.*, "TelcoFog: A Unified Flexible Fog and Cloud Computing Architecture for 5G Networks," *IEEE Commun. Mag.*, vol. 55, no. 8, pp. 36–43, 2017.
3. R. Mijumbi, J. Serrat, J.-L. Gorricho, N. Bouten, F. D. Turck, and R. Boutaba, "Network Function Virtualization: State-of-the-art and Research Challenges," *IEEE Communications Surveys & Tutorials*, vol. 18, no. 1, pp. 236–262, Q1 2016.
4. D. Soldani and A. Manzalini, "Horizon 2020 and Beyond: On the 5G Operating System for a True Digital Society," *IEEE Veh. Technol. Mag.*, vol. 10, no. 1, pp. 32–42, Mar. 2015.
5. S. Yi, C. Li, and Q. Li, "A Survey of Fog Computing: Concepts, Applications and Issues," in *Proc. ACM Workshop Mobile Big Data (Mobidata)*, 2015, pp. 37–42.
6. ETSI, "GS MEC 003 v.1.1.1 Mobile Edge Computing (MEC); Framework and Reference Architecture," Mar. 2016. [Online]. Available: http://www.etsi.org/deliver/etsi_gs/MEC/001_099/003/01.01.01_60/gs_MEC003v010101p.pdf
7. P. Mell and T. Grance, "The NIST Definition of Cloud Computing," Recommendations of the National Institute of Standards and Technology, Special Publication 800-145, Sep. 2011. [Online]. Available: http://nvlpubs.nist.gov/nistpubs/Legacy/SP/nistspecialpublication800-145.pdf
8. S. S. Manvi and G. K. Shyam, "Resource Management for Infrastructure as a Service (IaaS) in Cloud Computing: A Survey," *J. Netw. Comput. Appl.*, vol. 41, pp. 424–440, 2014.
9. D. Weerasiri, M. C. Barukh, B. Benatallah, Q. Z. Sheng, and R. Ranjan, "A Taxonomy and Survey of Cloud Resource Orchestration Techniques," *ACM Comput. Surveys*, vol. 50, no. 2, May 2017, art. 26.
10. ETSI, "GS NFV-INF v.1.1 Network Functions Virtualisation (NFV); Infrastructure Overview," Jan. 2015. [Online]. Available: http://www.etsi.org/deliver/etsi_gs/NFV-INF/001_099/001/01.01.01_60/gs_nfv-inf001v010101p.pdf
11. T. Taleb, K. Samdanis, B. Mada, H. Flinck, S. Dutta, and D. Sabella, "On Multi-Access Edge Computing: A Survey of the Emerging 5G Network Edge Cloud Architecture and Orchestration," *IEEE Communications Surveys & Tutorials*, vol. 19, no. 3, pp. 1657–1681, Q3 2017.

12. “The Openstack project.” [Online]. Available: https://www.openstack.org/
13. S. Thalanany and P. Hedman, “Description of Network Slicing Concept,” NGMN 5G P1 Requirements and Architecture, Work Stream End-to-End Architecture, v.1.0.8, Sep. 2016. [Online]. Available: https://www.ngmn.org/uploads/media/161010_NGMN_Network_Slicing_framework_v1.0.8.pdf
14. 3GPP, “Study on Management and Orchestration of Network Slicing for Next Generation Network,” TR 28.801, v.15.0.0, Sep. 2017.
15. J. O. nez Lucena, P. Ameigeiras, D. Lopez, J. J. R.-M. noz, J. Lorca, and J. Folgueira, “Network Slicing for 5G with SDN/NFV: Concepts, Architectures, and Challenges,” *IEEE Commun. Mag.*, vol. 55, no. 5, pp. 80–87, May 2017.
16. “The MATILDA Project.” [Online]. Available: www.matilda-5g.eu
17. P. Gouvas, A. Zafeiropoulos, C. Vassilakis, E. Fotopoulou, G. Tsiolis, R. Bruschi, R. Bolla, and F. Davoli, “Design, Development and Orchestration of 5G-ready Applications over Sliced Programmable Infrastructure,” in *Proc. 29th Int. Teletraffic Congress (ITC 29)*, Sep. 2017, pp. 13–18.
18. “Eucalyptus Cloud Platform.” [Online]. Available: github.com/eucalyptus/eucalyptus
19. “OpenNebula.” [Online]. Available: opennebula.org
20. “Apache CloudStack™.” [Online]. Available: cloudstack.apache.org
21. A. Vogel, D. Griebler, C. A. F. Maron, C. Schepke, and L. G. Fernandes, “Private IaaS Clouds: A Comparative Analysis of OpenNebula, CloudStack and OpenStack,” in *Proc. 24th Euromicro Int. Conf. on Parallel, Distributed, and Network-Based Processing (PDP)*, Feb. 2016, pp. 672–679.
22. S. Shahzadi, M. Iqbal, Z. U. Qayyum, and T. Dagiuklas, “Infrastructure as a Service (IaaS): A Comparative Performance Analysis of Open-Source Cloud Platforms,” in *IEEE 22nd Int. Workshop on Computer Aided Modeling and Design of Commun. Links and Netw. (CAMAD)*, Jun. 2017, pp. 1–6.
23. “Cloudify.” [Online]. Available: cloudify.co
24. “Juju Orchestrator.” [Online]. Available: jujucharms.com
25. “Openstack Heat.” [Online]. Available: wiki.openstack.org/wiki/Heat
26. SDxCentral, “2018 NFV Report Series Part I: NFV Infrastructure (NFVI) and VIM,” Market Report, Apr. 2018.
27. ——, “2018 NFV Report Series Part 2: Rising roles of MANO, LSO and Assurance,” Market Report, May 2018.
28. Linux Foundation, “Harmonizing Open Source and Standards in the Telecom World,” Networking and Orchestration White Paper, May 2017.
29. AGC Research, “Succeeding on an Open Field: The Impact of Open Source Technologies on the Communication Service Provider Ecosystem,” AGC Research Report, 2018.
30. “The ARCADIA Orchestrator.” [Online]. Available: www.ubitech.eu/ubitech-releases-the-first-version-of-the-arcadia-orchestrator
31. ETSI, “OSM Release FIVE Technical Overview,” OSM White Paper, Jan. 2019.
32. “The OpenVIM Project.” [Online]. Available: github.com/nfvlabs/openvim
33. “The ONAP Architecture.” [Online]. Available: wiki.onap.org/display/DW/Architecture
34. “Tacker.” [Online]. Available: wiki.openstack.org/wiki/Tacker
35. “OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA).” [Online]. Available: www.oasis-open.org/committees/tc_home.php?wg_abbrev=tosca
36. “Open Baton.” [Online]. Available: openbaton.github.io
37. “Gohan: API Gateway Server.” [Online]. Available: gohan.cloudwan.io
38. “The T-NOVA Project.” [Online]. Available: www.t-nova.eu
39. R. Bruschi, R. Bolla, F. Davoli, A. Zafeiropoulos, and P. Gouvas, “Mobile Edge Vertical Computing over 5G Network Sliced Infrastructures: An Insight into Integration Approaches,” *IEEE Commun. Mag.*, vol. 57, no. 7, pp. 78–84, Jul. 2019.
40. “Network and Computing Slice,” MATILDA Deliverable D4.2, Aug. 2019.
41. “The Ericsson Network Manager.” [Online]. Available: www.ericsson.com/en/portfolio/digital-services/automated-network-operations/network-management/network-manager

42. "OpenDaylight." [Online]. Available: www.opendaylight.org
43. E. Hernandez-Valencia, S. Izzo, and B. Polonsky, "How Will NFV/SDN Transform Service Provider OpEx?" *IEEE Netw.*, vol. 29, no. 3, pp. 60–67, May/June 2015.
44. R. Bolla, R. Bruschi, F. Davoli, C. Lombardo, J. F. Pajo, and O. R. Sanchez, "The Dark Side of Network Functions Virtualization: A Perspective on the Technological Sustainability," in *Proc. IEEE Intern. Conf. Commun.*, May 2017, pp. 1–7.
45. M. Aicardi, R. Bruschi, F. Davoli, P. Lago, and J. F. Pajo, "Decentralized Scalable Dynamic Load Balancing among Virtual Network Slice Instantiations," in *Proc. IEEE Global Telecommun. Conf.*, Dec. 2018, pp. 1–7.