Mariano Scazzariello

Isafjordsgatan, 22 – 164 40 Kista – Sweden
✉ mariano.scazzariello@ri.se • Updated: October 31, 2026

I am a Senior Researcher at RISE Research Institutes of Sweden. My work lies at the intersection of programmable networks, datacenter systems, and machine learning for networking.

Research Experience

RISE Research Institutes of Sweden — **Stockholm, Sweden**
Senior Researcher & RISE ICE Datacenter Manager — *Nov 2024–now*
Connected Intelligence Unit

The University of Texas at Austin — **Austin, Texas, US**
Visiting Researcher, UT Networked Systems Research Group — *Sep 2026*
Research on semantic-aware networking for LLM inference, hosted by Prof. Daehyeok Kim.

KTH Royal Institute of Technology — **Stockholm, Sweden**
Postdoctoral Researcher, "ULTRA" ERC-funded project — *Dec 2022–Nov 2024*
Networked Systems Laboratory (NSLab)

Roma Tre University — **Rome, Italy**
Ph.D. in Computer Science — *2019–2022*
Advisor: Prof. Giuseppe Di Battista
Degree Thesis: "Hyper-scale Datacenters: Physical Speed-up and Virtual Testing"

KTH Royal Institute of Technology — **Stockholm, Sweden**
Visiting Ph.D. Student — *Jan 2022–Feb 2022*
Design and Implementation of Ribosome [c6].

Roma Tre University — **Rome, Italy**
B.sc. and M.sc. in Computer Science — *2013–2019*
M.sc. Advisor: Prof. Giuseppe Di Battista
M.sc. Degree Thesis: "Virtualization and Orchestration of Large Network Infrastructures"
Final Grade: 110/110 with honors

Publications

International Conference Publications

[c1] **Scazzariello**, Caiazzi, Ghasemirahni, Kostić, and Chiesa. "Queue-Mem: Energy-Efficient Hardware Storage for Advanced Network Function Acceleration". In: *USENIX Symposium on Networked Systems Design and Implementation (**NSDI**)*. 2026.

[c2] Ghasemirahni, Farshin, **Scazzariello**, Chiesa, Maguire Jr, and Kostić. "FAJITA: Stateful Packet Processing at 100 Million pps". In: *ACM Conference on emerging Networking EXperiments and Technologies (**CoNEXT**)*. 2024.

[c3] Wang, **Scazzariello**, Farshin, Ferlin, Kostić, and Chiesa. "NetConfEval: Can LLMs Facilitate Network Configuration?" In: *ACM Conference on emerging Networking EXperiments and Technologies (**CoNEXT**)*. 2024.

[c4] Polverini, Cianfrani, Caiazzi, **Scazzariello**, Abdelsalam, Filsfils, and Camarillo. "Achieving Best-path Selection at Line Rate through the SRv6 Live-Live Behavior". In: *IEEE/IFIP Network Operations and Management Symposium (**NOMS**)*. 2024.

[c5] Caiazzi, **Scazzariello**, and Chiesa. "Millions of Low-latency State Insertions on ASIC Switches". In: *ACM Conference on emerging Networking EXperiments and Technologies (**CoNEXT**)*. 2023.

[c6] **Scazzariello**, Caiazzi, Ghasemirahni, Barbette, Kostić, and Chiesa. "A High-Speed Stateful Packet Processing Approach for Tbps Programmable Switches". In: *USENIX Symposium on Networked Systems Design and Implementation (**NSDI**)*. 2023.

[c7] Caiazzi, **Scazzariello**, Alberro, Ariemma, Castro, Grampin, and Di Battista. "Sibyl: a Framework for Evaluating the Implementation of Routing Protocols in Fat-Trees". In: *IEEE/IFIP Network Operations and Management Symposium (**NOMS**)*. 2022.

[c8] **Scazzariello**, Ariemma, Di Battista, and Patrignani. "Megalos: A Scalable Architecture for the Virtualization of Network Scenarios". In: *IEEE/IFIP Network Operations and Management Symposium (**NOMS**)*. 2020.

International Journal Publications

[j1] Polverini, Cianfrani, Listanti, Caiazzi, and **Scazzariello**. "In-Network Q-Learning-Based Packet Forwarding for Delay Sensitive Applications". In: *IEEE Network* (2025).

[j2] Polverini, Cianfrani, Caiazzi, and **Scazzariello**. "SRv6 Meets DetNet: A New Behavior for Low Latency and High Reliability". In: *IEEE Journal on Selected Areas in Communications* (2025).

[j3] **Scazzariello**, Ariemma, Di Battista, and Patrignani. "Megalos: A Scalable Architecture for the Virtualization of Large Network Scenarios". In: *Future Internet* (2021).

Workshop Papers, Demos, and Preprints

[o1] De Nadai, **Scazzariello**, Caiazzi, and Di Battista. "Proxy-based Architecture for Wide-Area RDMA". In: *ACM Conference on emerging Networking EXperiments and Technologies (**CoNEXT**)*. 2026.

[o2] **Scazzariello**, Rotman, Gavrilenko, Khashab, Shpiner, Kadosh, Chiesa, Kostić, and Silberstein. *Avoiding Cross-Datacenter Collective Congestion via Disaggregated Buffering*. `https://arxiv.org/abs/2605.11852`. 2026.

[o3] Siavashi, **Scazzariello**, Maguire Jr, Kostić, and Chiesa. *Blink: CPU-Free LLM Inference by Delegating the Serving Stack to GPU and SmartNIC*. `https://arxiv.org/abs/2604.07609`. 2026.

[o4] De Luca, De Nadai, **Scazzariello**, Caiazzi, Farshin, Chiesa, and Di Battista. "Aggregate Local, Sync Global: A Hierarchical Approach to Efficient Geo-Distributed LLM Training". In: *Workshop on Computation over Heterogeneous Networks (**NetCompute**)*. 2026.

[o5] Nordqvist, Wang, Ferlin, **Scazzariello**, and Chiesa. "RAG Against the Machine: Zero-Shot Software Vulnerabilities Classification using LLMs". In: *International Workshop on Large Language Models for Code (**LLM4Code**)*. 2026.

[o6] Wang, **Scazzariello**, Kostić, and Chiesa. "Toward Automated, Contamination-free Dafny Benchmark Generation". In: ***Dafny Workshop***. 2026.

[o7] Wang, **Scazzariello**, Alshnakat, Guanciale, Kostić, and Chiesa. *Dissect-and-Restore: AI-based Code Verification with Transient Refactoring*. `https://arxiv.org/abs/2510.25406`. 2025.

[o8] Langlet, **Scazzariello**, Luciani, Burocchi, Kostić, and Chiesa. *Can LLMs Forecast Internet Traffic from Social Media?* `https://arxiv.org/abs/2509.20123`. 2025.

[o9] Wang, **Scazzariello**, and Chiesa. "From Scientific Texts to Verifiable Code: Automating the Process with Transformers". In: *International Workshop on Large Language Models for Code (**LLM4Code**)*. 2025.

[o10] Puccioni, Farshin, **Scazzariello**, Wang, Chiesa, and Kostić. "Deriving Coding-Specific Sub-Models from LLMs using Resource-Efficient Pruning". In: *International Workshop on Large Language Models for Code (**LLM4Code**)*. 2025.

[o11] Cipollone, Wang, **Scazzariello**, Ferlin, Izadi, Kostić, and Chiesa. "Automating the Detection of Code Vulnerabilities by Analyzing GitHub Issues". In: *International Workshop on Large Language Models for Code (**LLM4Code**)*. 2025.

[o12] **Scazzariello**, Caiazzi, and Chiesa. "Deliberately Congesting a Switch for Better Network Functions Performance". In: *P4 Workshop in Europe (**EuroP4**)*. 2024.

[o13] Girondi, **Scazzariello**, Maguire Jr, and Kostić. "Toward GPU-centric Networking on Commodity Hardware". In: *International Workshop on Edge Systems, Analytics and Networking (**EdgeSys**)*. 2024.

[o14] Ghasemirahni, Farshin, **Scazzariello**, Chiesa, and Kostić. "Deploying Stateful Network Functions Efficiently using Large Language Models". In: *Workshop on Machine Learning and Systems (**EuroMLSys**)*. 2024.

[o15] Wang, **Scazzariello**, Farshin, Kostić, and Chiesa. *Making Network Configuration Human Friendly*. `https://arxiv.org/abs/2309.06342`. 2023.

[o16] Caiazzi, **Scazzariello**, Quinzi, Ariemma, Patrignani, and Di Battista. "Nesting Containers for Faithful Datacenters Emulations". In: *IEEE/IFIP International Workshop On Technologies For Network Twins (**TNT**)*. 2023.

[o17] Caiazzi, **Scazzariello**, and Ariemma. "VFTGen: a Tool to Perform Experiments in Virtual Fat Tree Topologies". In: *IFIP/IEEE International Symposium on Integrated Network Management (**IM**)*. 2021.

[o18] Ariemma, **Scazzariello**, and Caiazzi. "MRT#: a Fast Multi-Threaded MRT Parser". In: *IFIP/IEEE International Symposium on Integrated Network Management (**IM**)*. 2021.

[o19] **Scazzariello**, Ariemma, and Caiazzi. "Kathará: A Lightweight Network Emulation System". In: *IEEE/IFIP Network Operations and Management Symposium (**NOMS**)*. 2020.

Awards

- **IRTF/IETF Applied Research Networking Prize 2025**
  For our CoNEXT paper *"NetConfEval: Can LLMs Facilitate Network Configuration?"*
  Awarded to Changjie Wang to "recognise the best recent results in applied networking, interesting new research ideas of potential relevance to the Internet standards community".
- **Runner-up Best Paper at ACM CoNEXT 2024**
  For our paper *"NetConfEval: Can LLMs Facilitate Network Configuration?"*
  Selected among the top 3 best papers at the conference.
- **Best Paper at ACM CoNEXT 2023**
  For our paper *"Millions of Low-latency State Insertions on ASIC Switches"*
- **Best Demo at IFIP/IEEE IM 2021**
  For our demo *"MRT#: a Fast Multi-Threaded MRT Parser"*
- **Best Demo at IFIP/IEEE NOMS 2020**
  For our demo *"Kathará: A Lightweight Network Emulation System"*

Research Funding

Research Projects

Democratizing Trustworthy Software Development **2026-2029**
Funded by **Digital Futures** (Total Funding: **8.7M SEK**)
Role: Co-PI
Partners: KTH, RISE

SAFIR: Secure AI for Intelligent Resilience and Confidentiality in the Cloud **2026-2028**
Funded by **Vinnova** (Total Funding: **27M SEK**)
Role: Acted as one of the co-PIs of RISE
Partners: CanaryBit, Ericsson, KTH, Nvidia, RedHat, RISE, SAAB

EdgeWise: Fast & Explainable Next-Generation Robots for MCA **2026-2027**
Funded by **Digital Futures** as a Demonstration Project (Total Funding: **2M SEK**)
Role: Acted as one of the co-PIs of RISE
Partners: KTH, RISE

Individual Grants & Mobility

Individual Mobility within 6G **Sep 2026-Oct 2026**
Funded by **Vinnova** (Total Funding: ∼**210K SEK**)
Role: Principal Applicant
Research visit at The University of Texas at Austin

Professional Service

- **Program Commitee**
  2026: CoNEXT 2026, ICNP 2026, IWQoS 2026, EuroP4 2026, IPSN 2026
  2025: ICNP 2025, IWQoS 2025, ANRW 2025, NPC 2025, IPSN 2025
  2024: NSDI 2025, ICNP 2024 Demos & Posters, IWQoS 2024, NPC 2024, IPSN 2024
- **Organizing Commitee**
  2026: Netdev 0x1A
- **Award Commitee**
  2026: ACM CoNEXT Best Paper
- **Panel Evaluation**
  2026: Digital Futures postdoc grants
- **External Reviewer**
  IEEE/ACM Transactions on Networking (2024-2025)
  ACM SIGCOMM Computer Communication Review (2024)
  Elsevier Computer Networks (2024)
  IEEE Network Magazine (2023)

Teaching & Supervision Experience

Teaching Assistant **KTH Royal Institute of Technology**
IK2227: Network Systems with Edge or Cloud Datacenters *Period III, 2025*
Practical lab sessions using Kathará, covering the following topics: data center routing, EVPN BGP & VXLAN, Kubernetes, RDMA, and Anycast & DNS.

Ph.D. Students Supervision **RISE Research Institutes of Sweden**
2025 – current: Francesco De Nadai (Roma Tre University, industrial co-supervisor)

Master Students Supervision **RISE Research Institutes of Sweden**
2026: Joris Köster, Lorenzo Sacco, Jacopo Veronese
2025: Francesco De Nadai, Francesco De Luca, Luis Jira, Zhicheng Dong, Daniél Laki

Master Students Supervision **KTH Royal Institute of Technology**
2024: Daniele Cipollone, Laura Puccioni
2023: Changjie Wang, Sepehr Javid, Jingyi Pan, Chenming Du

Bachelor & Master Students Supervision **Roma Tre University**
2022: Gabriele Alecci, Alessandro Volpe, Andrea Marrocco
2021: Marta Burocchi, Tommaso Caiazzi
2020: Alex Cocciatelli

Thesis Reviewer **Roma Tre University**
Read, reviewed, and graded 11 external B.sc./M.sc. theses *2019–2022*

Selected Talks

IFIP TMA (PhD School) — **Delft, Netherlands**
From Network Topologies to Experiments: A Practical Introduction to Kathará — *June 2026*

USENIX NSDI — **Renton, WA, US**
Queue-Mem: Energy-Efficient Hardware Storage for Advanced Network Function Acceleration — *May 2026*
Video: `https://www.usenix.org/conference/nsdi26/presentation/scazzariello`

Nvidia Research — **Virtual**
Queue-Mem: Energy-Efficient Hardware Storage for Advanced Network Function Acceleration — *Oct 2025*

Roma Tre University — **Rome, Italy**
Invited Talk: "NetConfEval: Can LLMs Facilitate Network Configuration?" — *Jan 2025*

ACM CoNEXT — **Los Angeles, CA, US**
FAJITA: Stateful Packet Processing at 100 Million pps — *Dec 2024*
Video: `https://scazzariello.net/storage/media/Fajita-small2.mp4`

EuroP4 Workshop — **Charleroi, Belgium**
Deliberately Congesting a Switch for Better Network Functions Performance — *Oct 2024*

"Between 0x2 Nerds" Podcast — **Virtual**
BGP Security: ROSE-T (ROuting SEcurity Tool) — *Jan 2024*
Video: `https://www.youtube.com/watch?v=DLpzOmpRWCM`

ACM CoNEXT — **Paris, France**
Millions of Low-latency State Insertions on ASIC Switches — *Dec 2023*
Video: `https://www.youtube.com/watch?v=SXem7FXMUZo&t=2538`

RIPE 87 Meeting — **Rome, Italy**
ROSE-T: Making MANRS Compliance Simple (And Automatic!) — *Nov 2023*
Video: `https://ripe87.ripe.net/archives/video/1262/`

ACM SIGMETRICS — **Orlando, FL, US**
Network Monitoring on Multi-Pipe Switches — *Jun 2023*

IEEE/IFIP NOMS — **Miami, FL, US**
Nesting Containers for Faithful Datacenters Emulations — *May 2023*

Roma Tre University — **Rome, Italy**
Hyper-scale Datacenters: Physical Speed-up and Virtual Testing — *Apr 2023*

RIPE 85 Meeting — **Belgrade, Serbia**
Kathará: A Lightweight and Scalable Network Emulation System — *Oct 2022*
Video: `https://ripe85.ripe.net/archives/video/941/`

Polytechnic University of Marche — **Virtual**
Guest Lectures on Data Center Networking — *May 2020*

IEEE/IFIP NOMS — **Budapest, Hungary**
Sibyl: a Framework for Evaluating the Implementation of Routing Protocols in Fat-Trees — *Apr 2022*

IEEE/IFIP NOMS — **Virtual**
Megalos: A Scalable Architecture for the Virtualization of Network Scenarios — *Apr 2020*

IETF 107 Meeting - RIFT Working Group **Virtual**
Tools for Experimenting Routing in the Data Centers *Apr 2020*

Selected News & Articles

KTH News **Feb 2025**
How Can Large Language Models Make Network Configuration Human Friendly?

ACM Digital Library **Aug 2024**
Featured Article: "NetConfEval: Can LLMs Facilitate Network Configuration?"

APNIC Blog **May 2024**
Verify MANRS Compliance Automatically with ROSE-T

ACM Digital Library **Apr 2024**
Featured Article: "Millions of Low-latency State Insertions on ASIC Switches"

MANRS Blog **Mar 2024**
Verify MANRS Compliance Automatically with ROSE-T

KTH News **Feb 2024**
New cybersecurity analysis solution leads to significant reduction in energy consumption

Languages

Italian: Mothertongue
English: Proficient