Assurance-aware 5G Edge-Cloud Architectures for Intensive Data Analytics

MuseMI Seminar

Filippo Berto

Research Fellow @ SESAR Lab
Department of Computer Science
Università degli Studi di Milano

Feb. 26 2024

Motivation

Nowadays data intensive workflows are increasingly deployed in the Edge-Cloud continuum

Data is collected and preprocessed at the Edge and moved to the Cloud only when necessary

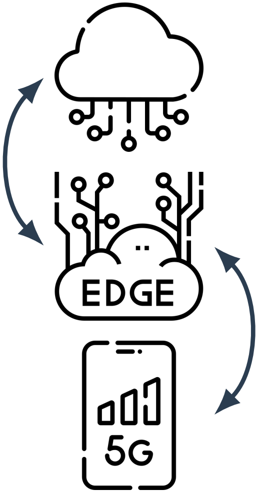

Nowadays data intensive workflows are increasingly deployed in the Edge-Cloud continuum

Data is collected and preprocessed at the Edge and moved to the Cloud only when necessary

5G technology is a fundamental enabler for the continuum, supporting private low-latency communication and advanced peripheral processing capabilities

Need for trustworthiness on the infrastructures and the services deployed on them

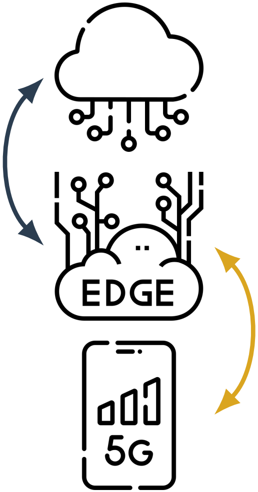

Challenges

- Current 5G standards are not fully ready for the edge-cloud continuum

Challenges

- Current 5G standards are not fully ready for the edge-cloud continuum
- Severe difficulties in handling security and privacy

Challenges

- Current 5G standards are not fully ready for the edge-cloud continuum
- Severe difficulties in handling security and privacy
- Current security and privacy assurance solutions are not well fitted for the dynamicity and heterogeneity of the continuum

Challenges

- Current 5G standards are not fully ready for the edge-cloud continuum
- Severe difficulties in handling security and privacy
- Current security and privacy assurance solutions are not well fitted for the dynamicity and heterogeneity of the continuum
    - Focus on application level, leaving strong expectations on the infrastructure

Literature Gaps

- **G1** Current 5G standards lack support for advanced security and QoS features, and integration with cloud environments

P. Ranaweera et al., M. Agiwal et al., R. Khan et al., F. Spinelli et al., T. Taleb et al., M. Anisetti et al.
ETSI, GS MEC 003 Multi-access Edge Computing (MEC); Framework and Reference Architecture. V3.1.1, 2022.
ETSI, ‘Multi-access Edge Computing (MEC); Framework and Reference Architecture’, ETSI ISG, ETSI GS MEC 003 V3.1.1, Mar. 2022.

Literature Gaps

- **G1** Current 5G standards lack support for advanced security and QoS features, and integration with cloud environments
- **G2** Research on distributed service workflows mainly focuses on FaaS, which does not fit well with modern data-intensive stateful workflows

W. Z. Khan et al., N. Rieke et al., S. Nastic et al., V. Casola et al., A. Orive et al., K. Fu et al

Literature Gaps

- **G1** Current 5G standards lack support for advanced security and QoS features, and integration with cloud environments
- **G2** Research on distributed service workflows mainly focuses on FaaS, which does not fit well with modern data-intensive stateful workflows
- **G3** Literature lacks a complete framework of non-functional properties, only performance-oriented ones are commonly recognized

S. Poojara et al., M. Glinz et al., M. Binkhonain et al., A. Bialas et al., M. Anisetti et al., E. Damiani et al., C. A. Ardagna et al.

Literature Gaps

- **G1** Current 5G standards lack support for advanced security and QoS features, and integration with cloud environments
- **G2** Research on distributed service workflows mainly focuses on FaaS, which does not fit well with modern data-intensive stateful workflows
- **G3** Literature lacks a complete framework of non-functional properties, only performance-oriented ones are commonly recognized
- **G4** Lack of non-functional aware workflow deployment solution for the continuum

Á. Santos et al., M. Anisetti et al., F. Giannone et al., C. Hebert et al., A. Brogi et al., V. Casola et al.

- Novel notion of 5G enabled edge-cloud continuum (G1)

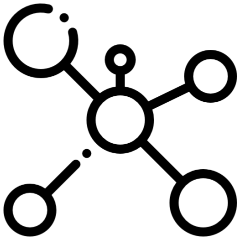

M. Anisetti, F. Berto, and M. Banzi. "Orchestration of data-intensive pipeline in 5G-enabled Edge Continuum". In: *2022 IEEE World Congress on Services (SERVICES)*. ISSN: 2642-939X. IEEE Computer Society, July 2022, pp. 2–10

Contributions

- Novel notion of 5G enabled edge-cloud continuum (G1)
- Realization of a complete continuum infrastructure including a fully functional simulated 5G stack (G2)

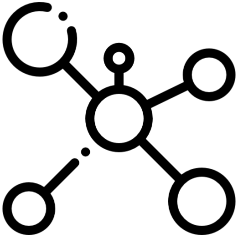

F. Berto, C. Ardagna, M. Torrente, D. Manenti, E. Ferrari, A. Calcante, R. Oberti, C. Fra', and L. Ciani. "A 5G-IoT enabled Big Data infrastructure for data-driven agronomy". In: *2022 IEEE Globecom Workshops (GC Wkshps)*. Rio de Janeiro, Brazil: IEEE, Dec. 2022, pp. 588–594

F. Berto, C. Ardagna, M. Torrente, D. Manenti, E. Ferrari, A. Calcante, R. Oberti, C. Fra', and L. Ciani. "A 5G-IoT enabled Big Data infrastructure for data-driven agronomy". In: *The 1st Italian Conference on Big Data and Data Science (ITADATA)*. Milan, Italy, Oct. 2022

Contributions

- Novel notion of 5G enabled edge-cloud continuum (G1)
- Realization of a complete continuum infrastructure including a fully functional simulated 5G stack (G2)
- Novel assurance methodology for modern distributed infrastructure (G3)

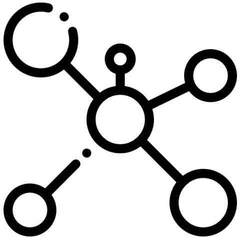

M. Anisetti, C. A. Ardagna, and F. Berto. "An assurance process for Big Data trustworthiness". In: *Future Generation Computer Systems* 146 (Sept. 2023), pp. 34–46

M. Anisetti, N. Bena, F. Berto, and G. Jeon. "A DevSecOps-based Assurance Process for Big Data Analytics". en. In: *2022 IEEE International Conference on Web Services (ICWS)*. Barcelona, Spain: IEEE, July 2022, pp. 1–10

M. Anisetti, C. A. Ardagna, F. Berto, and E. Damiani. "A Security Certification Scheme for Information-Centric Networks". en. In: *IEEE Trans. Netw. Serv. Manage.* 19.3 (Sept. 2022), pp. 2397–2408

M. Anisetti, C. A. Ardagna, F. Berto, and E. Damiani. "Security Certification Scheme for Content-centric Networks". In: *2021 IEEE International Conference on Services Computing (SCC)*. IEEE, Sept. 2021, pp. 203–212

Contributions

- Novel notion of 5G enabled edge-cloud continuum (G1)
- Realization of a complete continuum infrastructure including a fully functional simulated 5G stack (G2)
- Novel assurance methodology for modern distributed infrastructure (G3)
- Property-aware deployment solution for the assurance-focused continuum (G4)

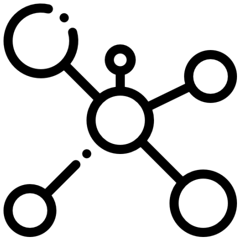

M. Anisetti, F. Berto, and R. Bondaruc. "QoS-Aware Deployment of Service Compositions in 5G-Empowered Edge-Cloud Continuum". In: *2023 IEEE 16th International Conference on Cloud Computing (CLOUD)*. ISSN: 2159-6190. IEEE, July 2023, pp. 471–478

Current notion of Edge-Cloud Continuum

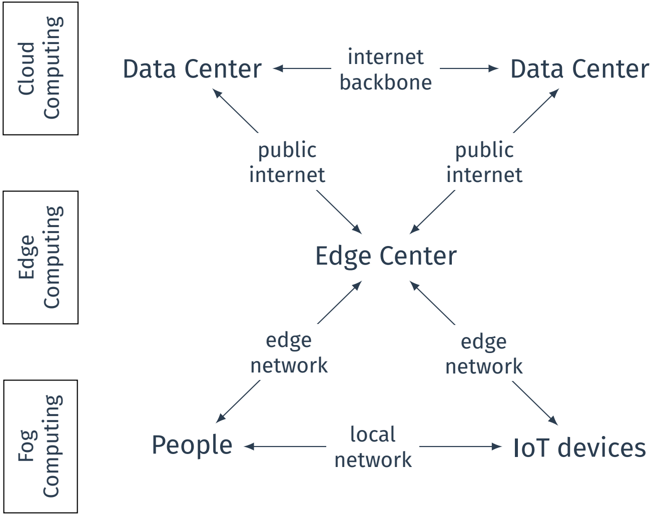

Our 5G enabled continuum

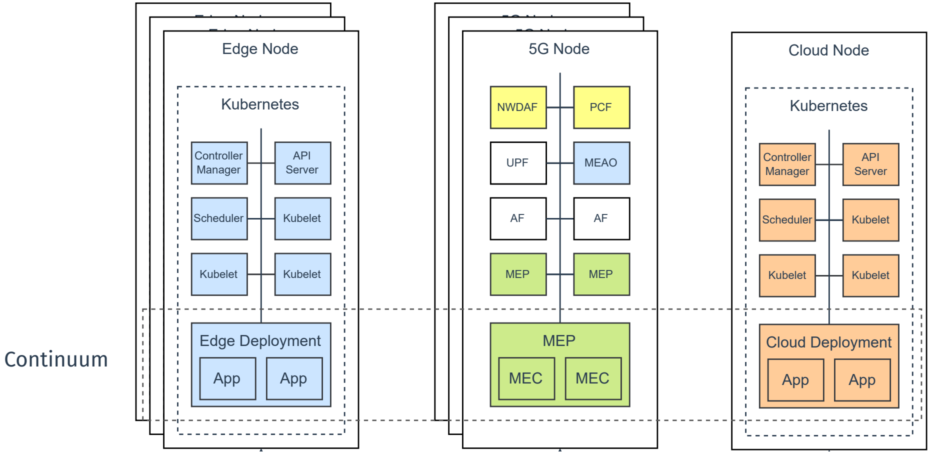

M. Anisetti, F. Berto, and M. Banzi. "Orchestration of data-intensive pipeline in 5G-enabled Edge Continuum". In: *2022 IEEE World Congress on Services (SERVICES)*. ISSN: 2642-939X. IEEE Computer Society, July 2022, pp. 2–10

Assurance Definition

"Way to gain justifiable confidence that IT systems will consistently demonstrate a (set of) security property and operationally behave as expected"

Infrastructure Assurance

Assurance allows the verification of properties on a system by inferring over evidence

Components expose state and configuration through monitoring endpoints

Evidence is collected my measuring the system state and configuration using metrics

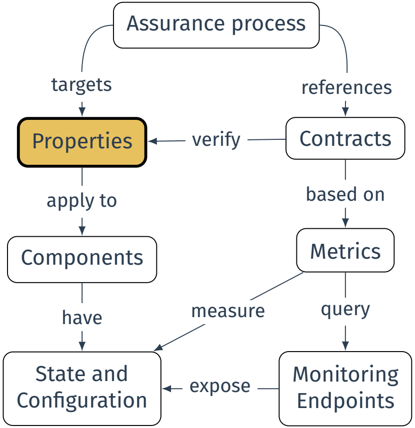

Infrastructure Assurance

Assurance allows the verification of properties on a system by inferring over evidence

Components expose state and configuration through monitoring endpoints

Evidence is collected my measuring the system state and configuration using metrics

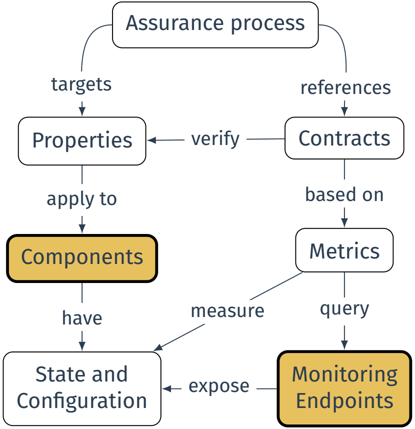

Infrastructure Assurance

Assurance allows the verification of properties on a system by inferring over evidence

Components expose state and configuration through monitoring endpoints

Evidence is collected my measuring the system state and configuration using metrics

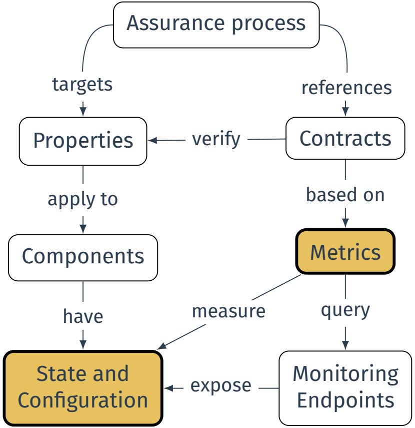

Infrastructure Assurance

Each property is associated with a contract that describes how to verify it in terms of metrics

Properties verification gives us guarantees on the behavior of the system

Assurance process continuously verifies contracts based on the collected evidence

Produce trustworthiness in infrastructures, as building blocks for distributed applications

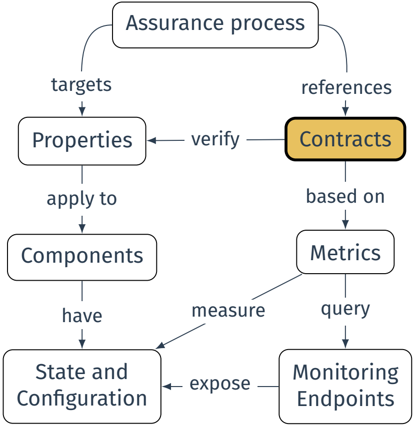

Infrastructure Assurance

Each property is associated with a contract that describes how to verify it in terms of metrics

Properties verification gives us guarantees on the behavior of the system

Assurance process continuously verifies contracts based on the collected evidence

Produce trustworthiness in infrastructures, as building blocks for distributed applications

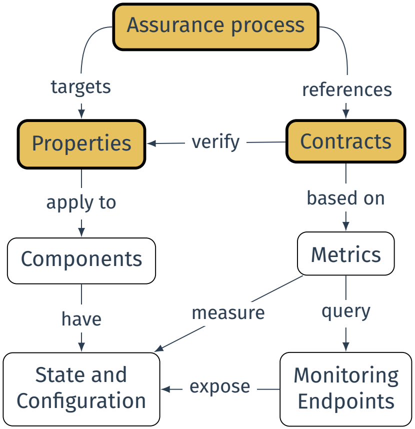

Applying assurance in modern distributed workflows

M. Anisetti, C. A. Ardagna, and F. Berto. "An assurance process for Big Data trustworthiness". In: *Future Generation Computer Systems* 146 (Sept. 2023), pp. 34–46

Assurance for Analytics Workflows

The target of the assurance process is a workflow τ composed of

- a set of tasks $t \in T$ implementing the processing workflow w
- a set of services $s \in$ implementing the ecosystem e and supporting the deployment and execution of the workflow

Our methodology is based on two abstractions

Assurance for Analytics Workflows

Abstract workflow defined via BNF as a sequence of steps (Input, Preparation, Analytics, Visualization)

Concrete workflow produced by instantiating each generic task $t \in w$ in an executable task with the form of a function call

$$
\begin{aligned}
w &::= \langle T_I \oplus P \oplus A \oplus T_V \rangle \\
P &::= \epsilon \mid T_P \mid P \oplus T_P \\
A &::= \epsilon \mid T_A \mid A \oplus T_A \\
T_I &::= \mathit{stream} \mid \mathit{fileSystem} \mid \mathit{DBMS} \mid \ldots \\
T_P &::= \mathit{cleaning} \mid \mathit{normalization} \mid \mathit{selection} \mid \ldots \\
T_A &::= \mathit{modeling} \mid \mathit{prediction} \\
T_V &::= T_I \mid T_I \oplus \mathit{visualization} \mid
\end{aligned}
$$

Assurance for Analytics Workflows

Abstract Service Ecosystem is a 5-tuple $\langle S_I, S_C, S_S, S_V, S_E \rangle$

- S_S is a set of storage services,
- S_C is a set of computational services,
- S_I is a set of ingestion services supporting data collection,
- S_V is a set of visualization services,
- S_E is a set of environmental services offering additional non-functional capabilities.

Running Example on Apache Spark

Tasks in p and $\hat{p}$		
	t	$\hat{t}$
T_I	t_1 = *fileSystem*	$\hat{t}_1$ = *loadFromHDFS()*
T_P	t_2 = *normalization*	$\hat{t}_2$ = *normalization(all)*
T_A	t_3 = *modeling(clustering)*	$\hat{t}_3$ = *k-meansModeling(k)*
T_V	t_4 = *fileSystem*	$\hat{t}_4$ = *saveToHDFS(model)*

Services in e and $\hat{e}$		
	s	$\hat{s}$
S_I	s_1 = *LoadFilesystem*	$\hat{s}_1$ = *Hadoop*
S_C	s_2 = *BatchProcessing*	$\hat{s}_2$ = *Spark*
S_C	s_3 = *Orchestration*	$\hat{s}_3$ = *Airflow*
S_S	s_4 = *StoreFilesystem*	$\hat{s}_1$ = *Hadoop*
S_V	$s_5 = \epsilon$	
S_E	s_6 = *AC: Authentication*	$\hat{s}_4$ = *Knox*
S_E	s_6 = *AC: Authorization*	$\hat{s}_5$ = *Ranger*

$$\Pi = \langle w, e \rangle \quad w = \langle t_1 \oplus t_2 \oplus t_3 \oplus t_4 \rangle$$

$$e = \langle s_1, [s_2, s_3], s_4, s_5, s_6 \rangle$$

$$I = \langle \hat{w}, \hat{e} \rangle \quad \hat{w} = \langle \hat{t}_1 \oplus \hat{t}_2 \oplus \hat{t}_3 \oplus \hat{t}_4 \rangle$$

$$\hat{e} = \langle \hat{s}_1, [\hat{s}_2, \hat{s}_3], \hat{s}_1, \epsilon, [\hat{s}_4, \hat{s}_5] \rangle$$

Requirements Annotations

The template is annotated with generic non-functional requirements to be addressed via two labeling functions

- λ assigns labels $\lambda(t_i)$ corresponding to workflow requirements in R_w
- γ assigns labels $\gamma(s_i)$ corresponding to service requirements in R_e

The instance is annotated with specific non-functional requirements to be address via two labeling functions

- θ assigns a label $\theta(\hat{t}_i)$ corresponding to workflow requirements in $R_{\hat{w}}$
- ψ assigns a label $\psi(\hat{s}_i)$ corresponding to workflow requirements in $R_{\hat{e}}$

Assurance Methodology

- The client chooses a workflow template and annotates it with generic requirements
- The annotated template is converted to an annotated workflow instance based on the concrete requirements that can be supported
- The workflow is checked using probes, verifying the annotated requirements, producing assurance confidence levels

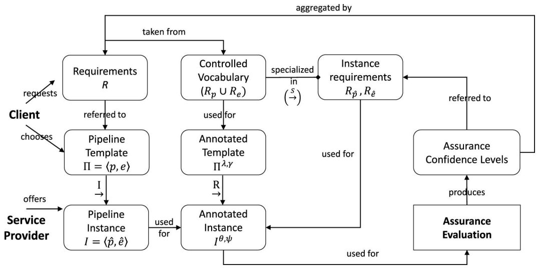

Running Example: Requirements and Assurance Probes

$\mathcal{R}$	Description
	Template requirements
r_1^λ	*Confidentiality at rest and in transit* for all the t in p
r_2^λ	*Authorization* for the $fileSystem$ task t in p
r_1^γ	*Confidentiality at rest and in transit* for all the s in e
	Pipeline Instance requirements derived from r_1^λ
r_1^θ	*No temporarily unprotected data storage* for all the $\hat{t} \in \hat{p}$
r_2^θ	*Avoid connection to external services* for all the $\hat{t} \in \hat{p}$
r_3^θ	*Avoid use of vulnerable code/libraries* for all the $\hat{t} \in \hat{p}$
r_4^θ	*Pipeline integrity* checking the correct ordering of tasks $\hat{t} \in \hat{p}$
	Pipeline Instance requirements derived from r_2^λ
r_5^θ	*Check Authorization* for ingestion task $\hat{t}_1 \in \hat{p}$
r_6^θ	*Check ownerships at rest* for visualization tasks $\hat{t}_4 \in \hat{p}$
	Ecosystem Instance requirements derived from r_1^γ
r_1^ψ	*Encrypted HDFS* for the $\hat{s}_1 \in \hat{e}$
r_2^ψ	*Inter-node communication security* for the $\hat{s}_1$ and $\hat{s}_2 \in \hat{e}$
r_3^ψ	*Orchestrator Confidentiality* for the $\hat{s}_3 \in \hat{e}$
r_4^ψ	*Communication channel security* for the $\hat{s}_4 \in \hat{e}$
r_5^ψ	*Authentication-enabled* for the $\hat{s}_4 \in \hat{e}$
r_6^ψ	*Authorization policies-enabled* for the $\hat{s}_5 \in \hat{e}$
r_7^ψ	*Vulnerability check* for all the services $\hat{s} \in \hat{e}$

Probe Name	Description
Task probes	
Code Inspection	search instructions/pattern
Dependency Check	find vulnerable dependencies
Code Vulnerability Check	search vulnerable code
Lineage	verify sequence of actions using logs
Pipeline probes	
Parameters Check	check tasks' actual parameters
Orchestration Check	check the workflow structure
Service probes	
Vulnerability Check	search for vulnerability
Configuration Check	parse and verify configuration
Ecosystem probes	
Infrastructure	targets lower layers such as OS (see [21])
General purposes probes	
Testing	perform specific test cases on a target
Monitoring	monitor a target ore a time frame

Running Example: Assurance Evaluation

Workflow tasks $\hat{t} \in \hat{T}$				
$\hat{t}$	$\mathcal{R}$	$P(r,\tau)$	$E(EV,r)$	$A_{\tau,\nabla}$
$\hat{t}_1$	r_1^θ	$P_1(r_1^\theta,\hat{t}_1)$	$[1.0]$	1.0
	r_2^θ	$P_2(r_2^\theta,\hat{t}_1), P_3(r_2^\theta,\hat{t}_1), P_4(r_2^\theta,\hat{t}_1)$	$[1.0, 1.0, 1.0]$	1.0
	r_3^θ	$P_5(r_3^\theta,\hat{t}_1), P_6(r_3^\theta,\hat{t}_1)$	$[0.75, 1.0]$	0.88
	r_5^θ	$P_8(r_5^\theta,\hat{t}_1)$	$[1.0]$	1.0
$\hat{t}_2$	r_1^θ	$P_1(r_1^\theta,\hat{t}_2)$	$[1.0]$	1.0
	r_2^θ	$P_2(r_2^\theta,\hat{t}_2), P_3(r_2^\theta,\hat{t}_2), P_4(r_2^\theta,\hat{t}_2)$	$[1.0, 1.0, 1.0]$	1.0
	r_3^θ	$P_5(r_3^\theta,\hat{t}_2), P_6(r_3^\theta,\hat{t}_1)$	$[0.75, 1.0]$	0.88
$\hat{t}_3$	r_1^θ	$P_1(r_1^\theta,\hat{t}_3)$	$[1.0]$	1.0
	r_2^θ	$P_2(r_2^\theta,\hat{t}_3), P_3(r_2^\theta,\hat{t}_3), P_4(r_2^\theta,\hat{t}_3)$	$[1.0, 1.0, 1.0]$	1.0
	r_3^θ	$P_5(r_3^\theta,\hat{t}_3), P_6(r_3^\theta,\hat{t}_1)$	$[0.75, 1.0]$	0.88
$\hat{t}_4$	r_1^θ	$P_1(r_1^\theta,\hat{t}_4)$	$[1.0]$	1.0
	r_2^θ	$P_2(r_2^\theta,\hat{t}_4), P_3(r_2^\theta,\hat{t}_4), P_4(r_2^\theta,\hat{t}_4)$	$[1.0, 1.0, 0.0]$	0.66
	r_3^θ	$P_5(r_3^\theta,\hat{t}_4), P_6(r_3^\theta,\hat{t}_1)$	$[0.75, 1.0]$	0.88
	r_6^θ	$P_9(r_6^\theta,\hat{t}_4)$	$[1.0]$	1.0
$\hat{p}$	r_4^θ	$P_7(r_4^\theta,\hat{p})$	$[1.0]$	1.0

Ecosystem services $\hat{s} \in \hat{S}$				
$\hat{s}$	$\mathcal{R}$	$P(r,\tau)$	$E(EV,r)$	$A_{\tau,r}$
$\hat{s}_1$	r_1^ψ	$P_{10}(r_1^\psi,\hat{s}_1)$	$[0.1]$	0.1
	r_2^ψ	$P_{11}(r_2^\psi,\hat{s}_1)$	$[0.1]$	0.1
$\hat{s}_2$	r_2^ψ	$P_{12}(r_2^\psi,\hat{s}_2)$	$[0.1]$	0.1
$\hat{s}_3$	r_4^ψ	$P_{13}(r_4^\psi,\hat{s}_3)$	$[0.1]$	0.1
$\hat{s}_4$	r_4^ψ	$P_{14}(r_4^\psi,\hat{s}_4)$	$[1.0]$	1.0
	r_5^ψ	$P_{15}(r_5^\psi,\hat{s}_4)$	$[1.0]$	1.0
$\hat{s}_5$	r_6^ψ	$P_{16}(r_6^\psi,\hat{s}_5)$	$[1.0]$	1.0
$\hat{s}_5$	r_7^ψ	$P_{17}(r_7^\psi,\hat{s}_5)$	$[0.57]$	0.57

Assurance levels $A_{\tau,\gamma}$ = Frequency of positive evaluation multiplied by the average of the positive evaluations

Running Example: Assurance Evaluation

Workflow tasks $\hat{t} \in \hat{T}$				
$\hat{t}$	$\mathcal{R}$	$P(r,\tau)$	$E(EV,r)$	$A_{\tau,\nabla}$
$\hat{t}_1$	r_1^θ	$P_1(r_1^\theta,\hat{t}_1)$	[1.0]	1.0
	r_2^θ	$P_2(r_2^\theta,\hat{t}_1), P_3(r_2^\theta,\hat{t}_1), P_4(r_2^\theta,\hat{t}_1)$	[1.0, 1.0, 1.0]	1.0
	r_3^θ	$P_5(r_3^\theta,\hat{t}_1), P_6(r_3^\theta,\hat{t}_1)$	[0.75, 1.0]	0.88
	r_5^θ	$P_8(r_5^\theta,\hat{t}_1)$	[1.0]	1.0
$\hat{t}_2$	r_1^θ	$P_1(r_1^\theta,\hat{t}_2)$	[1.0]	1.0
	r_2^θ	$P_2(r_2^\theta,\hat{t}_2), P_3(r_2^\theta,\hat{t}_2), P_4(r_2^\theta,\hat{t}_2)$	[1.0, 1.0, 1.0]	1.0
	r_3^θ	$P_5(r_3^\theta,\hat{t}_2), P_6(r_3^\theta,\hat{t}_1)$	[0.75, 1.0]	0.88
$\hat{t}_3$	r_1^θ	$P_1(r_1^\theta,\hat{t}_3)$	[1.0]	1.0
	r_2^θ	$P_2(r_2^\theta,\hat{t}_3), P_3(r_2^\theta,\hat{t}_3), P_4(r_2^\theta,\hat{t}_3)$	[1.0, 1.0, 1.0]	1.0
	r_3^θ	$P_5(r_3^\theta,\hat{t}_3), P_6(r_3^\theta,\hat{t}_1)$	[0.75, 1.0]	0.88
$\hat{t}_4$	r_1^θ	$P_1(r_1^\theta,\hat{t}_4)$	[1.0]	1.0
	r_2^θ	$P_2(r_2^\theta,\hat{t}_4), P_3(r_2^\theta,\hat{t}_4), P_4(r_2^\theta,\hat{t}_4)$	[1.0, 1.0, 0.0]	0.66
	r_3^θ	$P_5(r_3^\theta,\hat{t}_4), P_6(r_3^\theta,\hat{t}_1)$	[0.75, 1.0]	0.88
	r_6^θ	$P_9(r_6^\theta,\hat{t}_4)$	[1.0]	1.0
$\hat{p}$	r_4^θ	$P_7(r_4^\theta,\hat{p})$	[1.0]	1.0

Ecosystem services $\hat{s} \in \hat{S}$				
$\hat{s}$	$\mathcal{R}$	$P(r,\tau)$	$E(EV,r)$	$A_{\tau,r}$
$\hat{s}_1$	r_1^ψ	$P_{10}(r_1^\psi,\hat{s}_1)$	[0.1]	0.1
	r_2^ψ	$P_{11}(r_2^\psi,\hat{s}_1)$	[0.1]	0.1
$\hat{s}_2$	r_2^ψ	$P_{12}(r_2^\psi,\hat{s}_2)$	[0.1]	0.1
$\hat{s}_3$	r_4^ψ	$P_{13}(r_4^\psi,\hat{s}_3)$	[0.1]	0.1
$\hat{s}_4$	r_4^ψ	$P_{14}(r_4^\psi,\hat{s}_4)$	[1.0]	1.0
	r_5^ψ	$P_{15}(r_5^\psi,\hat{s}_4)$	[1.0]	1.0
$\hat{s}_5$	r_6^ψ	$P_{16}(r_6^\psi,\hat{s}_5)$	[1.0]	1.0
$\hat{s}_5$	r_7^ψ	$P_{17}(r_7^\psi,\hat{s}_5)$	[0.57]	0.57

Assurance levels $A_{\tau,\gamma}$ = Frequency of positive evaluation multiplied by the average of the positive evaluations

Running Example: Assurance Evaluation

Performance of the assurance process on the example workflow in different scenarios:

- Contextual changes (CC)
- Workflow changes (PC)
- Ecosystem changes (EC)
- Instance changes (IC)
- No changes (NC)

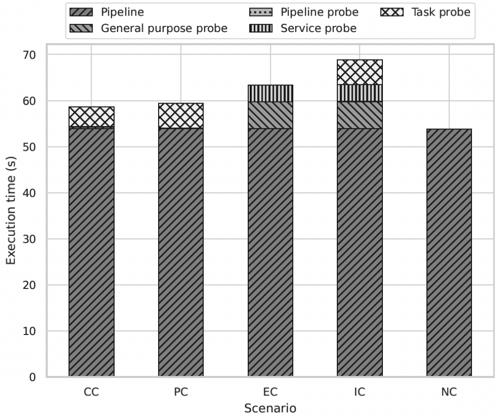

Assurance-aware deployments in the continuum

M. Anisetti, F. Berto, and R. Bondaruc. “QoS-Aware Deployment of Service Compositions in 5G-Empowered Edge-Cloud Continuum”. In: *2023 IEEE 16th International Conference on Cloud Computing (CLOUD)*. ISSN: 2159-6190. IEEE, July 2023, pp. 471–478

Deployment in Network and Computing infrastructures

The non-functional properties of the deployment targets influence the deployed application's properties

The heterogeneity of the continuum ecosystem reflects on the deployment targets

Exploiting infrastructure peculiarities to support NFP-based Service Level Agreements while deploying workflows in the continuum

Deployment Methodology

- Continuum Service Providers generate an annotated graph of their facilities

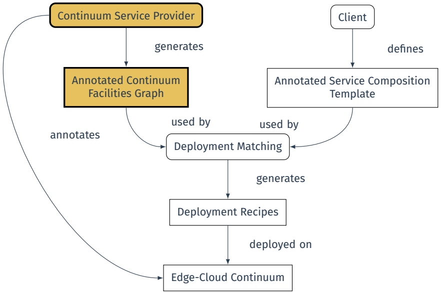

Deployment Methodology

- Continuum Service Providers generate an annotated graph of their facilities
- The client defines an annotated template for service composition

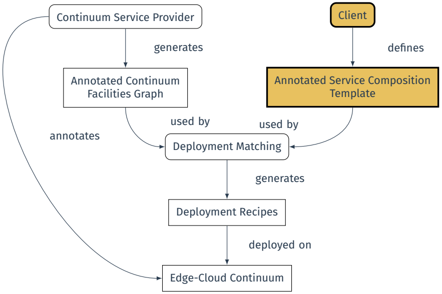

Deployment Methodology

- Continuum Service Providers generate an annotated graph of their facilities
- The client defines an annotated template for service composition
- The deployment matching process searches for a suitable match of services and deployment facilities

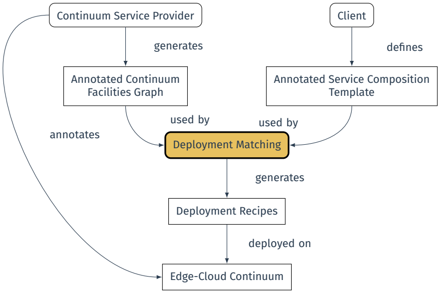

Deployment Methodology

- Continuum Service Providers generate an annotated graph of their facilities
- The client defines an annotated template for service composition
- The deployment matching process searches for a suitable match of services and deployment facilities
- If a match is found, the system generates deployment recipes for the Edge-Cloud continuum

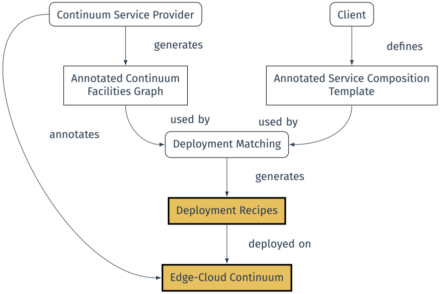

Deployment Matching

- The client identifies a list of services and annotate them with non-functional requirements

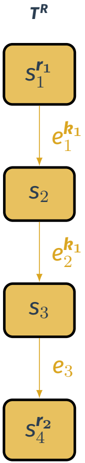

Deployment Matching

- The client identifies a list of services and annotate them with non-functional requirements
- The Continuum Service Provider describe the available deployment facilities and their non-functional capabilities

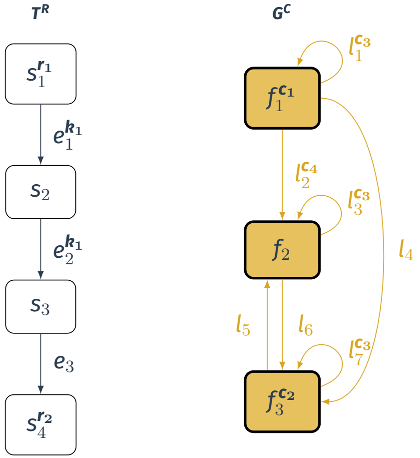

Deployment Matching

- The client identifies a list of services and annotate them with non-functional requirements
- The Continuum Service Provider describe the available deployment facilities and their non-functional capabilities
- The deployment matching process finds a suitable configuration considering services requirements and facilities capabilities

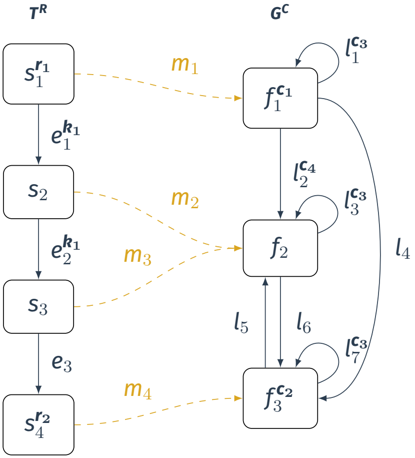

Handling assurance failure

- The node f_2 loses a non-functional capability that is required by S_2 and S_3

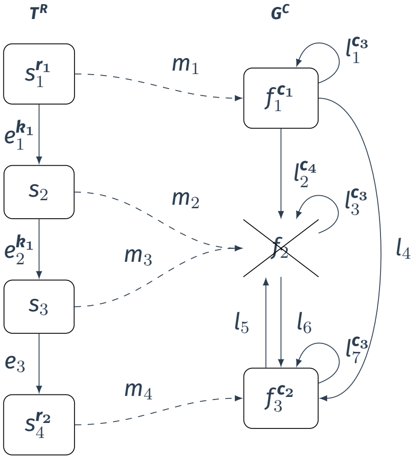

Handling assurance failure

- The node f_2 loses a non-functional capability that is required by S_2 and S_3
- The deployment process finds a new suitable matching and generates new deployment recipes

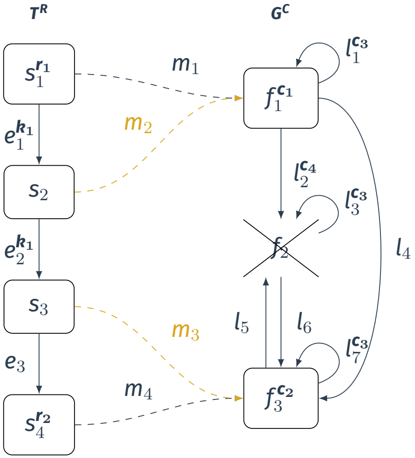

Performance Evaluation

Performance of the Matching Process varying the number of services and facilities

Services requirements and Facilities capabilities are synthetically generated

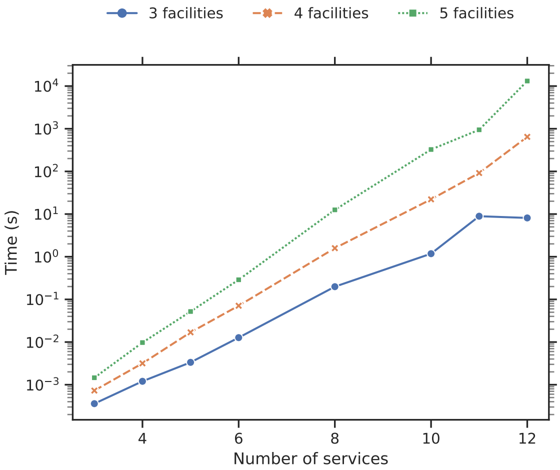

Workflow Deployment in the Continuum

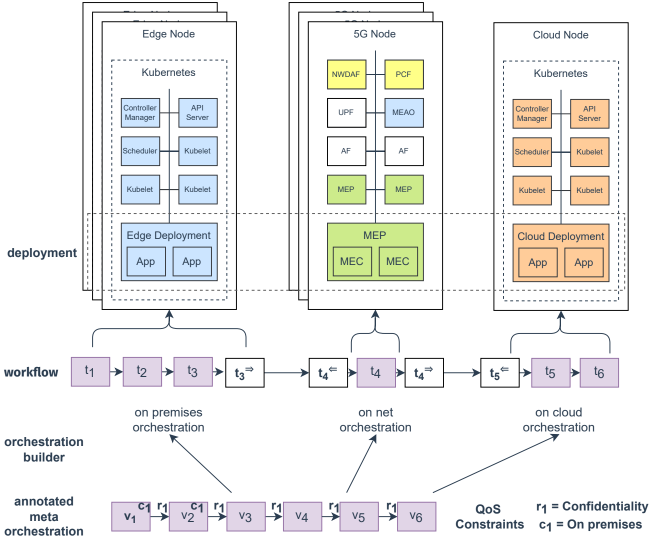

Assurance for Content Distribution Networks

M. Anisetti, C. A. Ardagna, F. Berto, and E. Damiani. “A Security Certification Scheme for Information-Centric Networks”. en. In: *IEEE Trans. Netw. Serv. Manage.* 19.3 (Sept. 2022), pp. 2397–2408

M. Anisetti, C. A. Ardagna, F. Berto, and E. Damiani. “Security Certification Scheme for Content-centric Networks”. In: *2021 IEEE International Conference on Services Computing (SCC)*. IEEE, Sept. 2021, pp. 203–212

Assurance for Content Distribution Networks

CDN based on Named Data Networking

Alternative network stack to TCP/IP

Advanced content cache system focused on in-protocol caching and security

Contents security and privacy by default

Can be used as an application layer in other protocols

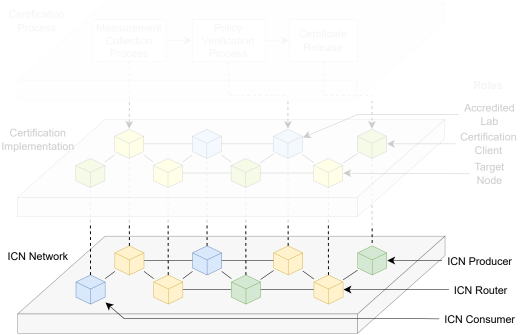

Assurance for Content Distribution Networks

Target Nodes expose metrics as NDN contents

Accredited Labs implement the assurance process

- collect evidence from Target Nodes
- verify properties of Target Nodes on-demand of Certification Clients
- issue signed certificates containing the verification results

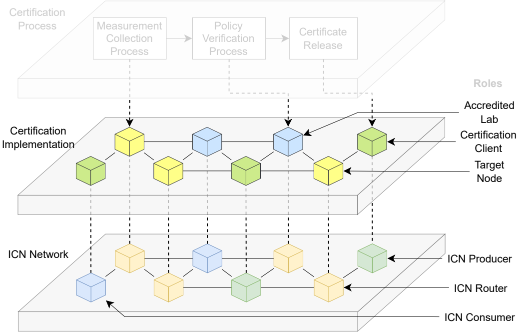

Assurance for Content Distribution Networks

Previously issued certificates by trusted AL can be reused to speed-up the verification process

Accredited Labs can collaborate sharing signed certificates and evidence

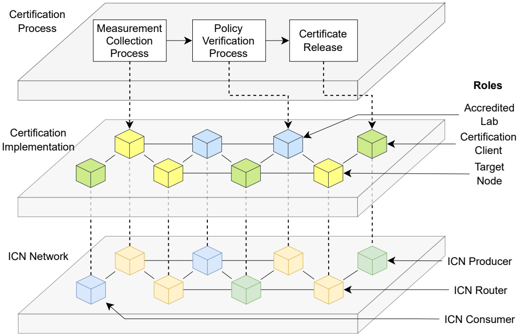

Centralized VS Decentralized Certification

The peculiar capabilities of NDN allowed us to develop two certification solutions

A centralized certification process, which is the standard implementation

A collaborative and decentralized certification process using NDN caching model

ALs share evidence and certificates with other ALs in the network, maintaining confidentiality and non-repudiability

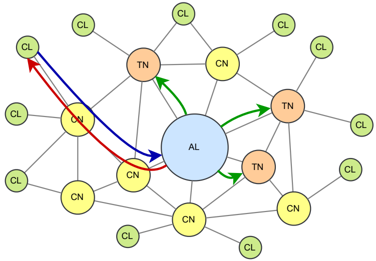

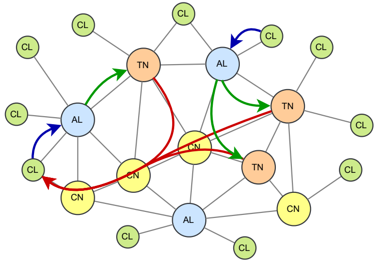

Performance evaluation

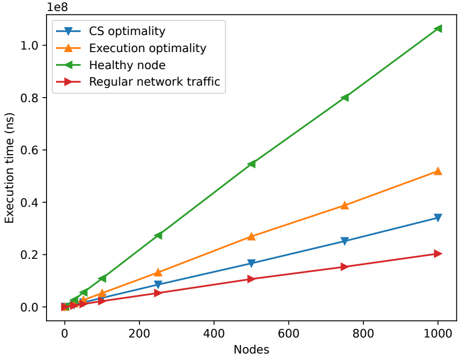

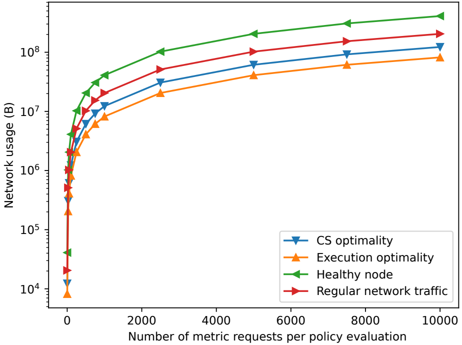

TIM Industrial Scenario

F. Berto, C. Ardagna, M. Torrente, D. Manenti, E. Ferrari, A. Calcante, R. Oberti, C. Fra', and L. Ciani. "A 5G-IoT enabled Big Data infrastructure for data-driven agronomy". In: *2022 IEEE Globecom Workshops (GC Wkshps)*. Rio de Janeiro, Brazil: IEEE, Dec. 2022, pp. 588–594

F. Berto, C. Ardagna, M. Torrente, D. Manenti, E. Ferrari, A. Calcante, R. Oberti, C. Fra', and L. Ciani. "A 5G-IoT enabled Big Data infrastructure for data-driven agronomy". In: *The 1st Italian Conference on Big Data and Data Science (ITADATA)*. Milan, Italy, Oct. 2022

TIM Industrial Scenario: Workflow

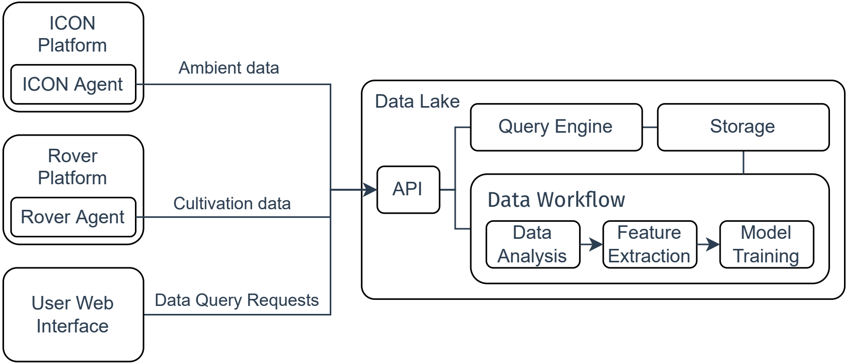

TIM Industrial Scenario: Performance

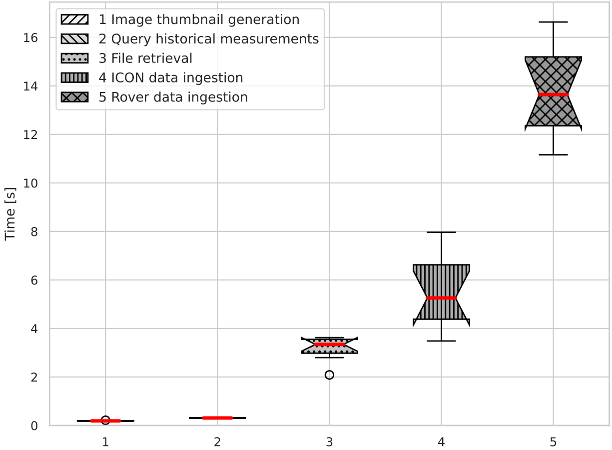

Conclusions

Introduction of a novel fully-functional 5G-empowered continuum architecture

The infrastructure assurance methodology has been developed and verified in integrated industrial-ready scenario

Proposed a solution for supporting intensive computation and smart deployment

Assurance for data intensive workflows exploiting big data platforms

Novel seamless deployment solution for workflows in the continuum

Future Research Directions

Intent-driven continuum empowered by assurance

Satellite based continuum

Lightweight assurance for unreliable networks

AI-based assurance methodology

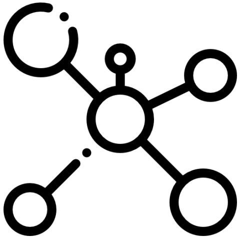

Thanks

The work in this thesis resulted in

- 2 journal article (Q1 according to Scimago)
- 7 conference papers
- 1 chapters in international books

References I

[1] M. Anisetti, C. A. Ardagna, and F. Berto. "An assurance process for Big Data trustworthiness". In: *Future Generation Computer Systems* 146 (Sept. 2023), pp. 34–46.

[2] M. Anisetti, C. A. Ardagna, F. Berto, and E. Damiani. "A Security Certification Scheme for Information-Centric Networks". en. In: *IEEE Trans. Netw. Serv. Manage.* 19.3 (Sept. 2022), pp. 2397–2408.

[3] M. Anisetti, C. A. Ardagna, F. Berto, and E. Damiani. "Security Certification Scheme for Content-centric Networks". In: *2021 IEEE International Conference on Services Computing (SCC)*. IEEE, Sept. 2021, pp. 203–212.

[4] M. Anisetti, N. Bena, F. Berto, and G. Jeon. "A DevSecOps-based Assurance Process for Big Data Analytics". en. In: *2022 IEEE International Conference on Web Services (ICWS)*. Barcelona, Spain: IEEE, July 2022, pp. 1–10.

[5] M. Anisetti, F. Berto, and M. Banzi. "Orchestration of data-intensive pipeline in 5G-enabled Edge Continuum". In: *2022 IEEE World Congress on Services (SERVICES)*. ISSN: 2642-939X. IEEE Computer Society, July 2022, pp. 2–10.

References II

[6] M. Anisetti, F. Berto, and R. Bondaruc. "QoS-Aware Deployment of Service Compositions in 5G-Empowered Edge-Cloud Continuum". In: *2023 IEEE 16th International Conference on Cloud Computing (CLOUD)*. ISSN: 2159-6190. IEEE, July 2023, pp. 471–478.

[7] C. A. Ardagna, E. Damiani, and F. Berto. "Script Language Security". In: *Encyclopedia of Cryptography, Security and Privacy*. Ed. by S. Jajodia, P. Samarati, and M. Yung. Berlin, Heidelberg: Springer Berlin Heidelberg, Sept. 2023, pp. 1–3.

[8] F. Berto, C. Ardagna, M. Torrente, D. Manenti, E. Ferrari, A. Calcante, R. Oberti, C. Fra', and L. Ciani. "A 5G-IoT enabled Big Data infrastructure for data-driven agronomy". In: *The 1st Italian Conference on Big Data and Data Science (ITADATA)*. Milan, Italy, Oct. 2022.

[9] F. Berto, C. Ardagna, M. Torrente, D. Manenti, E. Ferrari, A. Calcante, R. Oberti, C. Fra', and L. Ciani. "A 5G-IoT enabled Big Data infrastructure for data-driven agronomy". In: *2022 IEEE Globecom Workshops (GC Wkshps)*. Rio de Janeiro, Brazil: IEEE, Dec. 2022, pp. 588–594.

[10] H. Badir et al. "Where We are in Handling IoT and Robotic's Data for Agro-ecology Applications? An Architectural View". In: *11th International Conference on New Technologies, Artificial Intelligence and Smart Data (INTIS2023)*. Tangier, Morocco, May 2023.