Towards the Medicine of the Future in Bavaria and Germany, One Heartbeat at the Time With Confidential Computing

Florent Dufour

Deadline: 2022-12-15

Since 2018, the Bavarian ministry of health has invested 24.5 million euros in the DigiMed Bayern project[1] with the ambition to create the lighthouse that will guide Germany towards the medicine of the future. By developing a legal framework and a secure environment powered by confidential computing technologies, over one hundred researchers, clinicians, lawyers, and tinkerers from academia and industry across 14 institutions have found a sovereign computing environment to collaborate on sensitive multi-omic medical data. With the common goal to advance research on heart disease, they already published more than 50 scientific publications and developed large-scale studies like Vroni[2] and smart wearable technologies like HerzFit[3].

In the first part of this talk, we will focus on the Bavarian Cloud for Health Research (BCHR) which is the cornerstone of the project. Architected around confidential computing technologies and hosted at the top-tier Leibniz Supercomputing Centre (LRZ) in Munich[4], we will present how the Big Data and Artificial Intelligence team has engineered the BCHR with security and performance for AI/ML workloads in mind. We will showcase heterogeneous workloads running on the OpenStack-based cloud with hundreds of cores at the petabyte scale, as well as its prospective integration in the European cloud GAIA-X[5].

In the second part of this talk, we will focus on a new axis of research opened by confidential computing in the area of Privacy-Preserving AI. While approaches like Differential Privacy, Secure Multiparty Computation, or Homomorphic Encryption allow parties to collaborate on confidential data, they come at the expense of the model's utility. We will discuss how TEEs can be repurposed for AI workloads and allow to train models privately, at high velocity, and without reducing the model's accuracy[6]. The emphasis will be put on computer vision applications with convolutional neural networks, secure inference in TEEs, hardware acceleration with GPUs, and remote attestation of the privacy guarantee.

Keywords: Data-driven medical research, Artificial Intelligence, OpenStack, GPU, Quobyte, AMD-SEV/SNP, GAIA-X, Leibniz Supercomputing Centre, Technical University Munich.

References

[1] The DigiMed Bayern Consortium. *DigiMed Bayern.* Oct. 2018. URL: https://www.digimed-bayern.de/.

[2] Veronika Sanin et al. "Population-based screening in children for early diagnosis and treatment of familial hypercholesterolemia: design of the VRONI study". In: *European Journal of Public Health* 32.3 (June 1, 2022), pp. 422–428. ISSN: 1101-1262. DOI: 10.1093/eurpub/ckac007. URL: https://doi.org/10.1093/eurpub/ckac007.

[3] Deutsche Herzstiftung e.V. *Die neue HerzFit-App der Deutschen Herzstiftung.* Apr. 22, 2022. URL: https://youtu.be/fcs2SCPZfTs.

[4] Heinz-Gerd Hegering, Dietmar Täube, and Victor Apostolescu. *50 Jahre LRZ: das Leibniz-Rechenzentrum der Bayerischen Akademie der Wissenschaften ; Chronik einer Erfolgsgeschichte ; 1962 - 2012.* In collab. with Leibniz-Rechenzentrum. Garching: Leibniz-Rechenzentrum, 2012. 304 pp. ISBN: 978-3-00-038333-5.

[5] Federal Ministry for Economic Affairs and Energy (BMWi). "Project GAIA-X – A Federated Data Infrastructure as the Cradle of a Vibrant European Ecosystem". In: (Oct. 2019), p. 56.

[6] *SoK: Machine Learning with Confidential Computing.* Aug. 22, 2022. arXiv: 2208.10134[cs]. URL: http://arxiv.org/abs/2208.10134.