MATILDA 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 network service domains. To do so, it adopts a novel and holistic approach 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 network services 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 vertical application orchestrator (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.

During the **second year** of the project this service and functional separation concept has been clearly reflected into the final architecture, whose main elements, shown in Fig. 1A, are currently all under development. An original solution has been adopted for the integration of the 5G vApps into the 5G ecosystem at the VIM-level, by keeping 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. 1B.

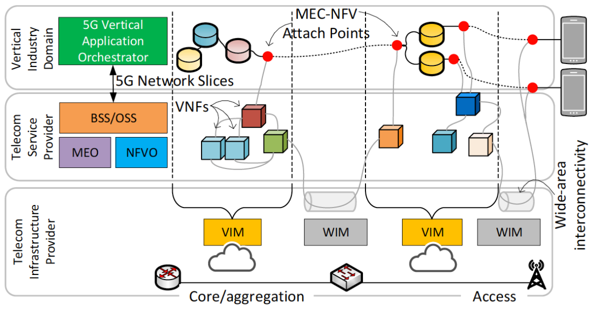

A)

Fig. 1. A) 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.

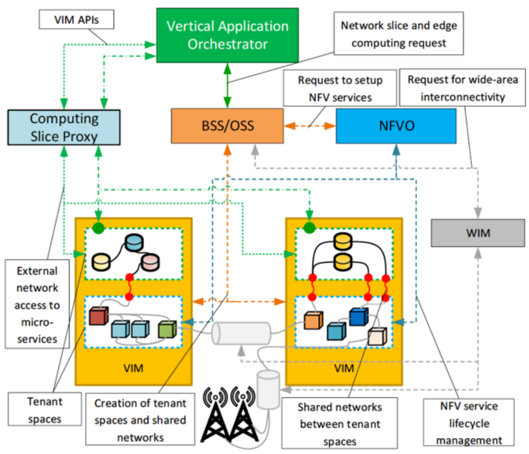

B)

Fig. 1.B) VIM-level integration.

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 their development is ongoing. The main components constituting the MATILDA vertical application orchestrator 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 vertical application orchestrator and the 5G programmable infrastructure management mechanisms.
A summary of the overall lifecycle of an application created with the MATILDA framework is represented in Fig. 2, highlighting the interaction among the different stakeholders and the usage of metamodels.

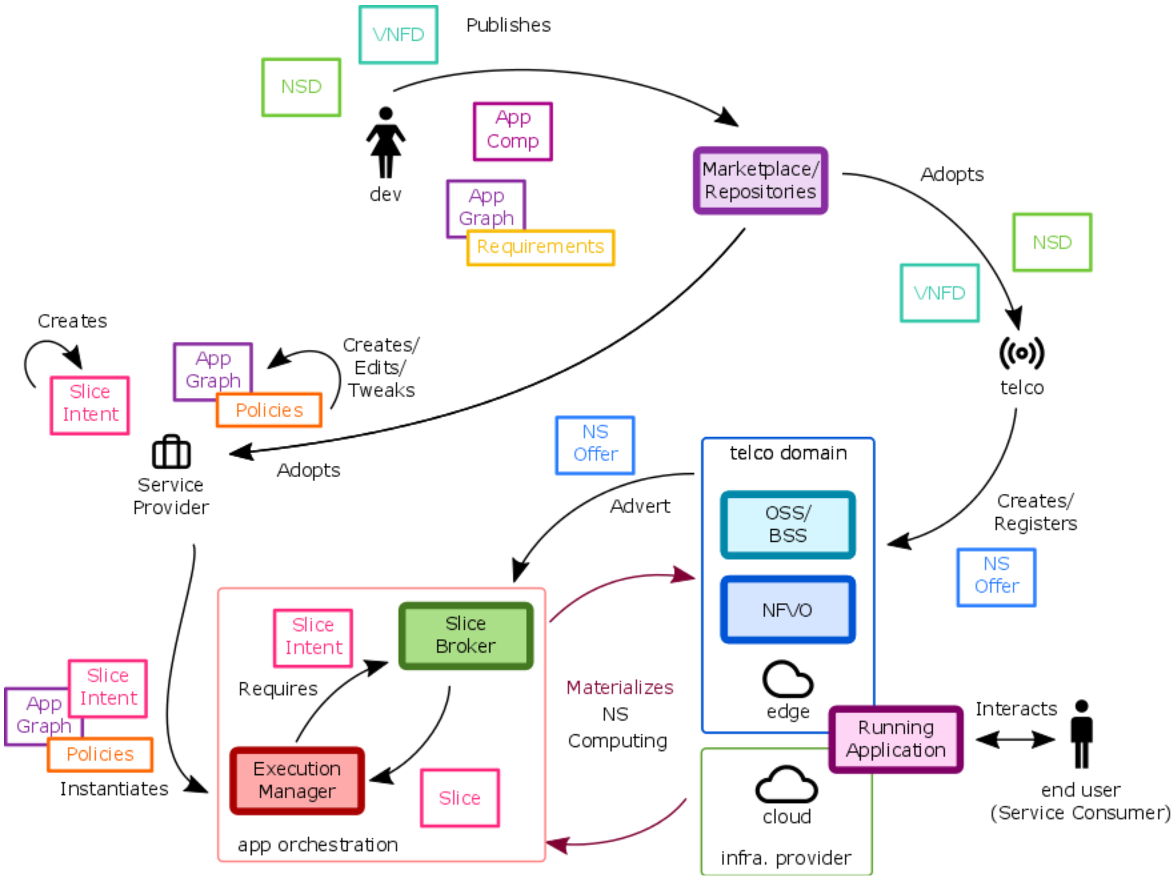

Fig. 2. MATILDA workflow highlighting the different stakeholders and metamodels.

Five vertical applications in different domains will be demonstrated on top of the MATILDA architectural layers and components. They are listed below, along with their Network and Operational KPIs.

- <u>*High Resolution Media on Demand Vertical with Smart Retail Venues' Integration*</u>, 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. *Network KPIs*: Device Density (~32 per small cell, ~50 per WiFi Hot Spot), Mobility (static users/low (0÷3 m/s)), Availability & Reliability (>99%), User Data Rate (~10 Mbps/user, depending on quality), End-to-end Latency (Max 1 s), Access Interoperability (must be available), Edge Computing (must be available), Storage at the Edge (must be available), Computing Acceleration at the Edge (must be available), Network Slicing Capability (must be available); *Operational KPIs*: 5GPACE App Deployment Time (~90 min), 5GPACE App On-Boarding Time (~15 min), Resource Usage Monitoring (must be available), 5GPACE App Component Scalability (must be available), Scaling Time (~20 s), Availability & Reliability (>99%), 5GPACE App Repository (must be available), Locality Awareness (must be available), Hardware Video Acceleration Management (must be available), Multi-site Management (must be available).
- <u>*Testing 4.0 - Distributed System Testing*</u>, based on FastWAN, an experimental communication technology that was developed as a solution for the enablement of geographically separated real-time industrial test benches. *Network KPIs*: Flexible

Bandwidth Allocation (up to 10 Mbps/node (FastWAN Unit)), Low Delay/Latency (inside Germany – ~50 ms in-node latency; inside Europe – ~100 ms in-node latency; Worldwide – ~200 ms in-node latency), Interoperability with Various Access Networks (WLAN, LTE, Ethernet). *Operational KPIs*: High Availability (99.99% of operational time), Resource Usage Monitoring (must be available), Component Scalability (must be available), Deployment Time (~90 min), On-Boarding Time (~15 min), Locality Awareness (must be available), Multi-Site Management (must be available).

- *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). *Network KPIs*: Availability & Reliability (>99.999%), Network Slicing Capability (must be available), End-to-end Latency for Interactive Applications (<20 ms), End-to-end Latency for Mission-Critical Applications (<1 ms), Bandwidth (~20 Mbps/user), Jitter (<1 ms), Packet Loss (<0.01%). *Operational KPIs*: iMON Dashboard Components On-Boarding Time (~15 min.), iMON Dashboard Component Deployment Time (~2 min), iMON Dashboard Application Graph Deployment Time (~5 min), Resource Usage Monitoring (must be available), iMON Dashboard Component Scalability (must be available), Scaling Time (~30 s), Availability & Reliability (>99.99%).
- *Industry 4.0 Smart Factory – Inter and Intra-Enterprise Integration*, focusing on a logistic application, which offers customers the possibility to track, change and prioritize their orders. *Network KPIs*: Device Density (~100 per LAN/WiFi Hot Spot), Bandwidth (up to ~10 Mbps/user, depending on quality), Availability & Reliability (WLAN, LTE, Ethernet), Delay/ Latency (distance calculation scenario: the system is expecting a response after 100 ms; pattern matching scenario: the system is expecting a response after 250 ms), Access Interoperability (must be available), Security & Privacy (must be available). *Operational KPIs*: Deployment Time (~90 min), Availability & Reliability (>99%), Resource Usage Monitoring (must be available), Component Scalability (must be available), Locality Awareness (must be available), Multi-Site Management (must be available).
- *Smart City Intelligent Lighting System*, deployed in Alba Iulia, a small- to middle-size city in Romania with about 70,000 inhabitants, in order 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 the intelligent city lighting solution. *Network KPIs*: Availability (>99.99%), Total Slice Bandwidth (~100 Mbps), End-to-End Latency (<300 ms), Jitter (~100 ms), Packet Loss (<0.1%). *Operational KPIs*: Device Status (100 Smart Light sensors), Service Availability (>99.99%), Device Bandwidth Capacity (~0.1 Mbps).

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 City testbed in Alba Iulia, Romania, integrating LTE/5G Lighting Sensors, radio access and VNFs hosted in the Orange Regional Datacentre, along with a Cloud middleware IoT platform.