Blockchain-based Zero Touch Service Assurance in Cross-domain Network Slicing

Vasileios Theodorou*, Alexios Lekidis*, Theodoros Bozios*, Kalman Meth†,
Adriana Fernández-Fernández‡, James Taylor§, Pedro Diogo¶, Pedro Martins¶, Rasoul Behravesh‖
*Intracom Telecom, Greece, † IBM, Haifa, Israel, ‡ i2CAT Foundation, Barcelona, Spain,
§ Bartr Group, Birmingham , England, ¶ Ubiwhere, Aveiro, Portugal, ‖ FBK, Trento, Italy
Email:*{theovas, alekidis, tmpo}@intracom-telecom.com , †meth@il.ibm.com ,
‡adriana.fernandez@i2cat.net, §jamestaylor@bartr.world , ¶{pdiogo, pmartins}@ubiwhere.com, ‖rbehravesh@fbk.eu

***Abstract*—The inclusion of resource sharing schemes within Network Function Virtualization (NFV) ecosystems allows for optimized usage of 5G infrastructure and extended capabilities of network slicing services. In such environments, marketplaces are formed to facilitate the exchange of NFV services across administrative domains, which may, however, belong to untrusted and unreliable entities. In this work, we propose a novel zero-touch approach for cross-domain network slicing service assurance, using enterprise blockchain technologies and employing an AI-driven closed-loop automation architecture. Our approach is based on the lifecycle management of Service Level Agreements (SLAs) using smart contracts—from service negotiation to service binding, monitoring, reconfiguration and decommissioning. Our closed-loop architecture is materialized using Cloud-Native operational Data Lakes and allows to continuously monitor the status and health of exchanged services and to detect or predict SLA violations so that immediate mitigation actions are taken to ensure service continuity. The proposed approach is applied in real Content Distribution Networks scenarios within the European project 5GZORRO and our experimental results demonstrate the ability of our system to accurately predict changes in service demand and to timely respond with preventive scaling actions.**

***Keywords*—service level agreement, NFV, Slicing, Blockchain**

I. Introduction

Network Function Virtualization (NFV) technologies play a major role in the composition of network service bundles—called *Network Slices* in the terminology of the fifth generation of mobile networks (5G)—to support mission-critical 5G applications with stringent requirements for bandwidth, reliability, latency and end-device density. The inclusion of resource and service sharing schemes within NFV systems, allows for the formulation of flexible network slices that can efficiently and concurrently cover the needs of multiple and diverse 5G verticals. Specifically, such schemes allow extending the capabilities of network slices with seamlessly introduced available services from multiple providers, while optimally utilizing underlying infrastructures towards satisfying the versatile requirements of 5G applications.

A recently introduced enabler for multi-stakeholder shareability schemes is the blockchain-based Marketplace (e.g., in [1][2][3]), where resource offers and leasing transactions are managed in a traceable and consistent manner in the formed ecosystems. Nevertheless, the presence of resources and services from multiple administrative domains reduces substantially the trust levels in network slices and adds complexity to their composition, hindering the management of their Quality of Service (QoS). Moreover, the presence of untrusted domains results in security threats, which might lead to potential compromises of 5G resources. Hence, important benefits of 5G networks, such as reliability and service assurance for end-users, cannot be guaranteed a priori in such ecosystems. This might lead to significant delays in the adoption of 5G technologies, as well as in their evolution to enable more open network service markets.

Current approaches for network slice service assurance are based on a closed-loop architecture [4] that enables advanced monitoring and allows to realize the automation of service management. This direction is also embraced by standardization bodies, such as the ETSI Zero-touch network and Services Management (ZSM) Industry Specification Group (ISG) [5]. However, important challenges need to be addressed when applying this architecture in multi-stakeholder environments, such as the complexity of cross-domain and cross-layer coordination and the lack of mechanisms not only for monitoring, but also for preventing violations in the agreements between service providers and consumers.

In this paper, we propose a novel approach that leverages enterprise blockchain technologies and in particular Smart Contracts, coupled with Cloud-Native operational Data Lakes to provide a zero-touch solution for the automated service assurance of multi-domain network slices. Our solution focuses on mechanisms for real-time detection as well as prediction of service agreement violations that are likely to cause service degradation. Specifically, Service Level Agreements (SLAs) are embodied within Smart Contracts and violations are evaluated based on set thresholds on a combined view of application-layer and network-layer monitoring metrics, which are collected and processed in an operational Data Lake. To proactively mitigate the occurrence of violations, proper actions and workflows are triggered on real-time, such as network slice scaling over infrastructure resources available in the marketplace. We use Artificial Intelligence (AI) prediction models that are trained to forecast violations and we employ Smart Contracts both for the transparent establishment of service-layer relationships and for the programmatically defined enactment of corrective actuation workflows.

The main contributions of this work are (i) a solution for automating service assurance of multi-domain network slices through a Smart Contracts-based, SLA-driven closed-loop architecture, (ii) the definition of AI-driven SLA breach prediction and mitigation mechanisms that are realized as modular Cloud-Native services and (iii) a prototype validation through a CDN scenario of a commercial solution deployed on a large-scale 5G testbed.

The rest of the paper is organized as follows. Section II presents an overview of the related work on blockchain technologies and marketplaces for multi-domain network slices. Section III presents our proposed Smart Contract-based SLA

lifecycle management approach and Section IV focuses on the automation of our solution through a data-driven architecture. This architecture is further enriched with elasticity aspects that are covered in Section V through the use of Cloud-Native Data Lakes. As a Proof of Concept, the entire integrated approach is tested in Section VI on a CDN scenario and finally, Section VII provides conclusions and perspectives for future work.

II. Related Work

Recently there have been several works introducing the use of blockchain technologies as a means to enable multi-stakeholder, End-to-End (E2E) network slicing scenarios. For instance, Nour *et al.* [2] employ blockchain technologies and introduce an entity named *Slice Provider*, for brokering and negotiating between verticals and network resources providers. Pursuing the same objective, Togou *et al.* [6], presents a method for distributing slice requests through a blockchain-based platform, to optimally bid for resources when slicing requirements cannot be fulfilled locally. Regarding radio access network resources, a secure and trusted blockchain-based spectrum-sharing mechanism is proposed by Gorla *et al.* [7] to minimize spectrum under-utilization in multi-operator scenarios. Similarly, in a unified view of multi-domain, multi-stakeholder resources (e.g., Virtualized Network Functionss (VNFs) and slicing services), Swapna *et al.* [3] adopts the concept of *"Slice as a Service"*, enforcing SLAs with the use of Smart Contracts and monitoring brokers.

When it comes to the distributed Management and Orchestration (MANO) of multi-stakeholder network services, Bondan *et al.* [8] propose a marketplace that addresses the lifecycle management of multi-vendor VNFs. In their approach, VNFs can be orchestrated in Service Function Chainings (SFCs) by a unified platform supporting multiple Virtualization Infrastructure Managements (VIMs). In the same direction, Rathi *et al.* [9] propose a blockchain-enabled scheme for the trustworthy coordination of multi-domain edge orchestrators that can provide E2E slices using resources acquired from multiple operators. Guerzoni *et al.* [10] show how an E2E MANO plane for 5G networks can be achieved using multi-domain awareness. Through a flexible, scalable and adaptable placement of VNFs, automated management of resources and slices is performed to fulfil multi-domain requirements. Also aiming at tackling scalability issues, Dieye *et al.* [11] propose a distributed multi-agent deep reinforcement learning approach and present a market-driven auction-based strategy for multi-party resource-sharing scenario.

Regarding QoS and SLA assurance of network slicing based systems, Papageorgiou *et al.* [12] propose a new SLA management solution that can perform SLA verification while dealing with higher system complexity of network slicing orchestration and virtualization layers. In the same direction, Xie *et al.* [4], extends the previous works in MANO, network slicing, and unified marketplaces by proposing an AI-based closed-loop service assurance architecture that can enhance both the network operation, resource sharing, and overall virtualized service management. Regarding the satisfaction of QoS requirements for cross-domain network slices, Theodorou *et al.* [13] propose a new network slicing mechanism optimized for concurrent support of multiple application domains with differential functional and non-functional requirements.

To our knowledge, our work is the first one to address AI-driven service assurance of complex multi-domain and multi-layer network slices in a holistic view. One important innovation that we introduce is the employment of blockchain technologies not only for the transparency and traceability of interactions in network service marketplaces, but also for the secure, proactive and streamlined MANO-based mitigation of SLA violations.

III. SLA Management using Smart Contracts

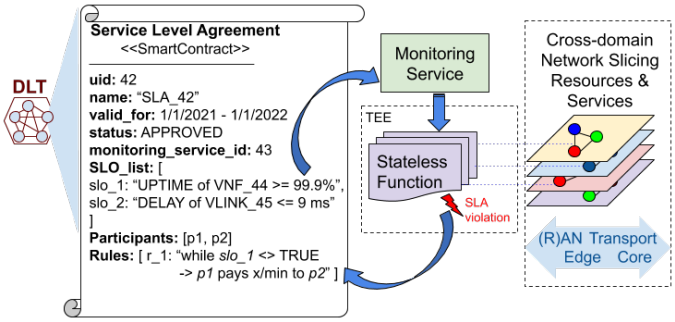

Fig. 1: SLA lifecycle management with DLT Smart Contracts

In the context of 5G mobile networks, SLA management automation is a key feature to deliver the promised levels of dynamic adaptation and network performance. As documented in 3GPP specifications [14], network operators provision customized network slice instances as E2E logical networks over common virtualized infrastructures to provide a customer's service request. Based on the specific requirements of the served request (e.g., throughput, latency, coverage), an SLA is associated to the slice instance, as an agreement between the operator and the customer, which specifies the slice's expected performance. Similarly, in case of NFV-based marketplaces, an SLA can be associated with a slice comprising of cross-domain resources, capturing the corresponding more complex agreements that are necessary to support the exchange of resources and services in such sharing environments.

Crucial to the realization of SLA management automation is the encapsulation of the machine-readable technical SLA representation with its governing rules, signed by all involved parties. To this end, we notice the direct mapping of such processes to primitives provided by Distributed Ledger Technologies (DLT) and in particular, we take advantage of the expressiveness and automation provided by Smart Contracts in enterprise blockchains (CORDA, Hyperledger etc. [15]). Hence, we employ Smart Contracts for the management of the lifecycle of SLAs, which are semi-automatically constructed using pre-defined templates. As shown in the example of Fig. 1, an SLA concisely reflects, in addition to generic contractual information (participants, validity period etc.), the committed QoS in terms of one or several Service-Level Objectives (SLOs). Given the virtualized nature of 5G environments, for a given active slice, these high-level (QoS-related) objectives must be tailored to low-level technical parameters, which more clearly identify the specific metrics and targets to be monitored from the considered NFV Infrastructure (NFVI), as well as the trusted monitoring services to provide such metrics. Complementing this information, reported status, incurred violations and agreed penalties (in form of rules) are also common attributes conforming SLA models [12].

The introduction of DLT-based Smart Contracts ensures that SLAs and the parties and obligations they encompass are automated in their lifecycle and guarantee consistency, trust and non-repudiation. Service providers compose SLA instances in a template-based manner, reflecting the terms they are willing to commit to for given resources. The machine readable SLA representation, combined with the human-readable legal prose that carries the semantic meaning of the contractual agreement and the signatures of all parties named on the agreement, conclude the definition of what is called a *Ricardian Contract*. By taking a template-driven approach that encapsulates the principles of a Ricardian Contract, we afford marketplace stakeholders the flexibility to define SLA agreement templates for classifications of resources and associated metrics to be measured. Marketplace templates (an example is depicted in Fig. 2) comprise the parameterized legal prose, backed by the technical model and, in our case, extended to incorporate programmable business logic, thus enabling the definition of truly smart SLAs. Templates are subject to a governance process and are globally resolvable thanks to the use of decentralized identifiers, before becoming available in the marketplace for a stakeholder to build commercial agreements around their product offers. Global collaborative repositories of available and reusable templates are formed incrementally along with these processes, covering versatile use cases and network slicing requirements. This collaborative approach gives rise to increased transparency and use-case coverage of contracts with obvious efficiency gains thanks to the reusable nature of SLA specifications within the marketplace ecosystem.

When it comes to monitoring services, they are prescribed from a repository of known services to specify the services that should be alerted of any changes to SLA definitions. Monitoring services are granted authorization to transact with the DLT through the issuance of verifiable credentials that needs to be present when submitting violation transactions. SLAs are associated with product offers by providers, before being published to a decentralized catalogue for other marketplace stakeholders to discover. A product order can then be signed between a provider and a consumer.

Once committed, the SLA and any interactions or events pertaining to it are governed solely by the Smart Contract. When the SLA becomes active, management of monitoring infrastructure can be automated according to the needs of the SLA measurements to be taken. A management entity in the service provider's domain, which we name *SLA Management*, takes care of such SLA lifecycle events, publishing SLA status updates to a common communication bus where monitoring services can subscribe, which we call the *Integration Fabric*. On consuming events, monitoring services handle the set-up and tear-down of monitoring and analytics pipelines, realized as stateless functions. One facet of the set-up of these pipelines is the deployment of any business logic encapsulated in the legal prose within a Trusted Execution Environment (TEE), such that it can be executed securely to determine the violation of SLAs. By executing this logic off-chain within a trusted environment, we achieve both a trusted outcome about the service health and an efficient use of DLT whereby only violations and state-changing logic are submitted to the ledger. After their determination in TEEs, violation events are signed by monitoring services and submitted to the SLA Management.

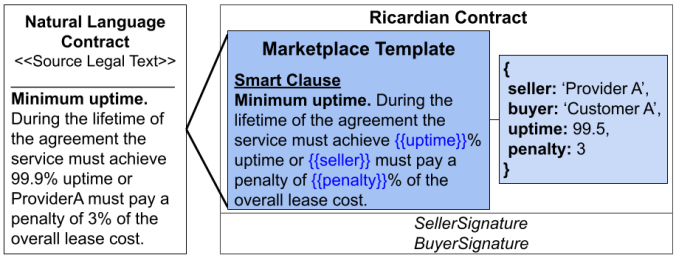

Fig. 2: Template-based generation of Ricardian Contracts

IV. Architecture for Zero-Touch Automation

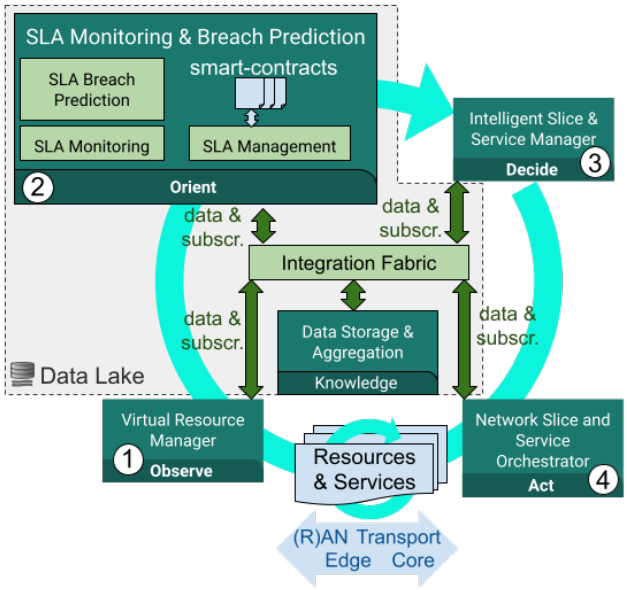

Fig. 3: Closed Loop SLA Assurance Architecture

In this section, we introduce our novel 5GZORRO architecture for service assurance, that follows the AI for IT Operations (AIOps) paradigm in a feedback-loop fashion. AIOps [16] exploits AI/Machine Learning (ML) techniques towards IT operations automation. This is achieved by constantly collecting, analyzing and correlating data across large and complex software systems and by providing realtime, actionable insights over system behaviours, as well as recommendations and corrective actions. We adopt an AIOps view that is entangled with a Closed Loop Architecture inspired by ETSI ZSM [5], so that the network slicing resources and services are managed in a "zero-touch" fashion. This feedback-loop-based view allows for high degrees of automation and intelligent service assurance, to respond to the high complexity of cross-domain and multi-layer network slicing resources and services.

Fig. 3 summarizes the AIOps cycle we employ for the data-driven automated management of SLA lifecycle. A feedback loop is logically composed of the following phases: (1) "observation", i.e., the gathering of monitoring data about managed multi-domain resources and services as provided by *Virtual Resource Managers*—the entities responsible for interfacing with the MANO layer(s) of each domain; (2) "orientation", i.e., aggregation and contextualization of the

raw data into SLA monitoring metrics by Monitoring and Breach Prediction components that analyze the information and produce notifications and alerts; (3) "decision-making", where intelligent agents process insights and notifications, assess the current operational status and decide how to handle detected or forecasted issues and anomalies; (4) and, finally, the "action" phase, where MANO elements for *Network Slice and Service Orchestration* are invoked to act upon the managed resources and services, to actuate the decision. Optionally, lower-level closed loops can take place within domains, as part of intra-domain orchestration.

All communication between components takes place via event-based, asynchronous interactions through the *Integration Fabric*, which exposes data publication and subscription services, as well as services for the installation of data pipelines. The latter are executed in the analytics and processing engine of the *Data Storage and Aggregation* component, which is also responsible for persisting data, acting as the Knowledge Base of the feedback loop. We refer to the combined services of data storage and aggregation, analytics and processing engine and notification services as a *Data Lake*, further described in Section V.

Delving deeper into the 5GZORRO SLA Monitoring and Breach Prediction components, the SLA Breach Prediction collects and analyses resource monitoring data and uses AI/ML techniques to 1) predict possible SLA violations as well as 2) provide timely awareness over service quality degradation, resource utilization changes and system misconfigurations. The monitoring data are supplied to the SLA Breach Prediction by the SLA Monitoring component, which collects raw data from the various *Virtual Resource Managers* of different domains, and translates (i.e., transforms and aggregates) them to SLO metrics of interest. This translation is governed by the *SLA Management* component, which follows the information contained within the Smart Contracts corresponding to SLAs, in the process that was described above in Sec. III. The sensitive nature of SLA Monitoring and Breach Prediction operations demands the usage of a tamper-free environment, where the data and computations cannot be altered. In such a multi-domain system, no inherent trust mechanisms are established and zero-trust hardware platforms, such as TEE, must be employed to establish a secure root-of-trust.

Upon the prediction of SLA violations, corrective mechanisms are proactively taken to maintain the desired QoS for offered services. Such mechanisms are linked to the amendment of network slices with resources from further administrative domains by using the *Intelligent Slice and Service Manager* (ISSM) module. Specifically, ISSM is a workflow manager that applies smart discovery mechanisms to identify the most prominent resource offerings, purchase them and then perform the adequate configurations for meeting the service SLAs. Finally, the slice amendment with the newly added resources is handled by the *Network Slice and Service Orchestration* module, which is responsible for the appropriate slice reconfigurations, such as slice scaling.

V. Cloud Native Data Lakes for Elastic SLA Monitoring

Operational Data Lake is one of the major enablers for the data-driven zero-touch automation in 5GZORRO. As depicted in Fig. 3, Data Lake plays a central role in the SLA lifecycle management loop, providing data services, notification services, and the event-driven platform for custom data processing and analytics pipelines. The Data Lake is designed in a Cloud-Native way, with a container orchestration engine, message queue, object storage and server-less workflow engine, to enable elastic allocation of computational resources to SLA management components and their data processing workflows. This elasticity feature is essential for the pragmatic management of complex multi-domain and multi-stakeholder environments, allowing for management components to dynamically scale and follow the heterogeneity and un-predicted evolution of the resources and services they monitor.

Resource Providers monitor the resources they manage and feed operational data to the Data Lake, through a resource metrics ingestion service. Service Providers lease these resources to create services and must fulfill terms of SLAs agreed upon with their customers. Upon service instantiation, 5GZORRO SLA Management services match up the metrics of each individual resource (from different providers) to the SLA that is impacted by that individual resource both in real-time and offline. This results in deploying custom data processing pipelines through the Data Lake exposed API, to monitor the service SLA continuously and to invoke the analytical, decision making, and actuating functions asynchronously in response to service lifecycle events and observed service health and performance.

Monitored resources and SLAs are assigned with unique identifiers which are used as keys in a platform-wide catalog within the Data Storage and Aggregation. The catalog is essentially a database that indexes all the information required for SLA lifecycle management and provides services to ingest, lookup, update and delete related data. Periodic monitoring reports sent by Resource Providers to Data Lake metrics ingestion service are parsed and forwarded to data catalog ingest service, that indexes the data to make it searchable on the resource unique identifier, timestamp, and data type fields. SLA Monitoring components use catalog services to manage the SLA data and to lookup the data for resources that the monitored SLAs depend upon.

VI. Evaluation

In order to evaluate our approach, we implemented a prototype of the SLA closed-loop architecture based on open-source frameworks and technologies. As a use case, we employed real application scenarios from the 5GZORRO project and specifically horizontal scaling scenarios of the fs|CDN® Anywhere[1] commercial CDN application, deployed as VNFs over 5G infrastructures—henceforth referred to as *vCDN*. As a blockchain technology for our prototype solution, we used the Corda[2] platform due to its full-featured smart contract support and its scalability and privacy benefits stemming from its inclusive consensus mechanism (i.e., only participating parties in a transaction have visibility over it and are involved in its consensus process).

In Fig. 4, we depict the System Architecture of our prototyping solution, set up on a testbed of the 5GZORRO Spanish Facility[3]. For the Network Service Management and Orchestration layers we follow the typical NFV MANO architecture, compatible with the ETSI NFV MANO framework. An

[1] fs|CDN® Anywhere Solution: http://www.intracom-telecom.com/en/-products/telco_software/iptv_multiplay/fs_cdn.htm

[2] Corda Blockchain Platform: https://www.corda.net/

[3] 5G Barcelona Open Digital Hub: https://5gbarcelona.org/

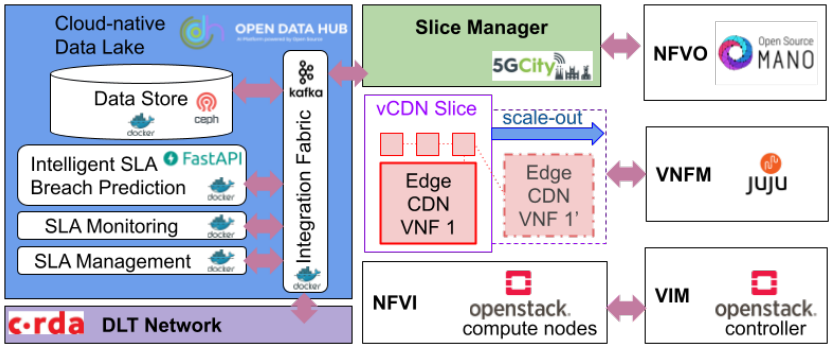

Fig. 4: System Architecture of prototype solution at the 5GZORRO Spanish Facility

Open-Source MANO (OSM) (release 8) installation undertakes the role of the NFV Orchestrator (NFVO) with Juju charms acting as the VNF Manager (VNFM) responsible for Day0/1/2 primitives over VNFs. An Openstack (version Ussuri) cluster is installed at the testbed for managing virtualized infrastructures. We employ a Slice Manager component developed within the H2020 Project 5G-City for interoperating with OSM and managing the lifecycle and horizontal scaling of Network Slices. The Slice Manager is integrated with the Cloud-Native Data Lake components (on the left side of Fig. 4) through a Kafka-based Integration Fabric. The Integration Fabric, as well as all components within the Cloud-Native Data Lake, are realized as Docker containers inside a Kubernetes cluster, following the Open Data Hub[4] architecture for elastic deployment. As a massively scalable Data Store we use a Ceph[5] cluster and for the Intelligent SLA Breach Prediction, we implemented a highly configurable ML and Analytics Platform using the FastAPI Framework.

The SLA Monitoring and SLA Management components are implemented as Python-based and Java-based micro-services, respectively. The SLA Management acts as a client node to the DLT network of our system. To this end, we implemented a Corda Distributed Application (CorDapp) consisting of Corda elements called *states* for modeling the facts of the ledger; *contracts* for defining acceptable State changes; and *flows* for embodying the processing logic for the transactions and consensus mechanisms. More specifically, we implemented the SLA as a *state* with the properties depicted in Fig. 1, accompanied by a corresponding *contract* with verification methods for the issuance of valid SLAs. Furthermore, we implemented *flows* for the agreement and termination of SLAs, as well as for the registration of monitoring events to the ledger. We followed a similar implementation rationale for the registration of resource and service offers to the formed marketplace of participant Corda nodes.

We conducted two sets of experiments to assess the feasibility and effectiveness of our approach. In the first set of experiments, we focused on the reaction time of our E2E service assurance mechanisms. We deployed a Network Slice of the vCDN application over the testbed NFVI and calculated the duration from the time that monitoring data are received by the Intelligent SLA Breach Prediction, all the way until it (A) predicts that excessive demand necessitates scaling action and sends a scaling request to the Slice Manager, (B) a transaction (*Tx*) commences on the blockchain for the issuance of an SLA between the network slice consumer and the infrastructure resource provider, and (C) the Slice Manager completes the Network Slice scaling with the addition and activation of a new Edge vCDN VNF. We used a VNF flavor of 8 CPU cores, 16GB RAM and 2TB HDD, which corresponds to operational requirements of the Edge vCDN component in production. Predictions were conducted on a container with 16G RAM and 4 cores of Intel Core i5 1.7GHz CPU. We ran 10 Corda nodes, but we note that all nodes were on the same local area network, disregarding potential networking delay of real-life scenarios. Average results from 50 runs for the time duration of the above three SLA assurance phases—correspondingly (A) *Prediction*, (B) *SLA Tx* and (C) *Slice scaling*—are shown in the Pie Chart of Fig. 5, where we show the percentage of each phase duration contributing to the E2E process. Slice scaling appears as the main bottleneck of this process (96.8% of total duration), due to Virtual Machine spin-up times. Individual predictions need roughly a second and DLT transactions some hundred milliseconds, which validates that our mechanisms take only a fraction of the time necessary for MANO activities.

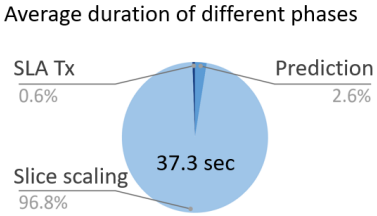

Fig. 5: Time duration of E2E SLA assurance phases

In the second set of experiments, we targeted the effectiveness and performance of our AI-driven intelligence mechanisms. To this end, we collected datapoints of per-minute observed bandwidth from running vCDN servers in production and we evaluated the performance of the Intelligent SLA Breach Prediction component. The time step of our monitoring data is in line with the reaction time captured in the first set of experiments, providing a time-wise validation of the compliance between prediction and reaction mechanisms. We used different methods to predict bandwidth in the near future, so as to trigger scaling actions when predictions exceed set thresholds. In Fig. 6 and Fig. 7 we show results from different

[4]Open Data Hub Platform for the Hybrid Cloud: http://opendatahub.io/
[5]Ceph Software Storage Platform: https://ceph.io/

configurations of LSTM and ARIMA on a training set of 11,000 and a test set of 3,500 consecutive observed values. In Fig. 6, we contrast the prediction of a (S)ARIMA(X) model with a multi-step LSTM model that receives 60 observed values as input and predicts the next 30. We refer to the latter LSTM model as *LSTM-60-30* and we define other LSTM models correspondingly (*LSTM-3-1*, *LSTM-20-10* etc.). Both models are trained on the complete training set and as we can see, the SARIMAX model appears to more accurately predict near-future values. From a Mean Absolute Prediction Error analysis (Fig. 7) we see that LSTM predictions have an error of less than 9% even for prediction of values 10 steps ahead. Additionally, we measured that training time for different LSTM variations range from 3 minutes for LSTM-3-1, to 18 minutes for LSTM-60-30 and 30 minutes for LSTM-100-1. We note here that due to the high configurability of our SLA Breach Prediction component, we were able to switch between algorithms, hyperparameter tunings, input/output datasets etc, simply with the parameterization of HTTP REST requests.

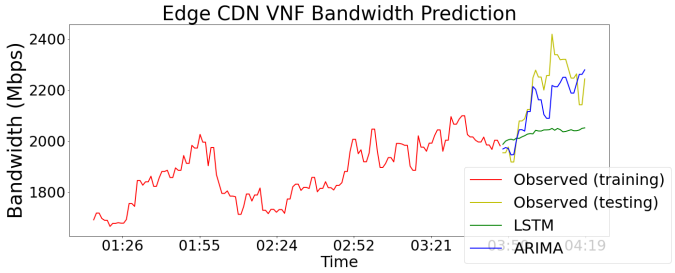

Fig. 6: Forecasts from Intelligent SLA Breach Prediction

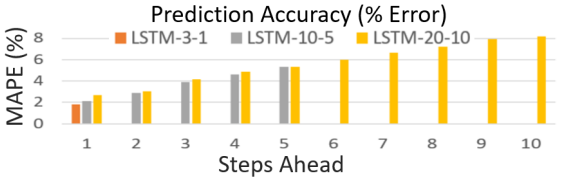

Fig. 7: Prediction accuracy of different forecasting algorithms

VII. Conclusion

This paper proposes a novel approach for providing service assurance in cross-domain network slices that are exchanging resources and services through network service marketplaces. The approach is based on a closed-loop architecture for lifecycle management of SLAs that are expressed through blockchain-based technologies and in particular smart contracts. To this end, lifecycle management includes a variety of mechanisms, including service binding, status monitoring, AI-driven detection and prediction of service violations as well as mitigation actions upon these violations. These mechanisms are 1) automated through the use of machine learning techniques and 2) highly elastic through the employment of Cloud-Native Data Lake technologies. The approach is applied for the resolution of real multi-domain network slicing challenges related to the deployment of CDN scenarios in 5GZORRO testbeds. The conducted experiments illustrate the feasibility and performance of our approach in complex environments, regarding the accuracy of short-term predictions as well as the automated corrective actions that are scheduled for ensuring the required CDN QoS.

As future work, we plan to investigate further machine learning mechanisms for longer-term predictions, designed and implemented for distributed deploy-ability on edge locations.

Acknowledgment

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 871533 (5GZORRO) and from CERCA Programme / Generalitat de Catalunya.

References

[1] J. Backman, S. Yrjölä, K. Valtanen, and O. Mämmelä, "Blockchain network slice broker in 5G: Slice leasing in factory of the future use case," in *IEEE IoT Business Models, Users, and Netw.*, 2017, pp. 1–8.

[2] B. Nour, A. Ksentini, N. Herbaut, P. A. Frangoudis, and H. Moungla, "A blockchain-based network slice broker for 5G services," *IEEE Networking Letters*, vol. 1, no. 3, pp. 99–102, 2019.

[3] A. Swapna, R. Rosa, C. Esteve Rothenberg, I. Sakellariou, L. Mamatas, and P. Papadimitriou, "Towards A Marketplace for Multi-domain Cloud Network Slicing: Use Cases," 09 2019, pp. 1–4.

[4] M. Xie, J. S. Pujol-Roig, F. Michelinakis, T. Dreibholz, C. Guerrero, A. Gallego Sanchez, W. Yi Poe, Y. Wang, and A. M. Elmokashfi, "AI-Driven Closed-Loop Service Assurance with Service Exposures," in *EuCNC 2020*. IEEE, 2020, pp. 265–270.

[5] C. Benzaid and T. Taleb, "AI-driven zero touch network and service management in 5G and beyond: Challenges and research directions," *IEEE Network*, vol. 34, no. 2, pp. 186–194, 2020.

[6] M. A. Togou, T. Bi, K. Dev, K. McDonnell, A. Milenovic, H. Tewari, and G.-M. Muntean, "DBNS: A Distributed Blockchain-Enabled Network Slicing Framework for 5G Networks," *IEEE Communications Magazine*, vol. 58, no. 11, pp. 90–96, 2020.

[7] P. Gorla, V. Chamola, V. Hassija, and N. Ansari, "Blockchain based framework for modeling and evaluating 5g spectrum sharing," *IEEE Network*, 2020.

[8] L. Bondan, M. F. Franco, L. Marcuzzo, G. Venancio, R. L. Santos, R. J. Pfitscher, E. J. Scheid, B. Stiller, F. De Turck, E. P. Duarte, A. E. Schaeffer-Filho, C. R. P. d. Santos, and L. Z. Granville, "FENDE: Marketplace-Based Distribution, Execution, and Life Cycle Management of VNFs," *IEEE Comm. Magaz.*, vol. 57, no. 1, pp. 13–19, 2019.

[9] V. K. Rathi, V. Chaudhary, N. K. Rajput, B. Ahuja, A. K. Jaiswal, D. Gupta, M. Elhoseny, and M. Hammoudeh, "A Blockchain-Enabled Multi Domain Edge Computing Orchestrator," *IEEE Internet of Things Magazine*, vol. 3, no. 2, pp. 30–36, 2020.

[10] R. Guerzoni et al., "Analysis of end-to-end multi-domain management and orchestration frameworks for software defined infrastructures: An architectural survey," *Transactions on Emerging Telecommunications Technologies*, vol. 28, 2016.

[11] M. Dieye, W. Jaafar, H. Elbiaze, and R. Glitho, "Market Driven Multi-domain Network Service Orchestration in 5G Networks," *IEEE Journal on Selected Areas in Communications*, 2020.

[12] A. Papageorgiou, A. Fernández-Fernández, L. Ochoa-Aday, M. S. Pelaez, and M. S. Siddiqui, "SLA Management Procedures in 5G Slicing-based Systems," in *EuCNC 2020*. IEEE, 2020, pp. 7–11.

[13] V. Theodorou, K. V. Katsaros, A. Roos, E. Sakic, and V. Kulkarni, "Cross-Domain Network Slicing for Industrial Applications," in *EuCNC 2018*. IEEE, 2018, pp. 209–213.

[14] 3GPP, "Study on Management and Orchestration of Network Slicing for Next Generation Network," 3rd Generation Partnership Project (3GPP), Technical Report (TR) 28.801, Jan. 2018, version 15.1.0.

[15] M. Swan, "Chapter Five - Blockchain for Business: Next-Generation Enterprise Artificial Intelligence Systems," *Adv. Comput.*, vol. 111, pp. 121–162, 2018.

[16] A. Masood and A. Hashmi, "AIOps: Predictive Analytics and Machine Learning in Operations," in *Cognitive Computing Recipes: Artificial Intelligence Solutions Using Microsoft Cognitive Services and TensorFlow*, 2019, pp. 359–382.