The 1st International Workshop on the Orchestration of the Serverless Edge-Cloud Continuum

July 31, 2024, Big Island, Hawaii, USA

- The **Serverless Function as a Service (FaaS)** model allows developers to write *stateless functions* that can be executed on-demand within containerised runtime environments.
- When a function is invoked, the FaaS platform handles resource provisioning, scaling, and operational aspects, billing users only for the actual time their code runs.
- This approach addresses the limitations of traditional Infrastructure as a Service (IaaS) Clouds, where resources must be provisioned in advance to meet workloads, offering cost savings during idle periods through a pay-per-use model and seamless scalability.

- *Edge Computing* shifts computation from centralised data centres to the edge, reducing latency and enabling real-time IoT processing.
- Edge nodes are typically resource-constrained, requiring efficient architectural designs for orchestrating distributed applications.
- The *Edge-Cloud continuum* provides near-user computation and access to cloud resources when needed.
- Centralised resource management and orchestration solutions used in the Cloud may not be suitable for the *dispersed* nature of the *Edge-Cloud continuum*.

The *Serverless FaaS* model can address these challenges by saving energy and compute resources, allowing users to share lightweight abstractions that can automatically be deallocated during idleness. However, current solutions feature central architectures not aligned with geo-distributed environments, making it challenging to harness the potential of distributed computing.

CALL FOR PAPERS

COHERENT (https://sites.google.com/view/coherent2024/) is focussed on the innovative application of the *Serverless Function as a Service* model to the *Edge-Cloud continuum*. It addresses various critical challenges, encouraging research into novel architectural designs for resource-constrained Edge environments, efficient Edge-Cloud distributed scheduling and resource allocation strategies. We invite researchers, practitioners, and experts in *Serverless Computing*, *Edge Computing*, and *Cloud Computing* to submit papers addressing the *integration of Serverless FaaS into the Edge-Cloud continuum*.

Topics of Interest (but not limited to)

- Innovative Serverless FaaS *architectural designs* for geo-distributed Edge-Cloud environments.
- Efficient Serverless FaaS *distributed scheduling* and *resource allocation* in the Edge-Cloud continuum.
- Approaches to guarantee *seamless scalability* of Serverless FaaS applications, minimising the *energy consumption* of resource-constrained edge nodes and considering applications SLAs.
- Solutions to enable the *orchestration of functions* in the Edge-Cloud continuum considering *hybrid models* and offloading to non-serverless Clouds.
- Efficient *execution* of Serverless FaaS *workflows* by trading off edge resources *pre-warming* and function *cold start* considering *data locality* and *geographic resource distribution*.
- Design of *strategies* and *mechanisms* to *manage application state* alongside functions for implementing *stateful applications* running in the Edge-Cloud continuum.

- Techniques enabling *real-time processing* and *automation* in IoT applications on the Edge-Cloud continuum using the Serverless FaaS model.
- *Federated management* and *deployment* of Serverless FaaS platforms in the heterogeneous *multi-provider* Edge-Cloud resource continuum.
- Solutions to address the *hardware heterogeneity* of the Edge-Cloud continuum, ranging from resource-constrained edge nodes to data centres, possibly including *hardware accelerators*.
- *Strategies* and *best practices* for ensuring *security* and *data protection* in Serverless FaaS applications at the edge.
- Autonomic *network management* of the Serverless FaaS Edge-Cloud Continuum.
- Real-world *case studies*, *use cases*, and *practical experiences* using the Serverless FaaS model in the Edge-Cloud continuum.

Instructions for Authors

Submitted manuscripts must be formatted in standard IEEE camera-ready format (double-column, 10-pt font) and submitted via Easychair at https://easychair.org/conferences/?conf=icccn2024 as PDF files (formatted for 8.5×11-inch paper). The manuscripts should be *no longer than 6 pages.* Submitted papers cannot have been previously published in or be under consideration for publication in another journal or conference. The workshop program committee reserves the right not to review papers that exceed the length specification or have been submitted or published elsewhere. Submissions must include a title, abstract, keywords, author(s) and affiliation(s) with postal and e-mail address(es).

Submitted papers will be reviewed by the workshop program committee and judged on originality, technical correctness, relevance, and presentation quality. An accepted paper must be presented at the *ICCCN 2024* venue (http://www.icccn.org/icccn24/) by one of the authors registered at the full registration rate. Each workshop registration covers up to two workshop papers by an author. Accepted and presented papers will be published in the *ICCCN proceedings* and submitted to IEEE Xplore as well as other Abstracting and Indexing (A&I) databases. IEEE reserves the right to exclude a paper from distribution after the conference, including IEEE Xplore® Digital Library, if the author does not present the paper at the conference.

We are working on organising a *Special Issue* for the *Journal of Grid Computing*, focusing on the topics covered by the workshop. Selected authors will be invited to submit *extended versions of their work* for consideration.

Important Dates

Paper submission deadline: March 4, 2024
Notification of acceptance: April 26, 2024
Camera-ready paper due: May 10, 2024 (Hard Deadline)
Workshop date: July 31, 2024

Workshop Co-Chairs

Francesco Tusa
University of Westminster, UK
Stuart Clayman
University College London, UK