TRANSPARENT APPLICATION DEPLOYMENT IN A SECURE, ACCELERATED AND COGNITIVE CLOUD CONTINUUM

Grant Agreement no. 101017168

Deliverable D7.14
Final report on standardization activities

Programme:	H2020-ICT-2020-2
Project number:	101017168
Project acronym:	SERRANO
Start/End date:	01/01/2021 – 31/12/2023

Deliverable type:	Report
Related WP:	WP7
Responsible Editor:	ICCS
Due date:	31/12/2023
Actual submission date:	29/12/2023

Dissemination level:	Public
Revision:	FINAL

This project has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No 101017168

Revision History

Date	Editor	Status	Version	Changes
08.11.23	Panagiotis Kokkinos (ICCS)	Draft	0.1	Initial Version of the document
04.12.23	Paraskevas Bourgos (INTRA)	Draft	0.2	Updates from INTRA.
11.12.23	Ralf Scheider (HLRS)	Draft	0.3	Updates from HLRS.
12.12.23	Adrian Spătaru (UVT)	Draft	0.4	Updates from UVT.
14.12.23	Stathis Karanastasis (INNOV)	Draft	0.5	Updates from INNOV.
18.12.23	Amelia Pakouline-Navarro (MLNX/NVIDIA)	Draft	0.6	Updates from MLNX/NVIDIA.
19.12.23	Anastasios Nanos (NBFC)	Draft	0.7	Updates from NBFC.
19.12.23	Márton Sipos (CC)	Draft	0.8	Updates from CC.
20.12.23	Kostas Siozios (AUTH)	Draft	0.9	Updates from AUTH.
20.12.23	Panagiotis Kokkinos (ICCS)	Draft	0.10	Updates from ICCS.
22.12.23	Javier Martín (IDEKO)	Draft	0.11	Updates from IDEKO.
22.12.23	Ferad Zyulkyarov (INB)	Draft	0.12	Updates from INB.
27.12.23	Panagiotis Kokkinos (ICCS)	Draft	0.13	Integrate review changes and final enhancements
29.12.23	Panagiotis Kokkinos (ICCS)	Final Version	1.0	Final version of the deliverable

Author List

Organization	Author
ICCS	Emmanouel Varvarigos, Aristotelis Kretsis, Panagiotis Kokkinos, Polyzois Soumplis
MLNX	J.J. Vegas Olmos, Yoray Zack, Amelia Pakouline-Navarro
CC	Marton Sipos, Daniel E. Lucani, Marcell Fehér
USTUTT/HLRS	Teona Macharadze, Dmitry Khabi, Ralf Scheider
AUTH	Kostas Siozios, Dimosthenis Masouros, Argyris Kokkinis
INTRA	Makis Karadimas, Paraskevas Bourgos
INB	Javier Castillo, Ferad Zyulkyarov
INNOV	Filia Filippou, Andreas Litke, Stelios Pantelopoulos, Stathis Karanastasis
IDEKO	Javier Martín, Elena Urkia
UVT	Silviu Panica, Gabriel Iuhasz, Ioan Dragan, Adrian Spătaru
NBFC	Anastassios Nanos, Charalampos Mainas

Internal Reviewers

Filia Filippou, INNOV

Elena Urkia, IDEKO

Abstract: The deliverable reports on SERRANO standardization activities, compliance with standards and activities and co-operations with other projects and initiatives.

Keywords: standards, cloud, edge

Table of Contents

Abbreviations

3GPP	3rd Generation Partnership Project
AI	Artificial iItelligence
API	Application Programming Interface
ARIB	Association of Radio Industries and Businesses
ASVS	Application Security Verification Standard
ATIS	Alliance for Telecommunications Industry Solutions
B5G	Beyond 5G
BRAINE	Big data pRocessing and. Artificial Intelligence at the Network Edge
CADF	Cloud Audit Data Federation
CAPIF	Common API Framework
CCSA	China Communications Standards Association
CCSC	Cloud Computing Standards Committee
CI/CD	Continuous Integration and Continuous Delivery
CIMI	Cloud Infrastructure Management Interface
CMM	Capability Maturity Model
CNC	Computerized Numerical Control
CPIP	Cloud Portability and Interoperability Profiles
CRUD	Create Read Update Delete
CSCC	Cloud Standards Customer Council
D	Deliverable
DC	Data Centre
DevSecOps	Development, Security, and Operation
DFDL	Data Format Description Language
DMTF	Distributed Management Task Force
DoW	Description of Work
DRMAA	Distributed Resource Management Application API
DSOMM	DevSecOps Maturity Model
EC	European Commission
ECMA	European Computer Manufacturers Association
EDOAL	Expressive and Declarative Ontology Alignment Language
ETSI	European Telecommunications Standards Institute
EU	European Union
IEC	International Electrotechnical Commission
IEEE	Institute of Electrical and Electronics Engineers
IoT	Internet of Things
ISG	Industry Specification Group
ISMS	Information Security Management Systems

ISO	International Organization for Standardization
JSON	JavaScript Object Notation
MEC	Multi-access Edge Computing
ML	Machine Learning
MQTT	Message Queuing Telemetry Transport
NATO	North Atlantic Treaty Organization
NG-RAN	Next Generation RAN
NIST	National Institute of Standards and Technology
OASIS	Organization for the Advancement of Structured Information Standards
OCCI	Open Cloud Computing Interface
OGF	Open Grid Forum
OPC	Open Platform Communications
O-RAN	Open RAN
OSI	Open Systems Interconnection
OVF	Open Virtualization Format
OWASP	Open Web Application Security Project
PM	Project Manager
PO	Project Officer
RAN	Radio Access Network
SA	Standards Association
SAMM	Software Assurance Maturity Model
SIIF	Intercloud Interoperability and Federation
SLA	Service level agreement
SOA	service oriented architecture
TSDSI	Telecommunications Standards Development Society, India
TTA	Telecommunications Technology Association
TTC	Telecommunication Technology Committee
UA	Unified Architecture
VM	Virtual Machine
VNF	Virtual Network Functions
WDL	Workflow Description Language

1 Executive Summary

Deliverable 7.14 (D7.14) is the final report on SERRANO related standardization activities. In particular, D7.14 reports: on standards and organizations related to cloud and edge, on SERRANO components' and activities' relation to standards, and the co-operation of the SERRANO project and its partners with other projects and initiatives.

2 Introduction

One of the objectives of the SERRANO project is to follow, align, and contribute were appropriate and possible to standardization activities that take place in related areas. Therefore, these standardization activities were closely monitored during the project, and the possible impact of SERRANO results was timely identified. Moreover, cooperation with other projects and initiatives within the scope of the project was pursued, identifying opportunities for cooperation and knowledge transfer.

2.1 Purpose of this document

The purpose of this document (D7.14) is to provide the final report on related standardization activities, on the SERRANO possible participation and compliance with these activities and the cooperation with other projects and initiatives. For completeness, D7.14 is partially based on D7.5 “Report on standardization activities” that was submitted in M18.

2.2 Document structure

The present deliverable is split into four chapters:

- Organizations and Standards
- Compliance with standards and activities
- Co-operation with other projects and initiatives
- The SERRANO open-source related activities

2.3 Audience

This document is public.

3 Organizations and Standards

A number of organizations create standards for edge and cloud computing. These cover various aspects of the edge and cloud technologies, offering recommendations and guidance for successful implementation.

The **National Institute of Standards and Technology – NIST**[1] is part of the U.S. Department of Commerce. Its purpose is to promote U.S. innovation and industrial competitiveness by advancing measurement science, standards, and technology in ways that enhance economic security and improve quality of life. Several cloud-related activities are (or have been) carried out by NIST:

- The NIST Cloud Computing Program has developed a Cloud Computing Technology Roadmap, as one of many mechanisms in support of United States Government secure and effective adoption of the cloud computing model to reduce costs and improve services. The NIST Cloud Computing Standards Roadmap Working Group has surveyed the existing standards landscape for interoperability, performance, portability, security, and accessibility. Where possible, new and emerging standardization work has also been tracked and surveyed. Using this available information, current standards, standards gaps, and standardization priorities are identified.
- NIST Cloud Computing Reference Architecture establishes a baseline cloud computing architecture. It defines services and relationships between cloud service providers, consumers, and other stakeholders.
- NIST Cybersecurity Framework is a set of guidelines and best practices designed to help organizations improve their cybersecurity strategies. This framework also applies to cloud services.
- NIST provides and updates the Definition of Cloud Computing. The NIST definition lists five essential characteristics of cloud computing: on-demand self-service, broad network access, resource pooling, rapid elasticity or expansion, and measured service. It also lists three "service models" (software, platform, and infrastructure) and four "deployment models" (private, community, public, and hybrid) that together categorize ways to deliver cloud services. The definition is intended to serve as a means for broad comparisons of cloud services and deployment strategies, and to provide a baseline for discussion from what is cloud computing to how to best use cloud computing.

The **Institute of Electrical and Electronics Engineers – IEEE** [2] is the world's largest technical professional organization dedicated to advancing technology. The IEEE Standards Association (IEEE SA) is a globally recognized standards-setting body within IEEE. It develops consensus standards through an open process that engages the industry and brings together a broad

[1] www.nist.gov

[2] www.ieee.org

stakeholder community. Also, the Cloud Computing Standards Committee (CCSC) promotes standards development in all aspects of the cloud computing ecosystem. It facilitates the development and use of standards-based choices by cloud computing ecosystem participants (cloud vendors, service providers, and users) in areas such as cloud application interfaces, cloud portability interfaces, cloud management interfaces, cloud interoperability interfaces, cloud file formats, and cloud operation conventions. Related standards include:

- IEEE P1935, IEEE Draft Standard for Edge/Fog Manageability and Orchestration
- IEEE P2301, IEEE Draft Guide for Cloud Portability and Interoperability Profiles (CPIP)
- IEEE P2302, IEEE Draft Standard for Intercloud Interoperability and Federation (SIIF)
- IEEE P2303, IEEE Draft Standard for Adaptive Management of Cloud Computing Environments

In cooperation with NIST, IEEE also set the Cloud Computing Standard that defines a functional model for a cloud federation. The intended technical benefit of the standard is to enable a dynamic infrastructure that can support evolving business models and ultimately facilitate the growth of cloud computing.

International Organization for Standardization – ISO [3] is an independent, non-governmental international organization that brings together experts to share knowledge and develop voluntary, consensus-based, market-relevant International standards that support innovation and provide solutions to global challenges. ISO works together with IEC (International Electrotechnical Commission) [4] on standards and guides. There a several ISO/IEC cloud standards that relate to cloud computing in general, cloud security, management, SLAs, data management and others, including the following:

- ISO/IEC 17789:2014, Information technology -- Cloud computing -- Reference architecture. This standard defines cloud computing roles, activities and functional components, as well as how they interact.
- ISO/IEC 18384:2016, Information technology -- Reference architecture for service oriented architecture (SOA). This standard defines the vocabulary, guidelines and general technical principles underlying SOA, which are often deployed in cloud platforms.
- ISO/IEC 19086-1:2016, Information technology -- Cloud computing -- Service level agreement (SLA) framework. This standard provides the framework for preparing SLAs for cloud services.
- ISO/IEC 19941:2017, Information technology -- Cloud computing -- Interoperability and portability. This standard specifies the interoperability and portability aspects of cloud computing.

[3] www.iso.org

[4] www.iec.ch

- ISO/IEC 19944-1:2020, Cloud computing and distributed platforms -- Data flow, data categories and data use. This standard describes how data moves among cloud service vendors and users of cloud services.
- ISO/IEC Technical Specification 23167:2020, Information technology -- Cloud computing -- Common technologies and techniques. This standard describes technologies and techniques used in cloud computing, including VMs, hypervisors and containers.
- ISO/IEC TR 23188:2020, Information technology — Cloud computing — Edge computing landscape. This standard examines the concept of edge computing, its relationship to cloud computing and IoT, and the technologies that are key to the implementation of edge computing.
- ISO/IEC 27002:2022, Information security, cybersecurity and privacy protection — Information security controls. This standard provides a reference set of generic information security controls including implementation guidance.

The **Cloud Standards Customer Council – CSCC**[5] is an end-user advocacy group dedicated to accelerate cloud computing adoption, focusing on cloud issues such as standards, security, and interoperability. The work of the Council consists in: contributing to lowering the barriers for a widespread use of cloud computing; elaborating best practices, case studies, guides, and standards roadmaps on cloud computing-related issues such as interoperability, cloud architectures, service agreements, security, and industry technologies; liaising with standards development organisations and contributing to standards development processes for new cloud standards; facilitating the exchange of real-world stories, practices, lessons and insights.

DMTF, formerly the **Distributed Management Task Force**[6], creates open manageability standards spanning diverse emerging and traditional IT infrastructures, including cloud, virtualization, network, servers, and storage. In particular, DMTF produces standards and whitepapers in the following areas: Cloud Infrastructure Management Interface (CIMI), Open Virtualization Format (OVF), Cloud Audit Data Federation (CADF) Data Format and Interface Definitions, Software Identification and Entitlement Usage Metrics.

European Telecommunications Standards Institute – ETSI [7] is an independent, not-for-profit, standardization organization that primarily develops telecommunications standards. Among its cloud-focused activities are the Technical Committee Cloud, the Cloud Standards Coordination initiative, the Multi-access Edge Computing (MEC) initiative, and the Global Inter-Cloud Technology Forum, each of which addresses cloud technology issues. The related activities include: the Identification of Cloud Computing users' needs, Cloud Computing standards and Open Source, Cloud Computing Interoperability and Security and other areas of interest. In particular, the Multi-access Edge Computing (MEC) initiative is an Industry Specification Group (ISG) within ETSI. ETSI ISG MEC specified a common and extensible

[5] www.cloud-council.org

[6] www.dmtf.org

[7] www.etsi.org

application enablement framework for delivering services, specific service-related APIs for information exposure and programmability, as well as management, orchestration, and mobility related APIs. ETSI ISG MEC is currently studying MEC federations to enable shared usage of MEC services and applications across MEC systems to support a multi-operator/multinetwork/multi-vendor environment.

Open Grid Forum – OGF[8] is an open community committed to drive the rapid evolution and adoption of applied distributed computing. OGF develops standards for grid computing, cloud, advanced digital networking, and distributed computing technologies. A selection of the most popular standards frameworks resulting from this OGF community activity is highlighted below:

- The Open Cloud Computing Interface (OCCI)[9] specification set defines a general protocol and API applicable to many different cloud resource management tasks.
- The Data Format Description Language (DFDL)[10] is a language for describing text and binary data formats.
- The WS-Agreement and WS-Agreement Negotiation[11] family of specifications provide a language and a protocol for the creation, management, and monitoring of automated machine-readable service agreements in real-time.
- The Distributed Resource Management Application API (DRMAA)[12] is a set of API specifications for tightly coupled and portable programmatic access to cluster, grid, and cloud systems.

The **Organization for the Advancement of Structured Information Standards – OASIS**[13] is a nonprofit consortium that works on the development, convergence, and adoption of open standards for cybersecurity, blockchain, Internet of Things (IoT), emergency management, cloud computing, legal data exchange, energy, content technologies, and other areas. OASIS's various cloud technical committees include OASIS Cloud Application Management for Platforms, OASIS Identity in the Cloud, and OASIS Topology and Orchestration Specification for Cloud Applications.

The **3rd Generation Partnership Project - 3GPP** [14] unites seven telecommunications standard development organizations (ARIB, ATIS, CCSA, ETSI, TSDSI, TTA, TTC). The project covers cellular telecommunications technologies, including radio access, core network, and service capabilities, which provide a complete system description for mobile telecommunications. The 3GPP specifications also provide hooks for non-radio access to the core network and

[8] www.ogf.org
[9] occi-wg.org
[10] www.ogf.org/ogf
[11] www.ogf.org/documents/GFD.193.pdf
[12] drmaa.org
[13] www.oasis-open.org
[14] 3gpp.org

interworking with non-3GPP networks. 3GPP has defined a set of releases for the new versions of the standards, each containing new functionality. In Release 17, 3GPP aims to provide native support for Edge Computing in 3GPP networks. These efforts include initiatives across several working groups in 3GPP including SA2, SA3, SA4, SA5, and SA6 that cover application layer architecture, core network enhancement, security, media processing, and management aspects respectively. In particular:

- 3GPP SA2: Defines the architecture for mobile core networks including 5G. In the context of edge computing, 3GPP SA2 defines how user traffic is routed to the appropriate application servers in the edge clouds. It also provides the means for applications to provision traffic steering rules.
- 3GPP SA5: It is responsible for management aspects of 3GPP networks. In the context of edge computing, 3GPP SA5 is in the process of specifying life-cycle management of application servers in the edge cloud.
- 3GPP SA6: Defines an architecture for enabling Edge Applications, specifically through the specification of an enabling layer to facilitate communication between application clients and applications deployed at the edge. The architecture also enables the Common API Framework (CAPIF) to be leveraged as a standardized means of providing and accessing APIs in the Edge Cloud.

The **Linux Foundation** [15] is a non-profit technology consortium founded in 2000 as a merger between Open Source Development Labs and the Free Standards Group to build sustainable ecosystems around open-source projects, accelerating technology development and commercial adoption. Several cloud and edge related activities run under or were initiated by the Linux Foundation:

- Linux Foundation (LF) Edge [16] is an umbrella organization that aims to establish an open, interoperable framework for edge computing independent of hardware, silicon, cloud, or operating system. By bringing together industry leaders, LF Edge will create a common framework for hardware and software standards and best practices, critical to sustaining current and future generations of IoT and edge devices.
- The Open Container Initiative (OCI) [17] is an open governance structure formed under the auspices of the Linux Foundation to create open industry standards around container formats and runtimes.
- The Cloud Native Computing Foundation (CNCF) [18], part of the Linux Foundation, is an open-source software foundation that promotes the adoption of cloud-native computing.

[15] linuxfoundation.org

[16] www.lfedge.org

[17] opencontainers.org

[18] www.cncf.io

The **Cloud Security Alliance (CSA)**[19] is committed to promote best practices for fostering awareness and ensuring security within cloud computing environments. Also, CSA provides tools and guidance that help entire industries and countries build their own cloud assurance ecosystem and enhance their security strategy.

[19] cloudsecurityalliance.org

4 Compliance with Standards and Activities

In what follows, we present the compliance of particular aspects of the SERRANO platform to standards, grouped based on partners' contributions. These include compliance:

- to security standards for CI/CD, DevSecOps operations and the Stream Handler component (INTRA).
- to communication related standards, required for gathering data from machines and IoT devices that relate to UC2 - Industry 4.0 (IDEKO).
- to virtualization related standards related to the container images (NBFC).
- to storage operations related to UC3 – Secure storage (CC).
- to secure storage protocols (MLNX).
- to abstraction, workflow description, model alignment, machine learning, and communication related standards for the ARDIA models and the AI-Enhanced Service Orchestrator component (INNOV).

Netcompany-Intrasoft's (INTRA) solutions provided in SERRANO based on CI/CD, DevSecOps methodologies and the Stream Handler Platform have followed and are aligned with the list of information security related standards below (to which Netcompany-Intrasoft is certified):

- ISO 27001 – Information Security Management: The ISO/IEC 27001 standard provides requirements for Information Security Management Systems (ISMS) and enables organizations of any kind to manage the security of assets such as financial information, intellectual property, employee details or information entrusted by third parties.
- ISO 22301 – Business Continuity: This ISO 22301 standard specifies requirements to implement, maintain and improve a management system to protect against, reduce the likelihood of the occurrence of, prepare for, respond to, and recover from disruptions when they arise.
- CMM – Capability Maturity Model: The CMM model allows organizations to identify and prioritize business improvement efforts and examine how closely the processes relate to best practices. It provides reliable, clear, consistent, and actionable focus on performance improvements that will impact the business most and help build and improve capability.
- EU – NATO Security Clearance.

In addition to the above, INTRA's internal methods and practices are aligned with the following standards:

- OWASP SAMM – Software Assurance Maturity Model: The OWASP SAMM model provides an effective and measurable way to analyse and improve the secure development lifecycle. SAMM supports the complete software lifecycle and is technology and process agnostic.
- OWASP ASVS – Application Security Verification Standard: The OWASP Application Security Verification Standard (ASVS) project provides a basis for testing web application technical security controls and provides developers with a list of requirements for secure development.
- OWASP DSOMM - DevSecOps Maturity Model: The DevSecOps Maturity Model shows security measures applied when using DevOps strategies and how these can be prioritized. With the help of DevOps strategies, security can also be enhanced, such as application libraries and operating system libraries in docker images that can be tested for known vulnerabilities. The DevSecOps Maturity Model allows appropriate principles and measures to be implemented to counteract attackers.

Nubificus (NBFC) complies with the Open Container Initiative (OCI) specifications and standards to provide:

- Interoperable application packaging based on container images throughout the SERRANO platform and the Cloud-Edge continuum.
- Interoperable execution by implementing OCI-compatible container runtimes to execute the container images built.

The main communication-related standards used by **IDEKO** for the SERRANO project are the following:

- MQTT (Message Queuing Telemetry Transport), last update March 2019, published by OASIS. A lightweight protocol designed to allow many devices to publish data on the network.
- MODBUS, date of publication 1979, by OSI. This standard is used to gather vibration data from the controller.
- OPC_UA, published in November 2017, by the OPC UA foundation. It is used to gather data related to the position of the machine axes from the CNC.
- JSON (ECMA 404) second edition published in 2017, by ECMA international. It is a lightweight, text-based, language-independent syntax for defining data interchange formats, used for the output data of the ball screw health assessment.

Chocolate Cloud (CC) offers a GDPR-compliant file storage and sharing solution through the SkyFlok service. The General Data Protection Regulation (GDPR) is a legal framework that sets guidelines for the collection and processing of personal information of individuals within the

European Union (EU). The requirements are strict, and the penalties for non-compliance are significant.

CC strictly follows the principles related to the processing of personal data:

- The right to be informed about how your data is being collected and how it is used;
- The right of access, which allows users to be aware of and verify the lawfulness of the processing of their data;
- The right to rectification – users can rectify their personal data if it is inaccurate or incomplete;
- The right to erasure – enables users to request the deletion or removal of personal data;
- The right to restrict processing – the data subject shall have the right to obtain from the controller restriction of processing;
- The right to data portability – allows users to obtain and reuse their personal data for their own purposes across different services;
- The right to object – users have the right to object to processing based on legitimate interests or the performance of a task in the public interest/exercise of official authority, direct marketing, as well as the right to object to processing for purposes of scientific/historical research and statistics;
- The right not to be subject to automated decision-making, including profiling.

In designing the SERRANO-enhanced Secure Storage Service, CC has aimed to follow the storage interface provided by Amazon Web Services (AWS) Simple Storage Service (S3). While technically not a standard per se, it is a widely used interface in cloud-based object storage systems. Most commercial solutions offer some level of compatibility with the S3 API, making it a popular choice for application developers. The storage service has compatibility with regard to the following features:

- AWS S3 Signature Version 4 authentication.
- All CRUD endpoints for objects and buckets, with most options supported.
- All endpoints for multipart uploads, with most options supported.
- Support for HTTP Range queries (IETF RFC7233) when retrieving objects. Features are limited to those supported by AWS S3.
- Error codes
- XML data schema

Mellanox Technologies LTD (MLNX) closely follows different protocol developments relevant to secure storage solutions. Several storage protocols use the advantage of InfiniBand and RDMA for performance reasons (high throughput, low latency, and low CPU utilization).

- SCSI RDMA Protocol (SRP) is designed to take full advantage of the protocol off-load and RDMA features provided by the InfiniBand architecture.
- iSCSI Extensions for RDMA (iSER) is an extension of the data transfer model of iSCSI, a storage networking standard for TCP/IP. It uses the iSCSI components while taking advantage of the RDMA protocol suite. ISER is implemented on various storage targets such as TGT, LIO, SCST and out of scope of this manual. For various ISER targets configuration steps, troubleshooting and debugging, as well as other implementation of storage protocols over RDMA (such as Ceph over RDMA, nbdX and more) refer to Storage Solutions on the Community website.
- Lustre is an open-source, parallel distributed file system, generally used for large-scale cluster computing that supports many requirements of leadership class HPC simulation environments.
- NVM Express™ over Fabrics (NVME-oF): NVME-oF is a technology specification for networking storage designed to enable NVMe message-based commands to transfer data between a host computer and a target solid-state storage device or system over a network such as Ethernet, Fibre Channel, and InfiniBand. Tunneling NVMe commands through an RDMA fabric provides high throughput and low latency. This is an alternative to the SCSi-based storage networking protocols. NVME-oF Target Offload is an implementation of the new NVME-oF standard Target (server) side in hardware. Starting from ConnectX-5 family cards, all regular IO requests can be processed by the HCA, with the HCA sending IO requests directly to an actual NVMe PCI device using peer-to-peer PCI communications. This means that excluding connection management and error flows, no CPU utilization will be observed during NVME-oF traffic.

MLNX is actively engaged in driving RDMA solutions and the integration of NVMe solutions within the standardization bodies and open-source communities.

INNOV-ACTS Limited (INNOV) is leading the activities regarding the development of the ARDIA framework and the AI-enhanced Service Orchestrator. The design of the respective abstraction models and services was driven by the identified user needs in the context of the SERRANO platform objectives, and by relevant standards and well-recognised, widely used specifications in the field, for ensuring the seamless and meaningful interaction among the components and services.

The design of the Abstraction Models that are part of the ARDIA framework was driven by the OASIS Topology and Orchestration Specification for Cloud Applications (TOSCA) standard[20]. In particular, the elements included in the Resource Model align with this standard so that an

[20] https://www.oasis-open.org/committees/tosca/

application can be easily deployed by a software component that supports the standard to the appropriate resources, taking into account the given parameters. Also, the design of the Telemetry Data Model was driven by the elements captured by widely used software tools in this field, such as Kubernetes[21]. Regarding the Application Model (high-level description of the application, the user's requirements and intent), an Object Oriented approach was followed for capturing the parameters of particular interest for the whole application or its particular components, also taking into account the organization of terms in the other two abstraction models (e.g., type of Resources attached to each Node The Workflow Description Language (WDL)[22] will be used to capture the actual flow of data and the interaction among the components.

The specification of mapping rules among the elements of the three abstraction models was driven by the Expressive and Declarative Ontology Alignment Language (EDOAL)[23] so that the source and target elements can be precisely determined along with additional parameters that may be necessary for the transition from one data representation to the other one. Hence, other software tools supporting this language can further process the correspondences specified. In SERRANO, these mapping rules are being used by the AI-enhanced Service Orchestrator to translate the application parameters (or constraints) to the appropriate corresponding ones that are expressed using Resource Model terms so that the Resource Orchestrator can use them accordingly. The design of the AI-enhanced Service Orchestrator and especially its internal component that is responsible for the introduction of machine learning techniques in the service orchestration process (Forecasting Mechanisms) is currently based on widely used Python libraries in this field, such as PySpark[24] and PyTorch[25] to name a few. Finally, the JSON[26] standard (ECMA-404 second edition, published in 2017) is used for the definition of the data interchange messages among the AI-Enhanced Service Orchestrator and other SERRANO platform components, while the YAML[27] human-friendly data serialization language is used for the definition of the application deployment description part.

[21] https://kubernetes.io/

[22] https://openwdl.org/

[23] https://moex.gitlabpages.inria.fr/alignapi/edoal.html

[24] https://spark.apache.org/docs/latest/api/python/

[25] https://pytorch.org/

[26] https://www.ecma-international.org/publications-and-standards/standards/ecma-404/

[27] https://yaml.org/

5 Open-source SERRANO Components and Other Available Services

SERRANO partners developed components that compose the SERRANO platform, are in large majority open-source. In the following table we report the repositories of these open-source components (Table 1).

Table 1: SERRANO Open-Source repositories per component.

Component	Maintained by	Repository
SERRANO SDK	INTRA	https://github.com/ict-serrano/serrano-sdk
Resource Orchestrator	ICCS	https://github.com/ict-serrano/Resource-Orchestrator
Resource Optimization Toolkit	ICCS	https://github.com/ict-serrano/Resource-Optimization-Toolkit
Telemetry Framework	ICCS	https://github.com/ict-serrano/Telemetry-Framework
HPC Interface	USTUTT/HLRS	https://github.com/ict-serrano/hpc-interface
AI-enhanced Service Orchestrator	INNOV	https://github.com/timchros/innov-ardia-aiso
A4C Orchestrator Plugin	UVT	https://github.com/adispataru/tufa-a4c-orchestrator
Service Assurance and Remediation	UVT	https://github.com/ict-serrano/service-assurance-ede
vAccel	NBFC	https://github.com/cloudkernels/vaccelrt
Container Runtimes	NBFC	https://github.com/nubificus/urunc

We also report on additional SERRANO-related services that are available (Table 2).

Table 2: Additional SERRANO-related services.

Service	Maintained by	Link
Cloud Monitoring	CC	https://www.skyflok.com/backend-performance

6 Co-operation with Other Projects and Initiatives

6.1 Projects

The **BRAINE EU project's**[28] overall aim is to boost the development of the Edge framework and, specifically, energy-efficient hardware and AI-empowered software systems, capable of processing Big Data at the Edge, supporting security, data privacy, and sovereignty. BRAINE's overall aim will be reached by targeting five fine-grained goals:

- Devising an EC infrastructure that offers control, computing, acceleration, storage, and 5G networking at the Edge and excels in scalability, agility, security, data privacy, and data sovereignty in Big Data and AI for low latency and mission-critical applications.
- Developing a future-proof Edge security framework and associated infrastructure based on the latest software and hardware security technologies.
- Developing a distributed and partly-autonomous system that takes data privacy and sovereignty into account on each and every decision regarding workload placement, data transfer, and computation, while guaranteeing interoperability with the environment.
- Developing a heterogeneous, energy-efficient Edge MicroDataCenter, suitable for stationed, mobile, and embedded autonomous applications, which goes beyond the current hardware and software architectures and offers Big Data processing and AI capabilities at the Edge.
- Testing and demonstrating the effectiveness and generality of the BRAINE approach by evaluating multiple real-world use cases and scenarios that exhibit the required scalability, security, efficiency, agility, and flexibility concerns.

SERRANO and BRAINE EU projects co-organized the workshop "Intelligent operations, security, and acceleration for edge computing" in the IEEE International Mediterranean Conference on Communications and Networking (MeditCom) 2021. The workshop sought to attract high-quality contributions covering both theory and practice over edge computing. In particular, the topics of interest included, but were not limited to the following areas:

- Security, privacy and data integrity in the edge
- Orchestration of edge resources
- Network and cloud telemetry
- Integrating AI with Edge Computing
- Machine Learning integration with Edge Computing
- Application of AI/ML at the edge

[28] www.braine-project.eu

- Edge intelligence
- Applications / VNFs deployed at the edge
- Acceleration of intensive workloads
- Networking programmability at the edge
- Verticals running in the edge

The **MARSAL EU project** [29] aims to provide an evolved architecture towards Beyond 5G (B5G), offering unprecedented degrees of flexibility and closed-loop autonomy at all tiers of the infrastructure, and significantly improved Spectral Efficiency via Cell-Free Networking. The overall concept of the MARSAL project is structured over three main pillars that in turn highlight the core activities of the project:

- The network design pillar: offers a combination of innovative cell-free and Hybrid Multiple Input Multiple Output technologies for the Radio Access Network (RAN) and Fronthaul domains, fully aligned with recent initiatives towards an Open RAN (O-RAN);
- The Elastic Edge Infrastructure pillar will support a fully Elastic Edge Cloud and dynamic slicing support for the wireless and optical domains, offering zero perceived latency to Multi-access Edge Computing (MEC) applications.
- The Network security pillar will focus on the security and privacy implications of multi-tenant infrastructures, offering a holistic framework for end users and tenants.

MARSAL adopts an evolved 3rd Generation Partnership Project (3GPP) Next Generation RAN (NG-RAN), which is extended with emerging Cell-Free technologies for network densification. Moreover, MARSAL contributes with innovations at the optical transport domain and significantly evolves the MEC system towards fully elastic Edge Computing. MARSAL will deploy a distributed Edge infrastructure with Data Centres (DCs) structured in 2 tiers, featuring Regional Edge and Radio Edge nodes. Radio Edge DCs will host the Network Functions of the (virtualized) RAN, which are fully aligned with the O-RAN specifications.

SERRANO and MARSAL projects through ICCS as a common partner, exchange knowledge regarding the use of edge and cloud technologies for serving "5G and beyond" networks' baseband processing requirements. This critical use case scenario investigated in MARSAL can be mostly benefited from the edge/cloud technologies developed in SERRANO.

The **PUZZLE project** [30] focuses on multi-dependency cyber-physical risk assessment, edge trust assurance services and remote attestation, distributed processing, programmable networking mechanism, cybersecurity analytics, deep analysis and distributed machine learning, threat intelligence and blockchain technologies. PUZZLE has a strong focus on security, and hence

[29] www.marsalproject.eu
[30] puzzle-h2020.com

there is a relevant match in relation to the SERRANO use case on secure storage and also about general concepts in relation to secure edge-cloud continuum access.

SERRANO and PUZZLE projects through MLNX as a common partner, exchange knowledge regarding the utilization of programmable networking mechanisms to enhance security and attempt to incorporate some of the security concepts developed on PUZZLE into SERRANO.

IoTAC[31] project emphasizes a multi-layered approach to IoT security, incorporating best practices and state-of-the-art technologies to create a comprehensive security framework. It includes advanced access control mechanisms, runtime protection features, and a Security-by-Design methodology that integrates security into every step of the software development life cycle. IoTAC also aims to provide a certification program based on international security standards, best practices, and research results of the project.

The IoTAC and SERRANO projects, while different in their specific focus, share commonalities in their approach to security, privacy, and deployment in the context of cloud-based, IoT, and edge computing environments. Information exchange took place between these two projects through INTRA as a common partner.

Areas for information and knowledge exchange between the two projects include:

1. **Security and Privacy**: Both projects emphasize robust security features. IoTAC benefits from SERRANO's approach to security and privacy in distributed computing and storage infrastructures. Similarly, SERRANO benefited from IoTAC's multi-layered security framework, particularly its emphasis on advanced access control and runtime protection.
2. **Software Development Lifecycle (SDLC)**: IoTAC's Security-by-Design methodology that integrates security into every step of SDLC provided valuable insights to SERRANO, especially in transparent application deployment.
3. **Resource Orchestration**: SERRANO's cognitive and automated orchestration of edge, cloud, and HPC resources offered valuable insights for IoTAC, particularly in terms of managing and securing IoT architectures.
4. **Standards and Certification**: IoTAC's focus on certification based on international standards and best practices benefited SERRANO, especially in terms of establishing a robust and recognized security framework.
5. **Integration and Deployment**: Both projects aim to provide smooth, rapid integration of their components into diverse architectures. Sharing best practices, tools, and technologies enhanced both projects' efficiency and ease of integration.

By exchanging knowledge, IoTAC and SERRANO enhanced their respective capacities to manage security and privacy, streamline SDLC processes, optimize resource orchestration, and improve overall system integration and deployment.

[31] https://iotac.eu/

CC and NBFCs also participate in the MLSysOps EU project[32]. The main objective of MLSysOps is to design, implement, and evaluate a complete AI-controlled framework for autonomic end-to-end system management across the full cloud-edge continuum. MLSysOps will employ a hierarchical agent-based AI architecture to interface with the underlying resource management and application deployment/orchestration mechanisms of the continuum. CC's and NBFC's SERRANO developments will also be partially used and further extended in this project.

ALLEGRO EU project[33] aims to design and validate a novel end-to-end sliceable, reliable, and secure architecture for next-generation optical networks, achieving high transmission/switching capacity. Among other key enabling innovations, ALLEGRO will develop a scalable AI/ML assisted control and orchestration system responsible for autonomous networking, dynamic and constrained service provisioning, function placement and resource allocation, leveraging devices increasing programmability and overall network softwarization. ICCS participates in the project and will utilize parts of the SERRANO platform for edge and cloud resource orchestration and workload (e.g., virtual network functions) allocation while interacting with optical network management systems and considering networking parameters.

CLEVER EU project[34] proposes a series of innovations in hardware accelerators, design stack, and middleware software that revolutionize the ability of edge computing platforms to operate federatedly, leveraging sparse resources that are coordinated to create a powerful swarm of resources. CLEVER will overcome traditional limitations of edge computing regarding limited resource availability by providing an effective framework for the seamless use of federated resources in the edge-cloud continuum. ICCS participates in this project, utilizing parts of the SERRANO platform to enable the management of computing resources and the allocation of applications and decomposed applications' (e.g., microservices) workloads.

6.2 Initiatives

6.2.1 EUCloudEdgeIoT

SERRANO participates in the activities (meetings, presentations, articles) of the EUCloudEdgeIoT.eu[35,36]. This initiative aims to assist in the understanding and development of the Cloud, Edge, and IoT (CEI) Continuum, promoting cooperation between a wide range of research projects, developers and suppliers, business users, and potential adopters of this new technological paradigm.

[32] https://mlsysops.eu/
[33] https://www.allegro-he.eu/
[34] https://www.cleverproject.eu/
[35] https://eucloudedgeiot.eu/
[36] https://eucloudedgeiot.eu/members/european-research-and-innovation-projects/

One of the main mechanisms the EUCloudEdgeIoT.eu has put in place in order to foster cooperation between the members of the community is the formation of six individual task forces (TF):

- TF1 - Strategic Laisons
- TF2 - Open Source Engagement
- TF3 - Architecture
- TF4 - Ecosystem Engagement
- TF5 - Market & Sectors
- TF6 - Communications

SERRANO participates in TF2 Open Source Engagement and TF3 Architecture. TF2 is lead by Eclipse and has the following key objectives:

- Develop a strategy for European digital autonomy in edge-to-cloud through Open Source
- Contribute to the definition of a common open architecture for the computing continuum
- Convince industry and research actors of the potential of Open Source to drive innovation and collaboration
- Train industry and research actors to embrace a long-term Open Source strategy

TF3 is lead by ATOS and has the following key objectives

- Enable the architectural discussion among projects in the area of IoT/Edge and Cloud to create a continuum.
- Identification of the thematic areas and building blocks.
- Understanding the contribution of each project to the thematic areas

Beyond the general participation in the task forces' activities, SERRANO described its architecture and technological developments and contributed to the mapping between the SERRANO architecture and the European reference architecture for the continuum. The definition of the reference architecture is one of the main objectives of TF3 within the EUCloudEdgeIoT initiative. In addition, SERRANO contributes in the definition of the reference architecture key build blocks, their features and functionalities, and communications workflows.

Finally, SERRANO contributes to the efforts of TF2 in defining a common Open Source Stack on the Cloud-to-Edge-to-IoT continuum by filling out a set of surveys prepared by the Eclipse Foundation. These surveys are focused on understanding the architecture and software dependencies for each EU-funded research project in the context of the cloud edge continuum.

6.2.2 Unikernels - Hackathon

Additionally, ICCS and NBFC co-organized a two-day hackathon[37] for unikernels. The format of the hackathon was the following:

- The first day, attendees had the opportunity to learn about Unikraft, a popular unikernel framework and to understand the use-cases of unikernels via three invited talks. ICCS and NBFC presented SERRANO and how unikernels can facilitate serverless function deployment.
- The second day, attendees were able to experiment with Unikraft, via specific challenges, coordinated by members of the Unikraft open-source community.

The hackathon had eighty (80) registered attendees. The first day attendance peaked at seventy (70), whereas the second day attendance peaked at forty (40). Twenty-eight (27) of the attendees completed the challenges and got participation certificates. Most attendees were students from universities across Greece, along with some professionals and interested third parties.

6.2.3 Unikernels Alliance

NBFC co-founded the Unikernel Alliance[38], a community-driven initiative to promote the use of unikernels. The Unikernel Alliance was recently formed to group a large set of open source unikernel (specialized VMs) projects, all on a mission to significantly improve cloud deployment efficiency, lower cloud resources spent, and run equivalent workloads on a much smaller infrastructure, leading to greener deployments. The first face-to-face meeting was held in RWTH, Aachen, Germany, in June 2023.

6.2.4 HiPEAC

SERRANO partners ICCS and NBFC co-organized a thematic session in the context of HiPEAC's Computing Systems Week in Oct 2021, where invited speakers elaborated on:

- workload management and orchestration
- efficient execution
- secure & trust in multi-tenant cloud & edge environments.

The session had 52 attendees from 34 institutions and 15 countries.

[37] https://unikraft.org/hackathons/2023-03-athens
[38] https://unikernelalliance.org

7 Conclusion

D7.14 reported on standardization activities from various bodies in the areas of cloud and edge computing. The compliance of particular aspects of the SERRANO platform to standards is also presented, and SERRANO's co-operation with other projects and initiatives is discussed.