ACROSS: Automated zero-touch cross-layer provisioning framework for 5G and beyond vertical services

Dimitris Giannopoulos*, Georgios Katsikas†, Kostis Trantzas*, Dimitrios Klonidis†, Christos Tranoris*, Spyros Denazis*, Lluis Gifre‡, Ricard Vilalta‡, Pol Alemany‡, Raul Muñoz‡, Anne-Marie Bosneag§, Alberto Mozo¶, Amit Karamchandani¶, Luis de la Cal¶, Diego R. López‖, Antonio Pastor‖, Ángela Burgaleta‖

*University of Patras, †UBITECH, ‡Centre Tecnològic de Telecomunicacions de Catalunya (CTTC/CERCA) (Spain), §Ericsson Ireland, ¶Universidad Politécnica de Madrid (Spain), ‖Telefónica I+D (Spain)

***Abstract*—As the demand for advanced and efficient network and service deployment continues to rise, the integration of multiple domains and the incorporation of AI technology are becoming essential. The ACROSS project is a Horizon Europe project, that aims to address this need by proposing an innovative end-to-end service deployment and management platform. This platform is designed to deliver unprecedented levels of automation, performance, scalability, and energy efficiency in the next-gen networks and services landscape. The platform will be built as a highly-distributed grid of domain-level orchestrators, spread across multiple geo-dispersed and potentially heterogeneous edge environments, all overseen by a cloud-managed multi-domain orchestrator. The use of standardised communication interfaces will promote separation of concerns and ensure compliance with ongoing standardization efforts, some of which include ETSI ZSM, ETSI NFV-OSM, TMF, ETSI TFS and ONF. The platform will also be enhanced with deep end-to-end telemetry, AI-driven intelligence, full-stack cross-domain zero-touch provisioning, and secure and trusted orchestration mechanisms.**

***Index Terms*—5G, 6G, service orchestration, zero-touch provisioning, AI, telemetry, trust and security**

I. Introduction

The growth of the Internet of Things (IoT) has driven innovation in both cloud computing and networking, leading to a rearchitecting of modern systems across the cloud continuum, from large hyperscalers to the network edge, with advancements in edge computing and 5G networks. Next generation 6G networks and Artificial Intelligence (AI) are poised to bring further transformations to industries through edge intelligence. Recent innovations in key domains, including the widespread adoption of IoT for data collection, powerful distributed computational power through Graphics Processing Units (GPUs) and programmable edge processors, and the maturity of AI with deep learning's ability to understand unstructured information, have brought practical edge intelligence closer than ever.

However, the growth and change in technology has brought both challenges and opportunities for organizations. Interoperability and integration are complicated due to a complex, multi-technology, multi-vendor and multi-tenant infrastructure managed by conflicting stakeholders. Hybrid compute environments also slow down development and adoption. The 6G network requires a Management & Orchestration (MANO) system capable of handling a huge number of events from numerous managed devices, but limited data plane visibility and a lack of AI in current systems hinder its development. Automation is still limited and the increase in data and events from 6G networks presents security and trust challenges.

The ACROSS project [1] aims to provide a solution for effective management in next-gen networks and services through a comprehensive architecture. It features a distributed grid of domain-level orchestrators managed by a cloud-based multi-domain orchestrator, with standard communication interfaces for compliance. The architecture includes AI intelligence, telemetry, zero-touch provisioning, and secure orchestration. The goal is to unlock potential in IoT and AI domains and drive the future of next-gen networks and services.

II. Section-related background

A. Ultra-scalable cross-domain service deployment and orchestration

Orchestration of application lifecycle mainly in cloud environments involves the use of modern service and network management platforms with tiered architecture across multiple layers. The layers are managed by different administrative entities, such as cloud infrastructure providers using Virtual Infrastructure Manager (VIM) software (e.g., OpenStack [2], Kubernetes [3]), and Software-Defined Network (SDN) controllers (e.g., ONOS [4], ODL [5], ETSI TeraFlowSDN (TFS) [6]) for server and network provisioning. Network Function Virtualization (NFV) Management and Orchestration platforms (e.g., ETSI OSM [7] or ONAP [8]) and overlay operations support system (OSS) may control multiple VIMs, SDN controllers, and application-level orchestrators with the aim to perform optimal service deployments. Despite standardization efforts by global bodies and consortia (i.e., ETSI, IETF, ONF, the Linux Foundation), end-to-end orchestration remains challenging due to limited orchestration APIs, single-objective orchestration loops, the need for AI-powered changes and closed loops, and scalability issues for future 6G platforms. Despite the great advancements, today's state-of-the-art centralized orchestration frameworks (ONF's Aether [9], OSM, and ONAP) still face these challenges.

B. Deep end-to-end telemetry on open programmable infrastructures

Monitoring in infrastructure refers to collecting and analyzing data to assess the state, detect potential incidents and prevent outages. Traditionally, infrastructure monitoring was complex and limited due to vendor-specific solutions and a lack of integration between network and compute resource monitoring. The advancements in cloud computing and the convergence towards software-defined principles have improved monitoring solutions, leading to the development of modern monitoring tools, such as Netdata [10], Prometheus [11], Sematext [12], Elastic Stack [13], and Zabbix [14], which provide a common view of compute and network resources. However, today's infrastructure monitoring systems face performance, accuracy, and generality challenges due to software-to-hardware interactions and a lack of visibility into microbursts and virtualized environments. For example, tools require deeper visibility into traffic, as packets may move from source to destination without ever leaving the server.

C. Real-time analytics and AI for next-generation end-to-end orchestration

The heterogeneity and diversity of beyond 5G mobile networks and the need for multi-tenant support pose challenges for network management, requiring automated solutions. AI-based methods are being explored to provide support for intelligent self-management and self-organization. 3rd Generation Partnership Project (3GPP) and other Standards Development Organizations (SDOs) are developing standardized architectures and protocols integrating AI, and standardizing data models and data management solutions to enable data collection and analysis to train AI models. New modules have been added to 5G core network architecture [15] to provide this capability. ETSI Experiential Networking Intelligence (ENI) is also incorporating data ingestion and processing components to facilitate data collection and analysis to build AI models. Moreover, the flexibility and programmability of 5G allow for creation of network segments with different Quality of Service (QoS) requirements, requiring new methods of network segment management and orchestration. Network Resource Elasticity (NRE) has emerged as a concept to dynamically adjust network resources and configuration to meet QoS and reduce operational costs (OPEX) [16]. NRE can be divided into three elements: (1) flexibility of computational resources for VNFs, (2) elasticity achieved through optimized VNF placement by orchestration, and (3) resource elasticity between slices. Combining these three elements provides a comprehensive solution for network resource management to meet the varying QoS requirements of various services. AI methods can play a key role in NRE by identifying patterns in data to predict trends and optimize resource allocation. AI algorithms can be applied in various phases of network slicing lifecycle. A capacity forecasting framework is proposed in [17] to dynamically adapt to traffic variations and reduce management costs. The anticipated increase in data traffic and usage may cause significant congestion, requiring advanced Congestion Control (CC) mechanisms. Traditional methods are not suitable for 5G's scale and dynamic nature, and new learning-based techniques, such as Machine Learning (ML) and Deep Learning (DL), are proposed to build generic CC schemes that dynamically learn and identify optimal actions. In addition, Reinforcement Learning (RL) agents can continuously monitor the network's state and react to the environment, applying a sequence of actions to minimize the congestion of the network to gain feedback in this interaction, which can be reused in its subsequent actions. For this reason, RL algorithms are more suitable for dynamic network environments as they can adjust their learning to remain robust to the variability or instability of the network over time. However, RL algorithms require a longer training time to converge to a near-optimal solution and are computationally more intensive during inference, which may limit their applicability in real-time scenarios [17]. Moreover, RL algorithms may not recommend safe options in a live network, which is why SafeRL options have been devised [18], [19].

D. Full-Stack cross-domain zero-touch provisioning (ZTP)

Automation has been a long-standing goal in telecom networks [20], aimed at improving operator efficiency and speed in network management. Currently, automation is limited to network device provisioning and firmware updates [21] or to limited domains such as SON. The emergence of edge computing and 5G has added hardware heterogeneity, requiring generic automation mechanisms [22]. Despite the term "zero-touch provisioning," most implementations are not truly zero-touch and often require manual steps [23]. SDN has drawn even more attention to automation as a key procedure [24], [25], but the networking community struggles to keep up with the rapid advancements in cloud computing, resulting in increased network programmability but also complexity. End-to-end automation of service and network management is crucial for 5G and beyond deployments, considering the extreme range of requirements imposed, such as ultra-low predictable latency, high throughput, dramatically improved customer-experience, and support for massive machine-to-machine communication, as well as the diversity of offered services, but current orchestration platforms fall short in fully addressing automation triggers and lack appropriate AI methods. The platform aims to introduce zero-touch operations in beyond 5G systems by addressing these issues.

E. Secure & trusted orchestration mechanisms

Secure and trusted orchestration mechanisms are addressed from 3 aspects: Trusted Computing, Safe AI, and Software-Defined Security.

Trusted Computing uses mechanisms, such as hardware-based enclaves, i.e., containers with dedicated memory regions, secured with on-chip encryption, allowing critical parts of applications to run in a secure environment. Remote attestation verifies the trusted environment's software and identity. Intel's Software Guard Extensions (SGX) platform supports sealing for state persistence across reboots, but rollback and forking attacks are challenges.

Safe AI is vulnerable to adversarial attacks such as evasion and poisoning that undermine system integrity. The research community is active in developing defensive techniques to make AI systems robust to these attacks [26]. The problem of cloud-edge based large-scale data processing deployments using trusted computing technologies has also only been partially addressed [27], [28]. In the context of Safe AI, advances in ML have enabled various applications like data analytics and digital forensics. Deep Neural Networks (DNNs) may exhibit high prediction accuracy, but, in reality, they may be vulnerable to adversarial attacks, which can impact system integrity and privacy [29], [30]. Adversarial samples are inputs crafted to fool a trained classifier, while poisoning attacks manipulate the training data or algorithm to "poison" the model [31], [32]. These vulnerabilities do not reflect a design flaw in ML algorithms but arise from deployment in adversarial scenarios. Defenses against evasion attacks have been proposed (such as defensive distillation [33], input transformation [34], and randomization [35]), but none are robust to adaptive attacks [36].

Meanwhile, Software-Defined Security (SD-SEC) [37] is a mechanism for protecting network slices against attacks [38] through Security-as-a-Service (SECaaS). SD-SEC brings security network functions such as Identity and Access Management, Network Segmentation, and Cryptographic services to the NFV infrastructure while enabling centralized software management, policy-driven orchestration, and real-time monitoring [39]. It inherits NFV's flexibility and scalability, and enables sophisticated security-aware automation workflows.

III. The ACROSS approach

The proposed architecture aims to be a scalable, end-to-end service orchestration platform consisting of domain-level orchestrators in a distributed grid, across heterogeneous edge environments, managed by a cloud-based multi-domain orchestrator. To achieve this, the proposed architecture will: (i) include a standardised cross-domain integration fabric with open, standard APIs for technology and vendor agnosticity, (ii) decouple today's orchestration loops from their objective (e.g., performance or scalability), thus introducing programmable orchestration objectives, including energy efficiency and cost reduction, (iii) use AI to turn telemetry into predicted actions as smart orchestration tasks, and (iv) differentiate the roles of domain and multi-domain orchestrators for scalable domain-level orchestration with end-to-end orchestration assistance.

The architecture aims to create an extensive in-band telemetry system that covers the entire processing range, from various network devices (e.g., 5G cells, switches, routers) to bare metal and hypervisor servers. Unlike current costly out-of-band monitoring systems, the proposed platform utilizes recent advancements in programmable networking hardware to provide high-performance and efficient packet-level visibility through hardware-offloaded in-band telemetry. The platform's telemetry system is: (i) scalable and efficient, as it realizes monitoring and reporting at nano-to-microsecond granularity and multi-hundred Gbps line-rate with zero CPU involvement, (ii) open and programmable, as it employs state-of-the-art technologies like P4, eBPF, OVS, Kubernetes/OpenStack, and utilizes open monitoring specifications [40], [41], [42], open APIs, and the support of major vendors (Intel, AMD Xilinx, Netronome, Netcope) through programmable hardware (SmartNICs, whitebox switches), and (iii) real-time, as it enables overlay analytics and AI systems for fine-grained packet-by-packet anomaly detection, congestion analysis, and other real-time analytics operations.

The ACROSS architecture aims to use standard AI/ML algorithms, such as Long-Short Term Memory (LSTMs), 1-D Convolutional Neural Networks (CNNs), and Graph Neural Networks (GNNs), to analyze the relationships between the VNFs in 5G networks. The selection of algorithms will vary depending on the network slice's lifecycle to achieve computational, orchestration-driven, and inter-slice resource elasticity [43]. The architecture will consider these three dimensions jointly to manage and orchestrate network resources to meet diverse QoS requirements. It will also explore AI-based traffic engineering (TE) techniques, based on methods that combine LSTM with 1D CNNs, to predict and react to network events. To ensure transparency and accountability, the architecture will integrate methods for interpretable models and will carefully evaluate the datasets used for model training, to mitigate potential biases [44].

The proposed architecture aims to achieve advanced automation beyond current device-level automation. This will be done through zero-touch automation support across multiple domains (RAN, edge, core cloud) and the entire hardware/software stack, from infrastructure to application layer. The platform will enable quick reactive automation by accommodating events from end devices and its telemetry system, and smart proactive automation through the platform's AI engine. Reactive and proactive automation cases will cover deployment, runtime management, and scaling zero-touch operations.

The proposed zero-touch architecture is designed to ensure secure orchestration by employing various mechanisms. Firstly, it enhances trusted computing technologies for large and varied cloud-edge deployments with the use of hardening trusted execution environments, protocols for secure enclave management, and support for open initiatives like Keystone for RISC-V CPUs. Secondly, the platform enhances the security of ML algorithms by developing defenses against data and model manipulation and resisting adaptive attacks. Lastly, the platform leverages SD-SEC for improved placement, orchestration, and monitoring of security VNFs through centralized management, automated monitoring, and enforcing KPIs in transport slice Security Service Level Agreements (SSLAs).

IV. Reference Architecture

ACROSS employs a (physically) distributed grid of domain-level ACROSS orchestrator instances, each dedicated to the management of a particular domain infrastructure. ACROSS domain orchestrators manage resources across the entire cloud continuum ranging from mobile radio access networks (RANs), transport and datacenter networks, as well as edge and core clouds. The ACROSS multi-domain end-to-end orchestrator is a logically centralised cloud-managed entity that ensures end-to-end service & network continuity across multiple - potentially geo-distributed - domains, constituting the entire edge-to-core cloud continuum. To bridge domain-level orchestrators with the multi-domain end-to-end orchestrator, ACROSS employs an intermediate integration fabric (see the top-right part in Fig. 1). This fabric ensures alliance with ongoing standardisation efforts (i.e., the ETSI ZSM integration fabric), but also fosters (i) separation of concerns through a scalable message bus between domain-specific and cross-domain functions and (ii) technology & vendor agnosticity through open standardised APIs. Atop the ACROSS multi-domain end-to-end orchestrator, ACROSS exposes a powerful front-end UI along with management APIs for orchestration, policy management, service management, AI, data generation, and telemetry towards a broad set of relevant stakeholders, such as operators, service providers, open-source communities, data scientists, and third parties.

The ACROSS orchestration approach has two control loop levels, represented by the yellow and purple circles in Fig. 1. At the domain level, the ACROSS domain orchestrator uses local telemetry and analytics (yellow step 1) to generate local intelligence, which drives zero-touch operations that interact with local orchestration loops to trigger necessary reconfigurations (yellow steps 2 and 3). The local orchestrator then communicates with the control block of the ACROSS domain orchestrator (yellow step 4) to carry out the reconfigurations on the underlying hardware. In a similar manner, the ACROSS multi-domain end-to-end orchestrator aggregates telemetry data across domains (purple steps 1 and 2), generates end-to-end analytics and intelligence (purple step 3), and triggers end-to-end automation functions that adjust inter-domain orchestration loops (purple step 4). To enforce end-to-end reconfigurations, the ACROSS multi-domain end-to-end orchestrator communicates with domain-level orchestrators through the ACROSS integration fabric (purple steps 5-7). The internal structure of both the domain-level orchestrators as well as the multi-domain end-to-end orchestrator is similar, as both contain seven logical blocks, in compliance with the ETSI GS ZSM 002 reference architecture [45].

V. Detailed overview of the ACROSS components

ACROSS Telemetry is a system that enhances visibility and control in modern infrastructures. It uses recent advancements in programmable networking hardware and cloud monitoring to create a deep end-to-end telemetry system. It offloads telemetry into whitebox switches and SmartNICs with hardware-offloaded netapps, providing detailed network state information without using host CPU. This information is combined with cloud monitoring tools to provide nano-to-micro-second level visibility into service state. This is achieved through distributed telemetry collectors, which store and provide the aggregated telemetry data to higher-layer blocks of ACROSS.

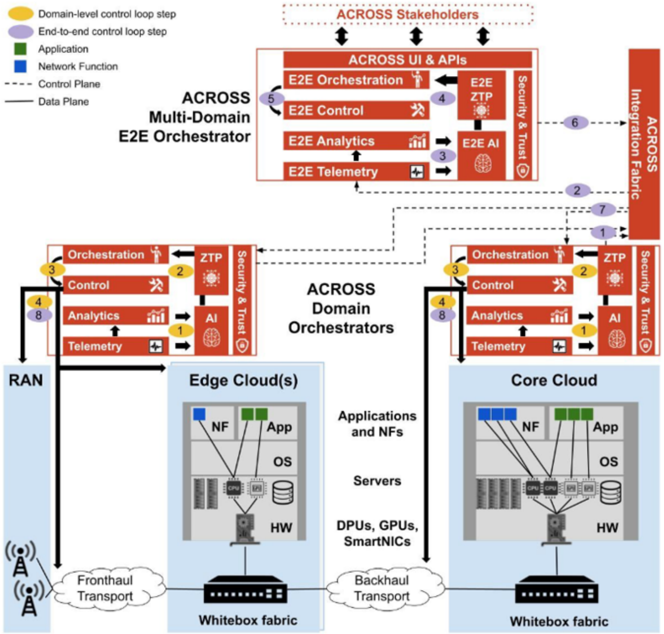

Fig. 1. Overview of the ACROSS zero-touch service management architecture

ACROSS Analytics uses telemetry data from the ACROSS system and environment to process and analyze data and events in real-time for (1) quick detection of service microbursts, (2) identification of congested flows for traffic engineering, (3) detection of unused resources for energy reduction, and (4) heavy-hitter detection for security and scaling. The system can use either traditional algorithms or AI-based algorithms and is designed to be flexible with an API that allows interchangeable use.

The ACROSS AI engine is based on an autonomous intelligent system for automating service deployment and management, using raw data and events from both the underlying infrastructure and the ACROSS telemetry. It goes beyond the data-driven approach of the ACROSS Analytics block by emulating human understanding and cognition, enabling learning, reasoning, and self-tuning of the infrastructure. The AI engine employs various AI schemes such as LSTM, 1D-CNN, NBeats, Transformers, etc., to turn data into intelligence. It operates both within a single domain and end-to-end, using policy-driven or data-driven triggers to select the appropriate AI algorithm. The AI engine stores the outcome and triggers zero-touch operations if necessary. ACROSS will also use synthetic data generated by (Generative Adversarial Networks) GANs to avoid privacy violations during ML training. The design and development of the AI engine are based on the concepts of Digital Map (DMap) and NDT-ready [46]. The DMap concept is a novel approach to build a complete Network Digital Twin (NDT) system in a modular way. This helps to avoid the high cost of planning, designing, and implementing a complete NDT system from scratch. The AI engine would be designed to utilize the DMap concept to build a complete NDT system incrementally and flexibly. The NDT-ready label would be incorporated to identify the fully interoperable nature of the AI Engine with the NDT. This would help the AI engine to standardize the communication channel between the physical network and NDT.

ACROSS Zero-Touch Provisioning is a part of the ACROSS Zero-touch orchestration engine, which automatically executes

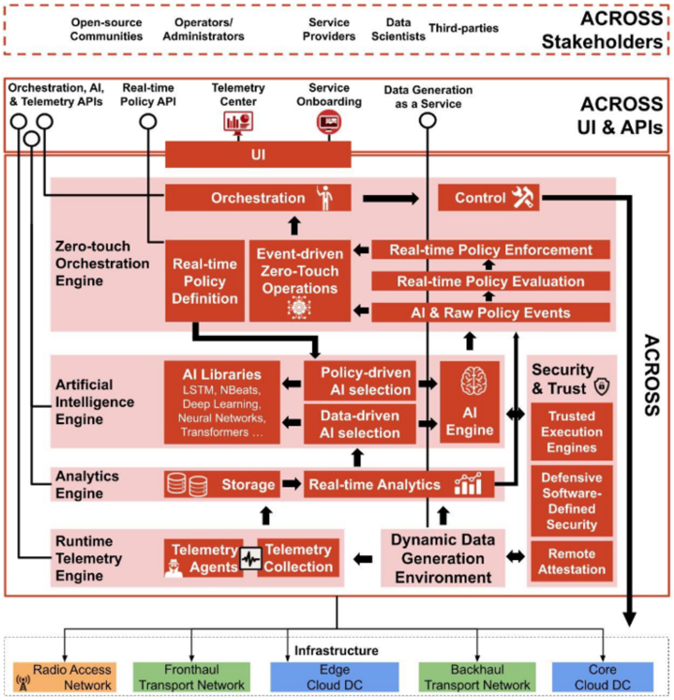
Open-source Communities
Operators/ Administrators
Service Providers
Data Scientists
Third-parties
ACROSS Stakeholders
Orchestration, AI, & Telemetry APIs
Real-time Policy API
Telemetry Center
Service Onboarding
Data Generation as a Service
ACROSS UI & APIs
UI
Orchestration
Control
Zero-touch Orchestration Engine
Real-time Policy Definition
Event-driven Zero-Touch Operations
Real-time Policy Enforcement
Real-time Policy Evaluation
AI & Raw Policy Events
ACROSS
Artificial Intelligence Engine
AI Libraries
LSTM, NBeats, Deep Learning, Neural Networks, Transformers ...
Policy-driven AI selection
Data-driven AI selection
AI Engine
Security & Trust
Trusted Execution Engines
Analytics Engine
Storage
Real-time Analytics
Defensive Software-Defined Security
Runtime Telemetry Engine
Telemetry Agents
Telemetry Collection
Dynamic Data Generation Environment
Remote Attestation
Infrastructure
Radio Access Network
Fronthaul Transport Network
Edge Cloud DC
Backhaul Transport Network
Core Cloud DC

Fig. 2. A detailed view of key ACROSS platform components

operations in response to events from the ACROSS AI engine. The engine analyzes ACROSS telemetry data and decides on tuning/reconfiguration operations to resolve detected/predicted issues or faults/anomalies. ACROSS conducts a study of common operations in production and experimental environments using service management and orchestration platforms from large operators and vendors to identify which tasks can be automated. ACROSS classifies these operations into 4 categories: (1) stakeholder-driven (from northbound ACROSS API or UI), (2) device-driven (from raw device events or telemetry data), (3) intelligence-driven (from the ACROSS AI engine), and (4) those requiring multiple operations to achieve a goal (e.g. new edge infrastructure).

ACROSS Orchestration is powered by two key components: ACROSS AI engine and ACROSS ZTP module. The AI engine processes telemetry data and makes decisions, while the ZTP module translates those decisions into actions taken by the infrastructure. The ACROSS orchestration engine includes a policy engine for stakeholders to create policies and establish automated control loops. The policies guide the AI engine to make decisions and trigger telemetry collection, then the ZTP module initiates actions based on the AI's reasoning. The orchestrator reconfigures resources through RAN, VIM, SDN, or WIM plugins. ACROSS's policy engine allows for multi-objective optimization of real-time policies, including energy and cost, which are crucial for future deployments to meet environmental and budget constraints.

ACROSS Control transforms tech-agnostic orchestration commands into technology-specific instructions for hardware, as shown in Fig. 2. It enhances popular infrastructure controllers, such as ETSI OSM and ONAP, to control programmable RANs, whitebox equipment, and manage energy-related resources in hardware. ACROSS brings support for these controllers and extends ETSI OSM to control energy consumption, e.g. through adjusting power states, enabling/disabling ports, autoscaling, hibernating containers [47], [48], and migrating services to less energy-intensive nodes.

ACROSS Security and Trust provides secure data management and advanced security techniques for edge-to-core cloud deployments, as shown in Fig. 2. It ensures trusted execution environments, protects telemetry and intelligence data with secure enclaves and secure storage, and focuses on secure intra-enclave and enclave-to-enclave communications and provides defensive SD-SEC mechanisms for the single-/multi- domain orchestrators.

ACROSS provides APIs and a UI to allow end-users and external systems to access its features. The Telemetry API provides visualization of telemetry data through the ACROSS Telemetry portal, which has real-time graphs of network and service parameters and highlights any issues. The Dynamic Data Generation API provides a way for AI engineers to increase the size of datasets for AI model training, while the Northbound AI API provides a view of telemetry data and supported AI libraries for AI development. The Northbound Real-Time Policy API allows for automatic creation, updating, and deletion of service and network policies with performance, energy, and cost constraints, while the Northbound Orchestration API facilitates the automated onboarding of application and network services and expands the ACROSS platform to new edge infrastructures. The ACROSS UI provides a wizard for quick service component onboarding and service composition, interfacing with the Orchestration API for an elevated user experience.

VI. Expected Outcome

The architecture is a secure, end-to-end network and service management solution designed to address the challenges of modern and future services in a rapidly evolving cloud processing environment. It offers a highly scalable service orchestration layer, comprised of a distributed network of domain orchestration instances managed by a multi-domain cloud-native service orchestrator. This allows for management of multi-faceted infrastructures across geographically dispersed edge-to-core deployments with enhanced security and trust. The platform leverages the growth of IoT, elevated data plane programmability, and AI to enhance service orchestration and outperform existing solutions. The proposed platform utilizes zero-touch mechanisms to optimize multi-objective service management while balancing performance, energy consumption, and cost. The platform promotes a democratized cloud ecosystem that benefits various market stakeholders and actors through principles such as event-driven microservices, open and standardized APIs, multi-faceted support, infrastructure agnosticity, abstracted service and business models, improved security and trust, and real-time automation through AI. The platform will set the standard for AI-driven, zero-touch service deployment and management, conforming to existing zero-touch standards and offering a roadmap for demonstration and advancement of these standards, particularly in light of the 6G requirements.

Acknowledgement

This work is supported by the Horizon Europe ACROSS Project (grant agreement No. 101097122).

References

[1] (2023, Feb.) 6G SNS STREAM A – SMART COMMUNICATION COMPONENTS, SYSTEMS AND NETWORKS FOR 5G MID-TERM EVOLUTION SYSTEMS. Accessed: 2022.

[2] OpenStack. [Online]. Available: https://www.openstack.org/

[3] Kubernetes. [Online]. Available: https://kubernetes.io/

[4] ONF Open Network Operating System (ONOS) SDN Controller. [Online]. Available: https://opennetworking.org/onos/

[5] OpenDaylight (ODL) SDN controller. [Online]. Available: https://www.opendaylight.org/

[6] ETSI Open Source Group for TeraFlowSDN (TFS) SDN controller. [Online]. Available: https://tfs.etsi.org/

[7] ETSI Open Source NFV Management and Orchestration (MANO). [Online]. Available: https://osm.etsi.org/

[8] The Linux Foundation Open Network Automation Platform. [Online]. Available: https://www.onap.org/

[9] ONF Aether 5G Connected Edge platform. [Online]. Available: https://opennetworking.org/aether/

[10] Netdata: Monitoring and troubleshooting transformed. [Online]. Available: https://www.netdata.cloud/

[11] Prometheus - Monitoring system time series database. [Online]. Available: https://prometheus.io/

[12] Sematext: Full Stack Infrastructure Monitoring. [Online]. Available: https://sematext.com/spm/

[13] Elastic Stack Monitoring. [Online]. Available: https://www.elastic.co/what-is/elasticsearch-monitoring

[14] Zabbix: The Enterprise-Class Open Source Network Monitoring. [Online]. Available: https://www.zabbix.com/

[15] Y. Wang, R. Forbes, C. Cavigioli, H. Wang, A. Gamelas, A. Wade, J. Strassner, S. Cai, and S. Liu, "Network management and orchestration using artificial intelligence: Overview of etsi eni," *IEEE communications standards magazine*, vol. 2, no. 4, pp. 58–65, 2018.

[16] D. Bega, M. Gramaglia, R. Perez, M. Fiore, A. Banchs, and X. Costa-Perez, "Ai-based autonomous control, management, and orchestration in 5g: From standards to algorithms," *IEEE Network*, vol. 34, no. 6, pp. 14–20, 2020.

[17] D. M. Gutierrez-Estevez, M. Gramaglia, A. De Domenico, N. Di Pietro, S. Khatibi, K. Shah, D. Tsolkas, P. Arnold, and P. Serrano, "The path towards resource elasticity for 5g network architecture," in *2018 IEEE wireless communications and networking conference workshops (WCNCW)*. IEEE, 2018, pp. 214–219.

[18] F. Vannella, G. Iakovidis, E. A. Hakim, E. Aumayr, and S. Feghhi, "Remote electrical tilt optimization via safe reinforcement learning," in *2021 IEEE Wireless Communications and Networking Conference (WCNC)*, 2021, pp. 1–7.

[19] A. Dalgkitsis, A. Chawla, A.-M. Bosneag, and C. Verikoukis, "Safeschema: Multi-domain orchestration of slices based on saferl for b5g networks," in *GLOBECOM 2022 - 2022 IEEE Global Communications Conference*, 2022, pp. 3435–3440.

[20] M. C. Huebscher and J. A. McCann, "A survey of autonomic computing—degrees, models, and applications," *ACM Computing Surveys (CSUR)*, vol. 40, no. 3, pp. 1–28, 2008.

[21] "Arista - Why is Zero Touch Provisioning (ZTP) Needed?"

[22] "Comarch - Zero-touch Network Provisioning: The Enabler of Automated Service Management in Telecoms."

[23] "Juniper - Zero Touch Provisioning."

[24] "How does Zero-Touch-Provisioning (ZTP) in Cisco SD-WAN work?"

[25] TP-link Omada. "Software Defined Networking (SDN) with Cloud Access. [Online]. Available: https://www.tp-link.com/in/omada-sdn/

[26] D. Lee, D. Kohlbrenner, S. Shinde, K. Asanović, and D. Song, "Keystone: An open framework for architecting trusted execution environments," in *Proceedings of the Fifteenth European Conference on Computer Systems*, 2020, pp. 1–16.

[27] J. Gu, Z. Hua, Y. Xia, H. Chen, B. Zang, H. Guan, and J. Li, "Secure live migration of sgx enclaves on untrusted cloud," in *2017 47th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2017, pp. 225–236.

[28] F. Alder, A. Kurnikov, A. Paverd, and N. Asokan, "Migrating sgx enclaves with persistent state," in *2018 48th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*. IEEE, 2018, pp. 195–206.

[29] B. Biggio, I. Corona, D. Maiorca, B. Nelson, N. Šrndić, P. Laskov, G. Giacinto, and F. Roli, "Evasion attacks against machine learning at test time," in *Machine Learning and Knowledge Discovery in Databases: European Conference, ECML PKDD 2013, Prague, Czech Republic, September 23-27, 2013, Proceedings, Part III 13*. Springer, 2013, pp. 387–402.

[30] C. Szegedy, W. Zaremba, I. Sutskever, J. Bruna, D. Erhan, I. Goodfellow, and R. Fergus, "Intriguing properties of neural networks," *arXiv preprint arXiv:1312.6199*, 2013.

[31] B. Biggio, B. Nelson, and P. Laskov, "Poisoning attacks against support vector machines," *arXiv preprint arXiv:1206.6389*, 2012.

[32] M. Jagielski, A. Oprea, B. Biggio, C. Liu, C. Nita-Rotaru, and B. Li, "Manipulating machine learning: Poisoning attacks and countermeasures for regression learning," in *2018 IEEE symposium on security and privacy (SP)*. IEEE, 2018, pp. 19–35.

[33] N. Papernot, P. McDaniel, X. Wu, S. Jha, and A. Swami, "Distillation as a defense to adversarial perturbations against deep neural networks," in *2016 IEEE symposium on security and privacy (SP)*. IEEE, 2016, pp. 582–597.

[34] C. Guo, M. Rana, M. Cisse, and L. Van Der Maaten, "Countering adversarial images using input transformations," *arXiv preprint arXiv:1711.00117*, 2017.

[35] C. Xie, J. Wang, Z. Zhang, Z. Ren, and A. Yuille, "Mitigating adversarial effects through randomization," *arXiv preprint arXiv:1711.01991*, 2017.

[36] A complete list of all (arxiv) adversarial example papers. [Online]. Available: https://nicholas.carlini.com/writing/2019/all-adversarial-example-papers.html

[37] G. Blanc, N. Kheir, D. Ayed, V. Lefebvre, E. M. de Oca, and P. Bisson, "Towards a 5g security architecture: Articulating software-defined security and security as a service," in *Proceedings of the 13th International Conference on Availability, Reliability and Security*, 2018, pp. 1–8.

[38] P. Alemany, D. Ayed, R. Vilalta, R. Muñoz, P. Bisson, R. Casellas, and R. Martínez, "Transport network slices with security service level agreements," in *2020 22nd International Conference on Transparent Optical Networks (ICTON)*. IEEE, 2020, pp. 1–4.

[39] J. Ortiz, R. Sanchez-Iborra, J. B. Bernabe, A. Skarmeta, C. Benzaid, T. Taleb, P. Alemany, R. Muñoz, R. Vilalta, C. Gaber *et al.*, "Inspire-5gplus: Intelligent security and pervasive trust for 5g and beyond networks," in *Proceedings of the 15th International Conference on Availability, Reliability and Security*, 2020, pp. 1–10.

[40] P4.org Applications Working Group, "In-band Network Telemetry (INT) Dataplane Specification version 2.1."

[41] P. Lapukhov and remy@barefootnetworks.com, "Data-plane probe for in-band telemetry collection," Internet Engineering Task Force, Internet-Draft draft-lapukhov-dataplane-probe-01, Jun. 2016, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/draft-lapukhov-dataplane-probe/01/

[42] F. Brockners, S. Bhandari, V. P. Govindan, C. Pignataro, H. Gredler, J. Leddy, S. Youell, T. Mizrahi, A. Kfir, B. Gafni, P. Lapukhov, and M. Spiegel, "VXLAN-GPE Encapsulation for In-situ OAM Data," Internet Engineering Task Force, Internet-Draft draft-brockners-ippm-ioam-vxlan-gpe-03, Nov. 2019, work in Progress. [Online]. Available: https://datatracker.ietf.org/doc/draft-brockners-ippm-ioam-vxlan-gpe/03/

[43] D. M. Gutierrez-Estevez, M. Gramaglia, A. de Domenico, N. di Pietro, S. Khatibi, K. Shah, D. Tsolkas, P. Arnold, and P. Serrano, "The path towards resource elasticity for 5g network architecture," in *2018 IEEE Wireless Communications and Networking Conference Workshops (WCNCW)*, Apr. 2018, pp. 214–219.

[44] R. Hamon, H. Junklewitz, and J. I. Sanchez Martin, "Robustness and explainability of artificial intelligence," *JRC Publications Repository*, Jan. 2020.

[45] P4.org Applications Working Group, "ETSI Zero touch network and service Management Reference Architecture."

[46] A. Mozo, A. Karamchandani, S. Gómez-Canaval, M. Sanz, J. I. Moreno, and A. Pastor, "B5gemini: Ai-driven network digital twin," *Sensors*, vol. 22, no. 11, p. 4106, Jan. 2022.

[47] Docker pause command. [Online]. Available: https://docs.docker.com/engine/reference/commandline/pause/

[48] Docker checkpoint command. [Online]. Available: https://docs.docker.com/engine/reference/commandline/checkpoint/