Non-Functional Certification of Modern Distributed Systems: A Research Manifesto

Claudio A. Ardagna Nicola Bena

firstname.lastname@unimi.it

https://homes.di.unimi.it/*lastname*/

Department of Computer Science, Università degli Studi di Milano, Milan, Italy

The work was partially supported by the projects *i)* MUSA - Multilayered Urban Sustainability Action - project, funded by the European Union - NextGenerationEU, under the National Recovery and Resilience Plan (NRRP) Mission 4 Component 2 Investment Line 1.5: Strengthening of research structures and creation of R&D “innovation ecosystems”, set up of “territorial leaders in R&D” (CUP G43C22001370007, Code ECS00000037); *ii)* SERICS (PE00000014) under the NRRP MUR program funded by the EU - NextGenerationEU.

Scenario

Modern distributed systems

- confluence of cloud-edge-IoT
- multi-layer structure
- ML-based services and infrastructure
- dynamic, non-deterministic, and unpredictable behavior

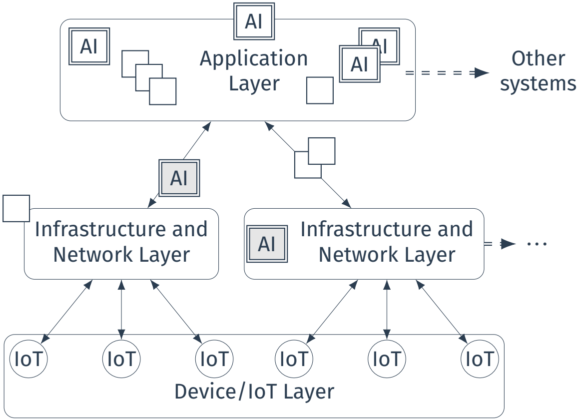

Scenario

Modern distributed systems

- impact of AI by 2030: $13 trillion[a]
- number of connected devices by 2023: 29.3 bln[b]
- economic impact of cloud-edge-IoT by 2025: $2.7–6.2 trillion[c]

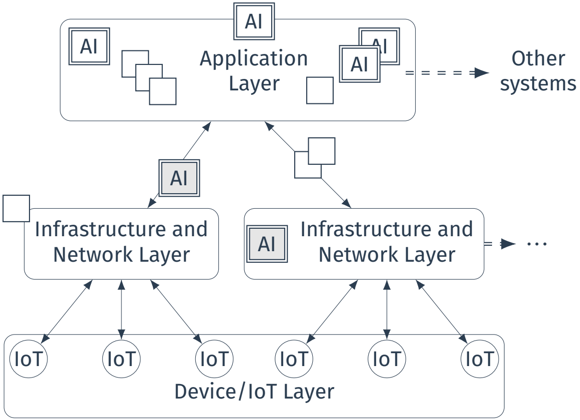

[a] Source: McKinsey
[b] Source: Cisco
[c] Source: McKinsey

Scenario

Modern distributed systems

- increasing pervasiveness
- increasing risk for security, safety, and privacy
- lack of trustworthiness
  - full/partial lose of control on data/applications
  - lack of evidence about service operation and effectiveness

$\Longrightarrow$ assurance based-certification to the rescue

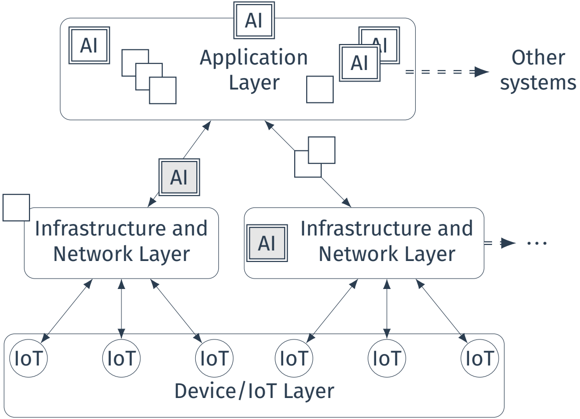

Certification

Certification scheme details the certification process verifying that a target system behaves as expected and demonstrates one or more non-functional properties

Certification

Certification scheme details the certification process verifying that a target system behaves as expected and demonstrates one or more non-functional properties

Software certification	Service certification	Cloud certification
• one time • lengthy and heavyweight	• mostly one time • model-based generation of test cases	• continuous and incremental • composition • semi-automatic or automatic

Certification

Certification scheme details the certification process, according to

- non-functional property
- target of certification
- evidence collection model
- certification model
- evidence
- certificate

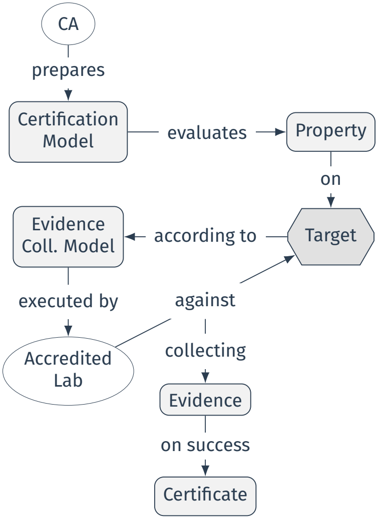

Certification

Certification scheme details the certification process, according to

- non-functional property
- target of certification
- evidence collection model
- certification model
- evidence
- certificate

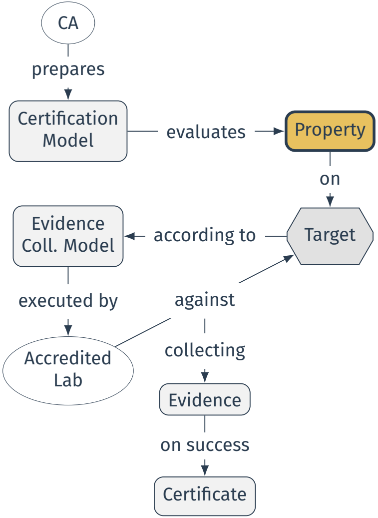

Certification

Certification scheme details the certification process, according to

- non-functional property
- target of certification
- evidence collection model
- certification model
- evidence
- certificate

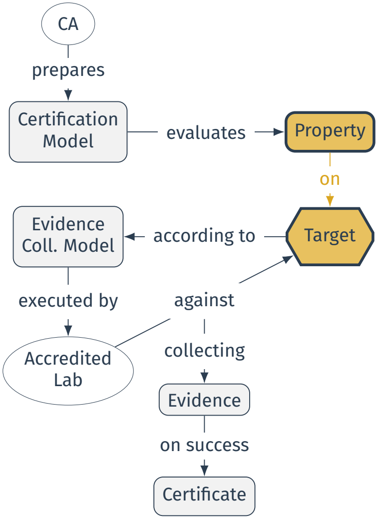

Certification

Certification scheme details the certification process, according to

- non-functional property
- target of certification
- evidence collection model
- certification model
- evidence
- certificate

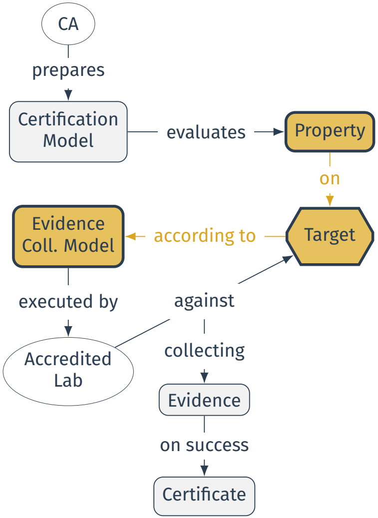

Certification

Certification scheme details the certification process, according to

- non-functional property
- target of certification
- evidence collection model
- certification model
- evidence
- certificate

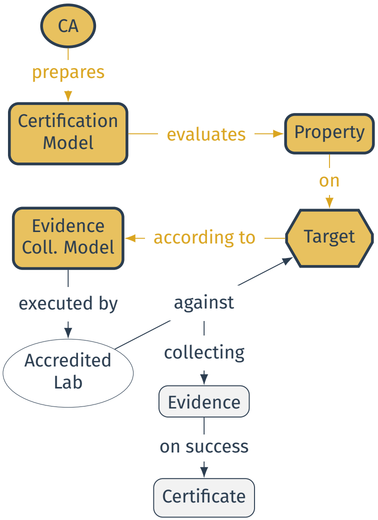

Certification

Certification scheme details the certification process, according to

- non-functional property
- target of certification
- evidence collection model
- certification model
- evidence
- certificate

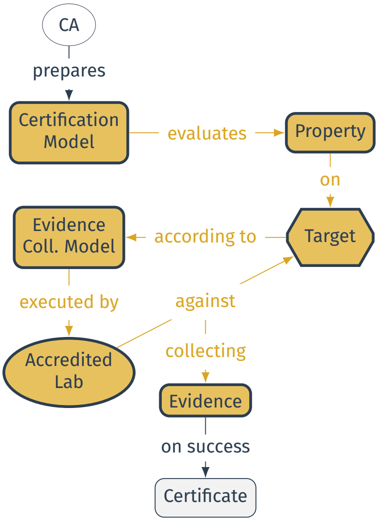

Certification

Certification scheme details the certification process, according to

- non-functional property
- target of certification
- evidence collection model
- certification model
- evidence
- certificate

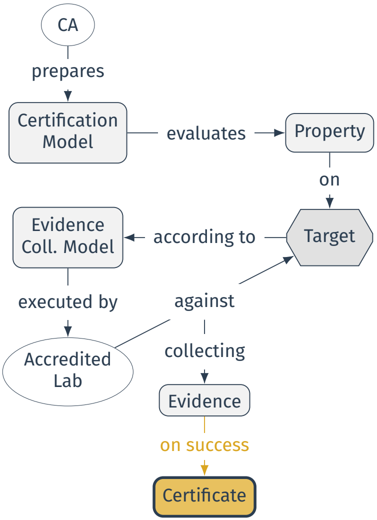

Certification: Example

Property Reliability $p_{rel}=(\widehat{p}_{rel}, \{replicas=2, replica\ zones=2\})$, where

- $\widehat{p}_{rel}$ is the name of the property (reliability)
- $replicas=2$ and $replica\ zones=2$ are attributes refining it

Target $s_1=\{c_{db}, c_{api}, c_{cross}\}$ is a set of components

Evidence collection model {`get-orchestrator`, `check-replicas`, `check-zones`}, where

- `get-orchestrator` checks the availability of the expected orchestrator and its configurations
- `check-replicas`, `check-zones` checks the deployment of the service

Certification

Software certification

- one time
- lengthy and heavyweight

Service certification

- mostly one time
- model-based generation of test cases

Cloud certification

- continuous and incremental
- composition
- semi-automatic or automatic

Can we adapt existing techniques to be applicable to modern distributed systems as we did in the past?

Certification

Software certification

- one time
- lengthy and heavyweight

Service certification

- mostly one time
- model-based generation of test cases

Cloud certification

- continuous and incremental
- composition
- semi-automatic or automatic

Can we adapt existing techniques to be applicable to modern distributed systems as we did in the past? NO!

Our Manifesto (1)

Our manifesto identifies the challenges, the corresponding research directions, and an implementation timeline, towards low-cost, trustworthy certification techniques at the basis of trustworthy modern distributed systems

Our Manifesto (2)

Research direction	Challenge	Timeline
RD1: Non-functional property definition	**C1.1**: Property definition	M
	C2.1: Multi-layer service composition	S, M
	C4.2: Certification-based system life cycle	M, L
RD2: Behavior-based certification	**C1.2**: Target modeling	M
	C2.1: Multi-layer service composition	S, M
	C2.3: Dishonest behavior	M
	C4.1: Increase automation	S, M
	C4.2: Certification-based system life cycle	M, L
RD3: Trustworthy evidence management	**C2.2**: Evidence lineage	M
	C4.3: Reduce reliance on blind trust	M

Research direction	Challenge	Timeline
RD4: Certification for ML	**C1.1**: Property definition	M
	C1.2: Target modeling	M
	C2.1: Multi-layer service composition	S, M
	C3.1: Property and target definition	M, L
	C3.2: Certification process modeling	M, L
	C3.3: ML pipelines	L
RD5: ML-based automation	**C1.1**: Property definition	M
	C1.2: Target modeling	M
	C2.3: Dishonest behavior	M
	C4.1: Increase automation	S, M
RD6: *DevCertOps* and beyond	**C1.3**: Integration of development and certification processes	M, L
	C2.1: Multi-layer service composition	S, M
	C4.1: Increase automation	S, M
	C4.2: Certification-based system life cycle	M, L

Research Direction: Behavior-Based Non-Functional Property

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Challenges

- Existing non-functional properties do not model system evolution over time
  - cannot be easily integrated with system life cycle
- Certification evaluation still relies on precise and human-made system modeling
  - but system boundaries are dynamic (lack of automation)
- Evidence management and collection still rely on static processes
  - no system behavior

Research Direction: Behavior-Based Non-Functional Property

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

- Flexible definition of properties based on system behavior
- Model expected system behavior and compare the retrieved behavior against it in a continuous fashion and adapting to system changes
- Trustworthy, human-readable evidence management and collection

M. Anisetti, C. A. Ardagna, and N. Bena. "Multi-Dimensional Certification of Modern Distributed Systems". In: *IEEE TSC* (2022); M. Anisetti, C. A. Ardagna, E. Damiani, and G. Polegri. "Test-Based Security Certification of Composite Services". In: *ACM TWEB* 13.1 (2019)

Research Direction: Behavior-Based Non-Functional Property

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

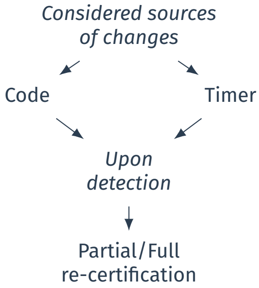

State of art

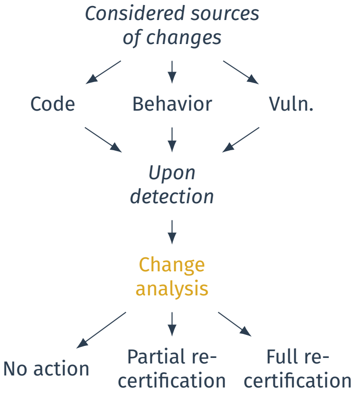

Beyond the State of art

Research Direction: Certification of ML

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Challenges

Certification schemes are designed for deterministic systems that can be inspected or tested

- cannot model and certify a ML-based service whose behavior is unpredictable
- cannot be limited to run-time model evaluation

Research Direction: Certification of ML

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Novel building blocks for the certification of ML-based systems

- novel definition of non-functional property
- evaluation based on observed predictions or explainability
- along the complete ML pipeline and towards the complete ML-based system

Research Direction: Certification of ML

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Multi-factor certification: jointly evaluate the ML-based service across multiple factors

- data: dataset used for training
- process: training process
- model: run-time model

Each factor has its own independent life cycle

M. Anisetti, C. A. Ardagna, N. Bena, and E. Damiani. "Towards Certification of Machine Learning-Based Distributed Systems". In: *arXiv preprint arXiv:2305.16822* (2023)

Research Direction: Certification of ML

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Ex.: certification of a malware detector trained on real-world and synthetic data (GAN) for property robustness

- data: verify that the distribution of the synthetic dataset is close enough to that of the real dataset
- process: verify that adversarial training is used to prevent adversarial (inference-time) attacks
- model: verify that adversarial data points are ineffective

M. Anisetti, C. A. Ardagna, N. Bena, and E. Damiani. "Towards Certification of Machine Learning-Based Distributed Systems". In: *arXiv preprint arXiv:2305.16822* (2023); M. Anisetti, C. A. Ardagna, N. Bena, V. Giandomenico, and G. Gianini. "Lightweight Behavior-Based Malware Detection". In: *Proc. of MEDES 2023*. To appear. Heraklion, Greece, May 2023

Research Direction: Certification of ML

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Ex.: certification of a malware detector trained on real-world and synthetic data (GAN) for property robustness

- data: verify that the distribution of the synthetic dataset is close enough to that of the real dataset
  - true
- process: verify that adversarial training is used to prevent adversarial (inference-time) attacks
  - false
- model: verify that adversarial data points are ineffective
  - false

M. Anisetti, C. A. Ardagna, N. Bena, and E. Damiani. "Towards Certification of Machine Learning-Based Distributed Systems". In: *arXiv preprint arXiv:2305.16822* (2023); M. Anisetti, C. A. Ardagna, N. Bena, V. Giandomenico, and G. Gianini. "Lightweight Behavior-Based Malware Detection". In: *Proc. of MEDES 2023*. To appear. Heraklion, Greece, May 2023

Research Direction: ML-Based Automation

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Challenges

Certification still relies on error-prone and expensive manual activities

- lack of automation
- reliance on precise and human-made system modeling
  - but system boundaries are dynamic

Research Direction: ML-Based Automation

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Use ML to boost the automation of certification activities

- automatically infer target system's behavior and properties
- automatically derive the corresponding evaluation

Research Direction: DevCertOps and Beyond

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Challenges

Certification is still seen as one-time, post-deployment activity

- lack of tight integration within the system life cycle
- lack of *usage* of certificates after their issuing

Research Direction: DevCertOps and Beyond

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

Integration of system development life cycle and certification life cycle

- certify all development/deployment artifacts
  - *shift certification to the left*
- certification part of the process driving system evolution

Research Direction: DevCertOps and Beyond

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

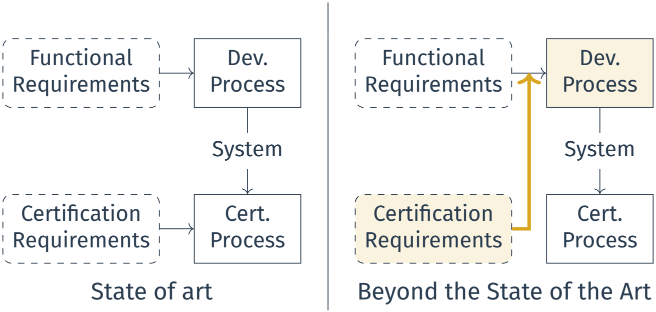

C. A. Ardagna, N. Bena, and R. M. de Pozuelo. "Bridging the Gap Between Certification and Software Development". In: *Proc. of ARES 2022*. Vienna, Austria, Aug. 2022

Research Direction: DevCertOps and Beyond

- **RD1**: non-functional property definition
- **RD2**: behavior-based certification
- **RD3**: trustworthy evidence management
- **RD4**: certification of ML
- **RD5**: ML-based automation
- **RD6**: *DevCertOps* and beyond

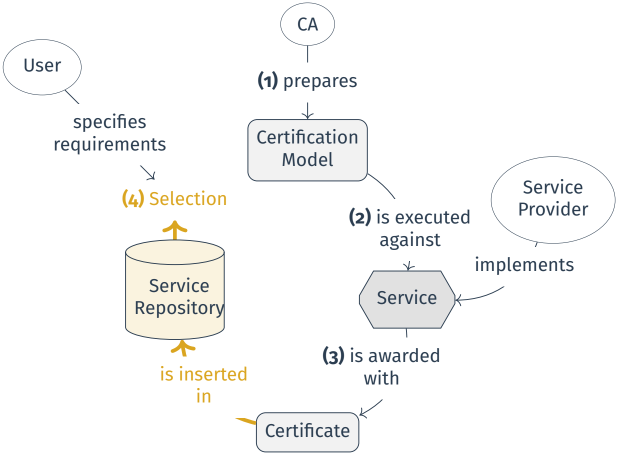

M. Anisetti, C. A. Ardagna, and N. Bena. "Multi-Dimensional Certification of Modern Distributed Systems". In: *IEEE TSC* (2022)

Conclusions

Certification is a pressing need
$\Longrightarrow$ certification as the preferred way to increase system trustworthiness

Existing static techniques make it practically unusable and with low value for modern distributed systems

- technical challenges and research directions in this roadmap
- policy makers and regulators have to do their part
  - e.g., legislative initiatives in EU (ENISA mandate on cybersecurity certification framework, The AI Act)

Main disruption: Machine Learning

- Certification for Machine Learning (*Cert4ML*)
- Machine Learning for Certification (*ML4Cert*)

Thanks! Questions?

Claudio A. Ardagna Nicola Bena

firstname.lastname@unimi.it

https://homes.di.unimi.it/*lastname*/

Department of Computer Science, Università degli Studi di Milano, Milan, Italy

The work was partially supported by the projects *i)* MUSA - Multilayered Urban Sustainability Action - project, funded by the European Union - NextGenerationEU, under the National Recovery and Resilience Plan (NRRP) Mission 4 Component 2 Investment Line 1.5: Strengthening of research structures and creation of R&D "innovation ecosystems", set up of "territorial leaders in R&D" (CUP G43C22001370007, Code ECS00000037); *ii)* SERICS (PE00000014) under the NRRP MUR program funded by the EU - NextGenerationEU.

References I

[1] M. Anisetti, C. A. Ardagna, F. Gaudenzi, and E. Damiani. "A Continuous Certification Methodology for DevOps". In: *Proc. of MEDES 2019*. Limassol, Cyprus, Nov. 2019.

[2] M. Anisetti, C. A. Ardagna, A. Balestrucci, N. Bena, E. Damiani, and C. Y. Yeun. "On the Robustness of Ensemble-Based Machine Learning Against Data Poisoning". In: *IEEE TSUSC* (2023). To appear.

[3] M. Anisetti, C. A. Ardagna, and N. Bena. "Continuous Certification of Non-Functional Properties Across System Changes". In: *ICSOC 2023*. Under review. 2023.

[4] M. Anisetti, C. A. Ardagna, and N. Bena. "Multi-Dimensional Certification of Modern Distributed Systems". In: *IEEE TSC* (2022).

[5] M. Anisetti, C. A. Ardagna, N. Bena, and R. Bondaruc. "Towards an Assurance Framework for Edge and IoT Systems". In: *Proc. of IEEE EDGE 2021*. Guangzhou, China, Dec. 2021.

[6] M. Anisetti, C. A. Ardagna, N. Bena, and E. Damiani. "Towards Certification of Machine Learning-Based Distributed Systems". In: *arXiv preprint arXiv:2305.16822* (2023).

[7] M. Anisetti, C. A. Ardagna, N. Bena, and A. Foppiani. "An Assurance-Based Risk Management Framework for Distributed Systems". In: *Proc. of IEEE ICWS 2021*. Chicago, IL, USA, Sept. 2021.

References II

[8] M. Anisetti, C. A. Ardagna, N. Bena, V. Giandomenico, and G. Gianini. "Lightweight Behavior-Based Malware Detection". In: *Proc. of MEDES 2023*. To appear. Heraklion, Greece, May 2023.

[9] M. Anisetti, C. A. Ardagna, F. Berto, and E. Damiani. "A Security Certification Scheme for Information-Centric Networks". In: *IEEE TNSM* 19.3 (2022).

[10] M. Anisetti, C. A. Ardagna, E. Damiani, and F. Gaudenzi. "A Semi-Automatic and Trustworthy Scheme for Continuous Cloud Service Certification". In: *IEEE TSC* 13.1 (2020).

[11] M. Anisetti, C. A. Ardagna, E. Damiani, and P. G. Panero. "A Methodology for Non-Functional Property Evaluation of Machine Learning Models". In: *Proc. of MEDES 2020*. Abu Dhabi, UAE, Nov. 2020.

[12] M. Anisetti, C. A. Ardagna, E. Damiani, and G. Polegri. "Test-Based Security Certification of Composite Services". In: *ACM TWEB* 13.1 (2019).

[13] M. Anisetti, N. Bena, F. Berto, and G. Jeon. "A DevSecOps-based Assurance Process for Big Data Analytics". In: *Proc. of IEEE ICWS 2022*. Barcelona, Spain, July 2022.

References III

[14] M. Anisetti, F. Berto, and M. Banzi. "Orchestration of data-intensive pipeline in 5G-enabled Edge Continuum". In: *Proc. of IEEE SERVICES 2022*. Barcelona, Spain, July 2022.

[15] C. Ardagna, R. Asal, E. Damiani, and Q. Vu. "From Security to Assurance in the Cloud: A Survey". In: *ACM CSUR* 48.1 (2015).

[16] C. A. Ardagna, R. Asal, E. Damiani, N. El Ioini, M. Elahi, and C. Pahl. "From Trustworthy Data to Trustworthy IoT: A Data Collection Methodology Based on Blockchain". In: *ACM TCPS* 5.1 (2021).

[17] C. A. Ardagna, N. Bena, C. Hebert, M. Krotsiani, C. Kloukinas, and G. Spanoudakis. "Big Data Assurance: An Approach Based on Service-Level Agreements". In: *Big Data* (2023).

[18] C. A. Ardagna, N. Bena, and R. M. de Pozuelo. "Bridging the Gap Between Certification and Software Development". In: *Proc. of ARES 2022*. Vienna, Austria, Aug. 2022.

[19] N. Bena, R. Bondaruc, and A. Polimeno. "Security Assurance in Modern IoT Systems". In: *Proc. of IEEE VTC 2022-Spring*. Helsinki, Finland, June 2022.