Dynamic and Quality-aware Network Slice Management in 5G Testbeds

Vincent Charpentier*, Nina Slamnik-Kriještorac*, Juan Brenes†, Andreas Gavrielides‡, Marius Iordache§, Georgios Tsiouris¶, Lian Xiangyu‖, and Johann M. Marquez-Barja*
* University of Antwerp - imec, IDLab - Faculty of Applied Engineering, Belgium
† Nextworks, Italy
‡ eBOS Technologies Ltd., Cyprus
§ Orange, Romania
¶ COSMOTE, Greece
‖ Telenet Group, Belgium

***Abstract*—The proliferation of 5G technology is enabling vertical industries to improve their day-to-day operations by leveraging enhanced Quality of Service (QoS). One of the key enablers for such 5G performance is network slicing, which allows telco operators to logically split the network into various virtualized networks, whose configuration and thus performance can be tailored to verticals and their low-latency and high throughput requirements. However, given the end-to-end perspective of 5G ecosystems where slicing needs to be applied on all network segments, including radio, edge, transport, and core, managing the deployment of slices is becoming excessively demanding. There are also various verticals with strict requirements that need to be fulfilled. Thus, in this paper, we focus on the solution for dynamic and quality-aware network slice management and orchestration, which is simultaneously orchestrating network slices that are deployed on top of the three 5G testbeds built for transport and logistics use cases. The slice orchestration system is dynamically interacting with the testbeds, while at the same time monitoring the real-time performance of allocated slices, which is triggering decisions to either allocate new slices or reconfigure the existing ones. In this paper, we illustrate the scenarios where dynamic provisioning of slices is required in one of the testbeds while taking into account specific latency/throughput/location requirements coming from the verticals and their end users.**

***Index Terms*—5G, network slicing, transport & logistics, VITAL-5G, third-party experimenters**

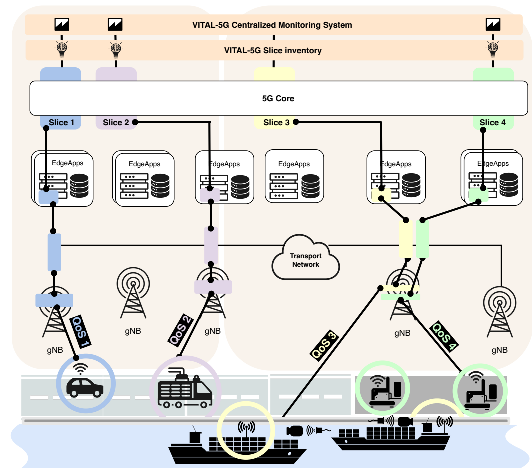

Fig. 1: End-to-end network slicing concept applied to 5G ecosystem with different verticals and their QoS requirements.

I. Introduction and Motivation

With the advent of 5G technology, vertical industries such as automotive, Transport & Logistics (T&L), and e-health have the opportunity to enhance their own services when making use of 5G technology. This is made possible since 5G-networks can achieve ultra-reliable (99.999%) low-latency (5 ms) network connectivity, including ultra-fast data rates that are 100x larger than in previous generations (4G) and with the density of 1 million connections/km^2. Taking into account Ericsson's projection that 45% of all global internet traffic will go over a 5G connection, thereby covering 65% of the world population [1], it is reasonable to expect that 5G systems will gain a prevailing role in modern production and distribution systems in the T&L sector [2] toward improving the efficiency and safety of T&L operations by automating and optimizing processes and resource usage [2]. Thus, to help accelerate this growth, adoption, and economic benefits, for verticals and in particular, for the T&L sector, the European Commission has co-financed the VITAL-5G project. The VITAL-5G project provides an enhanced 5G-empowered experimentation facility, with a portfolio of corresponding Network Applications, which enables T&L application developers to test and validate their T&L use cases within 5G ecosystems, in a user-friendly and intuitive manner, thus significantly reducing operational inefficiency and reducing market entry barriers prior to deploying their T&L service to live 5G networks [2].

One of the key technical aspects of 5G networks is *network slicing*, which allows operators to slice the whole network into several logical, virtualized, and high-performing networks, whose configuration and thus performance can be adjusted to fit the Quality of Service (QoS) requirements coming from verticals (Figure 1). With the main benefit that all slices are isolated from each other. Unlike 5G Non-standalone (NSA), 5G Standalone (SA) is deployed using the 5G Core functionalities, providing means for creating 5G slices in an end-to-end manner, i.e., stretching from the radio resources over edge and transport to the core, as illustrated in Figure 1. In the H2020 VITAL-5G project[1], we have focused on creating holistic 5G ecosystems based on 5G Release 16[2] in the port and warehouse environments to boost the T&L sector and explore the potential of 5G technology for improving the efficiency and safety of T&L operations (e.g., port dwell and waiting times).

In general, by creating different network slices for different types of traffic, network providers can ensure that each slice has the resources it needs to meet the needs of its users, which is beneficial in terms of more efficient use of resources, greater flexibility, cost savings, improved performance, increased privacy and new revenue streams. For Transport & Logistics (T&L) in particular network slicing brings: (i) improved efficiency, as operations such as real-time tracking of vehicles and cargo can be improved with high throughput on the uplink for HD

[1]VITAL-5G: https://www.vital5g.eu/.

[2]5G Release 16 https://https://www.3gpp.org/specifications-technologies/releases/release-16/

camera stream from vehicles/vessels, (ii) increased reliability, by leveraging Ultra-Reliable Low-Latency Communication (URLLC) slices that ensure safe network operations even in the case traffic congestion, and (iii) more flexibility, as T&L industry can offer more advanced services that, due to network connection requirements, were not possible with 4G (decreased latency and improved throughput due to network slicing and edge computing). To achieve the aforementioned benefits, i.e., to incentivize innovation and improve operational efficiency for verticals with the help of 5G and beyond technologies, vertical industries need to understand the behavior of the 5G systems and their potential, but there is a need to abstract the network complexity. One way to do it is through adopting the design and deployment of Edge Network Applications (EdgeApps) because they abstract the complexity of the underlying 5G infrastructure to T&L application developers [2], and allow them to define their requirements on a high-level. Thus, in this paper, we focus on the VITAL-5G solution for network slice management and orchestration, which is the main feature of the Slice inventory component within the VITAL-5G platform. We explain the features that this component brings and how it dynamically interacts with the three testbeds delivered in the project, thereby illustrating scenarios where dynamic provisioning of slices is required in one of the testbeds while taking into account specific latency/throughput/location requirements coming from the verticals and their end users.

For the purpose of facilitating the experience of experimenting with 5G systems, in the VITAL-5G project, we produced the VITAL-5G platform backend that operates on top of the distributed 5G facilities composed of the three VITAL-5G testbeds. The VITAL-5G platform provides functional elements and management backend services. While each VITAL-5G testbed offers a 5G infrastructure including cloud and edge resources where Network Applications or EdgeApps can be dynamically instantiated and orchestrated leveraging 5G network slices that are customized on the basis of the experiment requirements. These enhanced 5G-empowered experimentation facilities consist of the three well-established 5G testbeds, i.e., (1) Antwerp 5G testbed is located in the Port of Antwerp, for the purpose of deploying and testing Assisted Vessel Transport and similar use cases, (2) Galati 5G testbed is located in Danube area in Romania, providing 5G connectivity and data-enabled assisted navigation using Internet of Things (IoT) sensing and video cameras, and (3) Athens 5G testbed is located in Athens for enabling automation and remote operation of freight logistics.

In such complex environments with verticals demanding various network slice profiles for their services, the provisioning of network slices and their runtime management needs to be efficient and automated.

II. Related Work

Given the dynamicity and scalability that slicing brings to the 5G systems, Management and Orchestration (MANO) in multi-tenant scenarios is not a straightforward task. Dynamic provisioning of network slices is posing an additional challenge toward State-of-the-Art (SotA) MANO systems because they are not designed to respond to quick changes in resource allocation and traffic fluctuations imposed by an increased number of verticals using 5G. However, the work focusing on intelligent slice management mechanisms is still in the infancy stage, especially the application of theoretical concepts in real-life systems with 5G radio and core functions, and their validation within 5G testbeds. This is one of the crucial steps toward fulfilling the high demands of verticals in a dynamic and quality-aware manner. For example, the work presented by Zhang et al. [3] concluded that extensive work is needed for optimizing slice orchestration, thereby leveraging existing Artificial Intelligence (AI) concepts to realize an intelligent slice orchestration system [3]. The majority of current research is now focused on the different phases of network slice management and thus how network operators need to allocate network resources to maximize their benefits [4,5]. A recent survey on slicing management conducted by Wu et al. [6] reported that most survey papers in the current literature related to slicing management focus on the enabling technologies of network slicing e.g., Software-Defined Networking (SDN), Multi-access Edge Computing (MEC). Thus, only a few papers discuss in-depth the application requirements for network slicing [6]. To bridge this gap in network slicing surveys Wu et al. provided a survey of slicing management from the perspective of vertical applications, i.e., smart transportation, smart energy, and smart factory/manufacturing since these are the three key services from the backbone of Industrial Internet of Things (IIoT) a subdomain of IoT) [6]. Further, Wang et al. [5] proposed a simulation-based Machine Learning (ML) scheme for dynamic resource scheduling for networks slicing, aiming to achieve automatic and efficient resource optimization and End-to-End (E2E) service reliability [5]. However, despite the research potential to explore more innovative, intelligent (i.e., AI-empowered) network slice provisioning orchestration frameworks, the SotA is still limited to simulation work.

In the comparison of MANO solutions for the 5G era [7], Trakadas et al. presented a framework named SONATA (5G-TANGO project[3]), for managing and orchestration of Network Services (NSs). They compared SONATA in a real testbed with the two well-known open-source MANO frameworks, i.e., Open Source MANO (OSM) and Cloudify [7]. They concluded that Couldify was suitable for deployments that have no strict requirements like network slicing and Service-Level Agreement (SLA) contracts [7]. While OSM proved to be robust and mature [7], SONATA provides a complete toolchain for automated NS management including network slicing managers [7]. Some other projects like 5G-TRANSFORMER, SLICENET, and MATILDA-5G, worked on a similar topic. In particular, the SLICENET project designed and developed a framework for slice management and monitoring oriented to verticals, allowing greater control over the slice. While the SLICENET framework is capable of monitoring the performance of slices, it is not capable of predicting what may happen in the near future to the network requirements of the slice for the 5G applications that make use of the slice, since it does not have real-time access to the application metrics, which is the case of our VITAL-5G Slice Inventory. Furthermore, 5G-TRANSFORMER aimed to allow the verticals to request network slices via a high-level abstracted view, letting the underlying architecture take care of infrastructural decisions. The service orchestrator mechanisms take decisions based on the slice requirements imposed by the network context and the different verticals' service requests [8], unlike the VITAL-5G approach where the slices and EdgeApps are being monitored in real time, to guarantee the optimal utilization of the slice. Finally, MATILDA-5G [9] dealt with slices as abstractions for the underlying resources to orchestrate intelligently and uniformly the creation and maintenance of slices. Compared to VITAL-5G, MALITDA-5G did not retrieve the needed metrics of the 5G applications and Network slices over defined interfaces, nor tested the solution with the demanding needs of the verticals.

From such analysis, we can clearly identify some of the gaps in current research and development activities concerning network slicing in real-life systems. In particular, real-time monitoring of slices (e.g., latency, bandwidth, data rates) and Network applications (i.e., EdgeApps), is not considered when making slice configuration-related decisions. Such dynamic monitoring of network and application performance is crucial for intelligent and proactive operations, i.e., prediction of traffic

[3]5G-TANGO: https://www.5gtango.eu/.

demand and performance of slices and EdgeApps, with reference to the QoS requirements from verticals.

III. Slice Management Framework

The VITAL-5G slice management framework relies on the concept of Network Application (i.e., EdgeApps). The main objective behind the Network Application concept is to extend the typical Network Functions and Network Services to enable the seamless deployment of the applications in 5G-enabled virtualized infrastructures. The key enabler for this is the introduction of the Network Application blueprint, which complements the application orchestration descriptors allowing the application developers to establish the 5G slice profile characteristics (i.e. uplink/downlink data rates, end-to-end latency, etc) required by the application [2,10]. This information is used during the Lifecycle Management (LCM) of the services, to allocate the 5G slice for the service instance either by using a pre-provisioned slice or by dynamically provisioning one for the service instance.

A. Slice Inventory: Architecture and Features

The VITAL-5G Slice Inventory, illustrated in Fig. 2 is a component of the VITAL-5G platform, which is in charge of keeping track of the 5G slices across the distributed 5G testbeds, dynamically managing and provisioning the slices, and selecting the optimal slice for a given experiment based on its requirements. A more detailed overview of designed features is given below.

- **Dynamic retrieval of 5G slice profile features (Slice Manager and Storage):** After configuring the slice on the 5G testbed, either in a pre-configured or dynamic way, the slice profile features must be shared with the Slice Inventory. Afterward, all information related to the slices, such as configuration parameters, attached User Equipments (UEs), and testbed capabilities, is stored. The objective is to give a current view of the available slices to anyone who is interested in deploying EdgeApps using the 5G slices on a particular testbed.
- **Managing the lifespan of network slices (Slice Manager):** Based on the capabilities of the VITAL-5G testbeds, the Slice Inventory module uses its southbound interface to handle the lifespan of network slices (such as slice provisioning, attaching UE, and configuring slices dynamically), by sending management requests to the local slice plugins deployed on each VITAL-5G testbed, as shown in Figure 2.
- **Slice allocation and querying (Slice Manager):** Due to the performance requirements defined by verticals, it is crucial to select the appropriate slice for deploying the service, once the experiment is launched and the vertical service is onboarded [11]. The Slice Inventory provides interfaces for searching for the 5G slices currently deployed in the VITAL-5G testbeds, as well as assigning the 5G slice that can be used for a specific vertical service (dynamically or in a pre-defined way) [11]. Therefore, Service LCM, which is described in Section III-B, is sending the request for slice provisioning based on the EdgeApp requirements (Fig. 2 [11]). Slice Inventory is taking into account the real-time performance of slices deployed in the 5G testbeds, and makes a final decision about the slice selection. The UE will be attached to a slice if it meets the slice requirements and if the VITAL-5G testbed supports it [11]. If not, a new slice will be created or existing ones will be reconfigured dynamically to satisfy the provisioning requirements, only if the VITAL-5G testbed has the capability to do so [11].

B. VITAL-5G modules interacting with Slice Inventory

In Fig. 2, we illustrate the interaction between VITAL-5G components and the Slice Inventory, and below we detail more on each of the platform components that Slice inventory is interacting with.

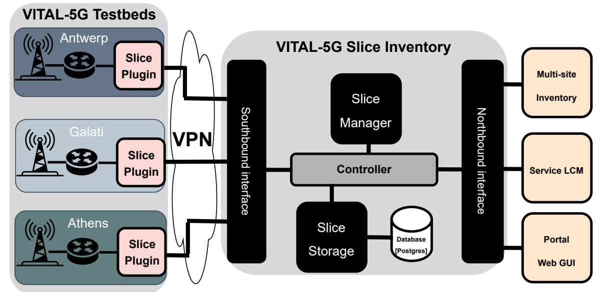

Fig. 2: The VITAL-5G slice inventory, including its interaction with the local slice plugins on the testbeds and the VITAL-5G Platform.

```
sliceinventory=# select * from urllc_slice;
 id | exp_data_ratedl | exp_data_rateul | jitter | latency |  site  |         slice_id          | traffic_type
----+-----------------+-----------------+--------+---------+--------+---------------------------+--------------
  1 |            1000 |            1000 |      1 |      10 | galati | Galati_merged_EMBB_URLLC  | critical
  2 |            1000 |            1000 |      1 |      10 | galati | Galati_URLLC              | critical
(2 rows)

sliceinventory=# select * from urllc_slice_imsi;
 urllc_slice_id |      imsi
----------------+-----------------
              1 | 226107900000077
              2 | 226107900000001
              2 | 226107900000009
              2 | 226107900000075
```

Fig. 3: Slice plugin output on the 5G testbed.

a) Service Life-Cycle Manager (LCM): is the module responsible for processing the incoming LCM requests, mainly interacting with the: (i) Testbed Network Function Virtualization Orchestrator (NFVO), for the LCM operations of the network services associated with the application, (ii) 5G Slice Inventory, for the allocation of the 5G slice associated to the service, and (iii) Centralized Monitoring Platform, for the configuration of the infrastructure and service metrics to be retrieved from the testbeds. In particular, for the allocation of the 5G slice to be attached to the service, the Service LCM uses the 5G slice profile information established in the Network Application blueprint to determine the 5G slice characteristics to be requested from the 5G Slice Inventory through its REST API.

b) The Portal Web GUI: is the interface that allows users to access the platform's back-end tools and functions through a user-friendly interface, enabling them to design experiments, choose experiments, EdgeApps, and service blueprints. The interface is intended for internal experimenters within the consortium and external experimenters who want to use the existing EdgeApps or onboard their own. The GUI provides individual account access, Network Application validation, Blueprint uploading, access to testbed assets, and visualization of testbed results in terms of key performance indicators. The GUI itself does not have exposed APIs but rather connects with the centralized monitoring platform that is in charge of collecting the Network and Service KPIs and displaying them on-demand for the user, the Experiment and Service LCM, the NetApps, Services, and Experiments Catalogue as well as the Blueprints Validation tool which is integral for Quality of Experience (QoE).

c) Multi-site Inventory: is the centralized repository of the testbed information, containing a catalog of (i) IoT devices available in the testbed facilities, (ii) Testbed NFVO endpoints, and (iii) Testbed 5G slicing capabilities and endpoints. The info about testbed slicing capabilities is used to determine if the site supports dynamic allocation of slices, dynamic modification of the slice profile characteristics and if the Slice management function of the testbed supports dynamic attachment of UEs to the already deployed slices. Moreover, this inventory may also store the endpoints to reach the testbed slice management functions.

C. Slice plugins on the 5G testbeds

The VITAL-5G Slice Inventory uses its southbound interface to communicate with the local slice plugins deployed on each VITAL-5G testbed, as shown in Figure 2. Thus, once a slice is available on a specific testbed, either in a pre-configured or dynamic manner, the slice profile characteristics can be communicated with Slice Inventory. Furthermore, the Slice Inventory southbound interface is also used to perform LCM operations of network slices e.g., slice provisioning, UE attachment, and dynamic slice configuration based on the decisions made by Slice inventory and the real-time performance of slices that Slice inventory monitors. This way, the local slice plugins represent a bridge between the VITAL-5G platform (i.e., Slice Inventory), and the 5G network (i.e., 5G Core), in order to translate the network slicing request coming from the VITAL-5G platform towards the network of the specific VITAL-5G testbed.

The slice plugin performs dynamic slice management and UE attachment to a specific slice according to the requirements imposed by verticals (e.g., EdgeApp developer for T&L services), and received from the VITAL-5G platform. In Fig. 3, we showcase the structure of slice and UE-related data that Slice inventory is retrieving from the 5G testbeds via slice plugins. In the case of all 5G testbeds, this plugin acts as a mediator and translates the VITAL-5G platform slice characteristics to QoS flow binding parameters, which are communicated to the 5G Core through a REST API. These are the 5G QoS Identifier (5QI) values and Allocation and Retention Priority (ARP) that control the QoS forwarding treatment of each flow by the 5G Core and the 5G New Radio. In particular, the 5QI is a scalar value with a one-to-one mapping to a standardized combination of QoS characteristics specified in TS 23.501 [12], while ARP defines the relative importance and controls the pre-emption capabilities and vulnerabilities of the QoS flow. In particular, the 5G Core marks the traffic on the N3 interface for downlink and N6 interface on uplink with a 6-bit Differentiated Services Code Point (DSCP) value, in order for the traffic to be suitably treated during its passing through the network elements that constitute the transport and switching network.

For example, in the case of the Galati testbed, the slice plugin is developed as a python application (using Flask Web framework) and it is implemented to work as a proxy for the Slice Inventory and Slice management components of the VITAL-5G platform for (1) slice data retrieval, (2) UE attachment, (3) UE deletion from a slice and (4) dynamic slice creation/reconfiguration. Similarly, in the case of Antwerp and Athens, a Java-based variant of the plugin is created, but compatibility has been ensured using the same REST API configuration. In particular, dynamic slice reconfiguration is supported by the Galati testbed, which means that Slice inventory can trigger quality-aware reconfiguration of existing slices based on the collected data and internal logic running in the Slice inventory component. In that case, real-time information related to slices and UEs can be retrieved from the Unified Data Management (UDM). Also, a new UE can be easily provisioned in the system and attached to an existing slice in the testbed.

IV. Mechanisms for Real-time Slice Performance Monitoring

The role of the VITAL-5G Centralized Monitoring Platform is to collect various types of metrics from the three VITAL-5G testbeds and to make them accessible to other components of the VITAL-5G platform, e.g., the Slice Inventory [13]. In the case of Slice inventory, it is of utmost importance to monitor the real-time performance of network slices that are already deployed on the 5G testbeds, in order to make LCM decisions in a quality-aware and dynamic manner.

This platform can monitor metrics that reflect on the (i) network slice performance e.g., latency, peak throughput, average throughput, (ii) service metrics like response time, UE measured throughput, and (iii) infrastructure metrics like CPU load [13]. To retrieve those metrics, the Centralized Monitoring Platform interfaces with the local monitoring systems that are deployed on each of the VITAL-5G testbeds as shown in Figure 4. The local monitoring systems publish collected metrics to a topic specified by the Centralized Monitoring Platform [13], which then subscribes to that topic and collects these metrics either in real-time or in a bulk format [13]. The collected metrics are visualized on the Portal Web GUI through dynamically created Grafana dashboards and panels, with each metric creating a new panel on the dashboard, allowing easy monitoring of KPIs and fluctuations [13]. The main features of the centralized monitor are discussed below:

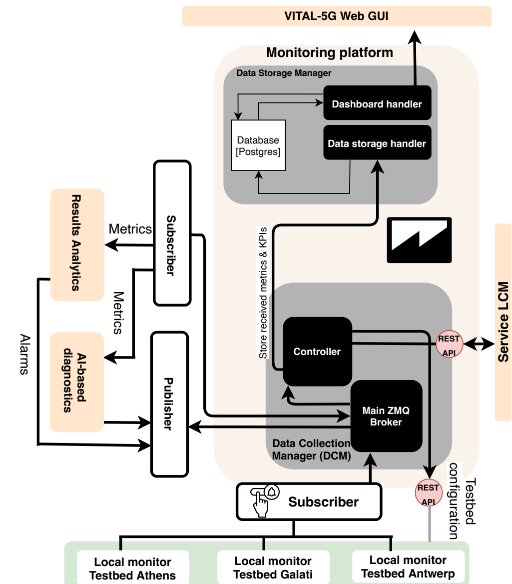

Fig. 4: Software architecture of the VITAL-5G centralized monitoring system.

- **Dynamic creation of monitoring topics:** As already mentioned above once a vertical service is created via the VITAL-5G platform, as a part of an ongoing experiment, the monitoring system needs to be triggered by the service LCM to create particular topics that will be used by a corresponding local testbed to report the metrics as also shown in Figure 5 [11]. The topic is composed of the type of metric that needs to be collected, also stating the source of metric (i.e., local testbeds such as Antwerp, Galati, and Athens) [11]. The communication between Service LCM and Centralized Monitoring Platform is REST-based as shown in Figure 5 [11].
- **Real-time collection of metrics from the 5G testbeds:** The VITAL-5G Centralized Monitoring Platform, triggered by the Service LCM, dynamically informs the relevant 5G testbed's local monitoring system of the metric type to be published, as shown in Figure 5. The local monitoring systems on the 5G testbeds are solely responsible for measuring metrics from the infrastructure, network, and services. The communication between the local monitoring system on a 5G testbed and the centralized monitoring system on the VITAL-5G platform is REST-based for topic configuration, while each 5G testbed utilizes a ZeroMQ publisher for collecting and publishing metrics to the Centralized Monitoring Platform.

- **Dynamic exposure of collected metrics towards external entities:** The internal ZeroMQ publisher on the Centralized Monitoring Platform makes all received metrics accessible to other VITAL-5G components e.g., Results Analytics and AI-based diagnostics to subscribe and collect the metrics for analysis and anomaly detection using the ZeroMQ subscriber (as seen in Figures 4 and 5) [11]. It is noteworthy that the collected data can be obtained not only by these two components but by any external entity willing to gather metrics from the 5G testbeds [11].
- **Dynamic visualization of the collected metrics:** The visualization of all metrics gathered from different 5G testbeds, through their local monitoring systems, is achieved by dynamically creating Grafana dashboards and panels [11]. Each collected metric results in a new panel on the Grafana dashboard that identifies the metric type and originating testbed, making it easier to view and monitor KPI changes [11].
- **Dynamic termination of data collection process:** Once the experiment is completed and the vertical service must be terminated, the collection for metrics for the service will also need to be terminated [11]. The Service LCM sends a REST-based termination request to the Centralized Monitoring Platform, which halts all active subscribers and deletes all Grafana dashboards to avoid any unused topics [11]. The Centralized Monitoring Platform informs the relevant testbeds and their local monitoring system that the service has been terminated, and there is no need for further pushing of metrics [11]. The metric collection is dynamically discontinued, but the previously collected metrics for this service and experiment will be kept on the Centralized Monitoring Platform as historical data and can be utilized for debugging or training machine learning models [11].

V. Orchestrating 5G slices across 5G testbeds: Proof-of-Concept

Given the description of the interaction between the VITAL-5G platform components and the Slice inventory in Section III-B, in Fig. 5 we visualize this interaction, showing the entire flow of the dynamic provisioning of slices. The flow refers to a specific experiment that includes a specific VITAL-5G testbed, e.g., the Antwerp testbed, while taking into account the specific latency/throughput/location requirements coming from the verticals and their end users.

A. Workflow of dynamic slice provisioning

In the first phase of the VITAL-5G Experimentation Procedure[4], the experimenter needs to create a vertical service for its vertical use case. This implies the creation of Network Applications (i.e., EdgeApps), which compose that specific vertical service. Besides the creation of the constituting EdgeApps, the experimenter also needs to specify the mobile network requirements for each EdgeApp, in the form of a Network Application blueprint. This way, the strict requirements related to the network can be provided to the VITAL-5G testbeds for each specific EdgeApp. Once all the EdgeApps are onboarded on the VITAL-5G platform (with all the necessary packages) in the third phase, the testbed will be configured accordingly to meet the required network performance for the T&L EdgeApps of the vertical service. This involves the provisioning or dynamic creation of network slices if the VITAL-5G testbed supports such configuration.The third phase is shown in Fig. 5, together with the interaction with the Centralized Monitoring Platform. Furthermore, as shown in Fig. 5, the Slice inventory obtains from the local slice plugins the information about all available slices per testbed. Thus, when then the Service LCM makes the request towards the Slice inventory, i.e., the third phase as mentioned above, with the request to attach a new UE to a specific slice on a specific testbed, the Slice inventory will then check if there is a slice that meets the requirements for the specific testbed the experimenter has chosen and if so, the UE will be attached to the slice. In case no slice is meeting the requirements, the Slice inventory will trigger the dynamic creation of a new slice, as shown in Fig. 5. The slice inventory and its orchestration mechanism manage all of the above-aforementioned logic. Furthermore, the slice orchestration mechanism is at the same time monitoring the real-time performance of the allocated slices and metrics from the EdgeApps of the vertical service itself. This way, proactive decisions can be triggered to either allocate new slices or reconfigure the existing slices to provide always the optimal network services for the T&L vertical users.

[4]https://www.vital5g.eu/elementor-1243/

B. Initial results

We have obtained initial experimentation results, which are shown in Fig. 6. The experiments reflect on the performance of the slice inventory, which is running on the VITAL-5G platform and not on the testbeds. The requests sent to the slice inventory are performed from any of the three testbeds. The first experiment refers to the calculation of the average time needed for a Slice Plugin on any of the testbeds to report its slices. This operation is shown in Fig. 5 flow line number 8, and on average it takes 0.239 seconds with a variance of 0.0013 seconds. The second experiment is the average time needed for obtaining and processing the requests from the Service LCM, as shown in figure 5 flow number line 9. This operation is performed on the Slice inventory side, and it includes the selection of an optimal slice for the T&L vertical service, it takes on average 0.245 seconds with a variance of 0.00578 seconds. We already expect when more intelligent slice Orchestrator algorithms are in place in the Slice Inventory, the request to select the most appropriate slice for the T&L vertical service will differ, and depending on the complexity of the algorithm, it might take longer. But due to the fact that Slice Inventory has access to real-time metrics of both the T&L EdgeApps of the vertical service and the network metrics of the slices of the testbeds, including all other network metrics of the testbeds, we expect that this proof-of-concept Slice Inventory will make more robust slice orchestration decisions over the long-term, and also make those decisions proactively so that the performance of the vertical service is not affected. This means that when a slice is provisioned, the decision about slice allocation and UE attachment will not change in a relatively short period, in order to meet the changing application requirements. In turn, this will save time compared to non-intelligent slice orchestrators, as there is a much higher chance that the slices will need to be swapped or reconfigured due to inefficient slice selection algorithms.

VI. Conclusion and Future work

In this paper, we addressed the need for dynamic and quality-aware network slice management for 5G testbeds. We provided a brief overview of SotA when it comes to slicing management techniques. We described in depth how the VITAL-5G slice management framework works, and how it interacts with other VITAL-5G components on the testbed, and on the VITAL-5G platform, with the goal of dynamically creating network slices and managing the existing ones. We also described how the real-time monitoring system works, and how it feeds the other VITAL-5G components with real-time network and service performance. Finally, we showed some preliminary results of our proof-of-concept solution for slice orchestration, studying the slice provisioning time on the Slice inventory component, and other impact factors such as retrieving slice information from the 5G testbeds.

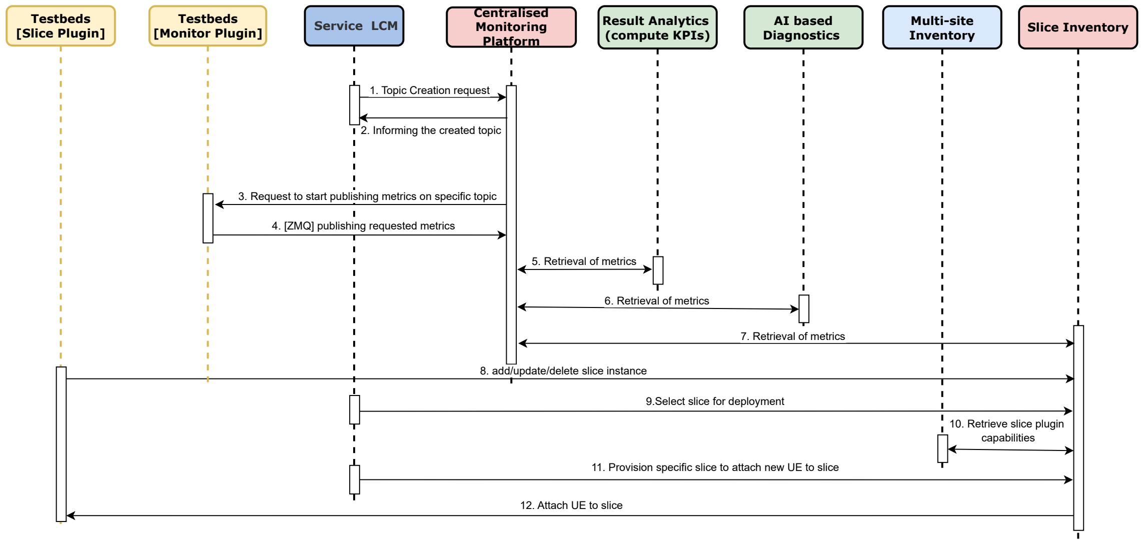

Fig. 5: Interactions between the service LCM, centralized monitoring platform, and the slice inventory.

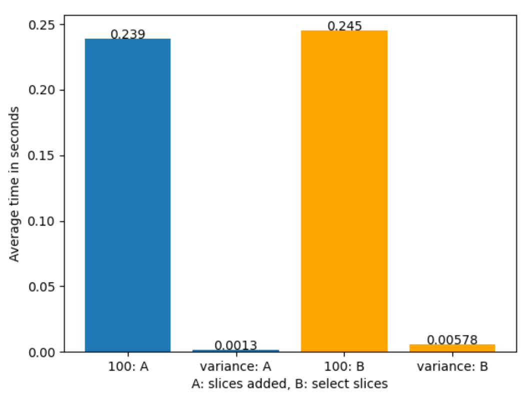

Fig. 6: Average time adding and selecting slices with the corresponding variance.

In future work, we are planning to further study more sophisticated ways of orchestrating network slices, applying native AI techniques, which will be the basis of future 6G systems as well. This approach is viable due to the fact that significant amounts of data can be collected for training since the Slice inventory component has real-time access to the network slice performance retrieved from the centralized monitoring system on the VITAL-5G platform.

VII. Acknowledgement

This work has been performed in the framework of the European Union's Horizon 2020 project VITAL-5G co-funded by the EU under grant agreement No. 101016567.

References

[1] Ericsson, "Ericsson Mobility Report," November 2019. [Online] Available: https://www.ericsson.com/assets/local/reports-papers/mobility-report/documents/2019/emr-november-2019.pdf, Last accessed on: January, 2023.

[2] N. Slamnik-Kriještorac, G. Landi, J. Brenes, A. Vulpe, G. Suciu, V. Carlan, K. Trichias, I. Kotinas, E. Municio, A. Ropodi, and J. M. Marquez-Barja, "Network applications (netapps) as a 5g booster for transport & logistics (t&l) services: The vital-5g approach," in *2022 Joint European Conference on Networks and Communications & 6G Summit (EuCNC/6G Summit)*, pp. 279–284, 2022. doi: http://dx.doi.org/10.1109/EuCNC/6GSummit54941.2022.9815830.

[3] S. Zhang, "An overview of network slicing for 5g," *IEEE Wireless Communications*, vol. 26, no. 3, pp. 111–117, 2019. doi: http://dx.doi.org/10.1109/MWC.2019.1800234.

[4] X. Li, M. Samaka, H. A. Chan, D. Bhamare, L. Gupta, C. Guo, and R. Jain, "Network slicing for 5g: Challenges and opportunities," *IEEE Internet Computing*, vol. 21, no. 5, pp. 20–27, 2017. doi: http://dx.doi.org/10.1109/MIC.2017.3481355.

[5] H. Wang, Y. Wu, G. Min, J. Xu, and P. Tang, "Data-driven dynamic resource scheduling for network slicing: A deep reinforcement learning approach," *Information Sciences*, vol. 498, pp. 106–116, 2019. doi: http://doi.org/10.1016/j.ins.2019.05.012.

[6] Y. Wu, H.-N. Dai, H. Wang, Z. Xiong, and S. Guo, "A survey of intelligent network slicing management for industrial iot: Integrated approaches for smart transportation, smart energy, and smart factory," *IEEE Communications Surveys & Tutorials*, vol. 24, no. 2, pp. 1175–1211, 2022. doi: http://dx.doi.org/10.1109/COMST.2022.3158270.

[7] Trakadas, Panagiotis and Karkazis, Panagiotis and Leligou, Helen C. and Zahariadis, Theodore and Vicens, Felipe and Zurita, Arturo and Alemany, Pol and Soenen, Thomas and Parada, Carlos and Bonnet, Jose and Fotopoulou, Eleni and Zafeiropoulos, Anastasios and Kapassa, Evgenia and Touloupou, Marios and Kyriazis, Dimosthenis, "Comparison of management and orchestration solutions for the 5G era," *JOURNAL OF SENSOR AND ACTUATOR NETWORKS*, vol. 9, no. 1, p. 18, 2020. doi: http://dx.doi.org/10.3390/jsan9010004.

[8] L. Valcarenghi, B. Martini, K. Antevski, C. J. Bernardos, G. Landi, M. Capitani, J. Mangues-Bafalluy, R. Martinez, J. Baranda, I. Pascual, A. Ksentini, C. F. Chiasserini, F. Malandrino, X. Li, D. Andrushko, and K. Tomakh, "A Framework for Orchestration and Federation of 5G Services in a Multi-Domain Scenario," in *Proceedings of the Workshop on Experimentation and Measurements in 5G*, EM-5G'18, (New York, NY, USA), p. 19–24, Association for Computing Machinery, 2018. doi: http://doi.org/10.1145/3286680.3286684.

[9] B. Rusti, H. Stefanescu, M. Iordache, J. Ghenta, C. Brezeanu, and C. Patachia, "Deploying smart city components for 5g network slicing," in *2019 European Conference on Networks and Communications (EuCNC)*, pp. 149–154, 2019. doi: http://dx.doi.org/10.1109/EuCNC.2019.8802054.

[10] VITAL-5G, "Deliverable D1.2," December 2021. [Online] Available: https://www.vital5g.eu/wp-content/uploads/2022/01/VITAL5G-D1.2-5G-system-specifications-and-architecture_Final.pdf,.

[11] VITAL-5G, "Deliverable D2.3," February 2023. [Online] Available: https://www.vital5g.eu/public-deliverables-2/,.

[12] ETSI TS 123 502 V16.12.0 (2022-03), "5G Procedures for the 5G System (5GS) (3GPP TS 23.502 version 16.12.0 Release 16)," March 2022. [Online] Available: https://www.etsi.org/deliver/etsi_ts/123500_123599/123502/16.12.00_60/ts_123502v161200p.pdf, Last accessed on: February, 2023.

[13] VITAL-5G, "Deliverable D2.1," February 2023. [Online] Available: https://www.vital5g.eu/public-deliverables-2/,.