

From Data Centers to Fog Computing: The Evaporating Cloud

Guillaume Pierre
Univ Rennes, Inria, CNRS, IRISA

Cloud computing data centers are composed of very powerful computing nodes connected by reliable backbone networks. However, these resources are concentrated in a small number of data centers. The latency between an end user and the closest available cloud data center comes in the range of 20-150 ms. A number of latency-sensitive applications (e.g., augmented reality) require extremely low end-to-end latencies and therefore cannot make use of traditional cloud platforms. Fog computing therefore aims to complement traditional cloud infrastructures with additional resources located extremely close to the user, within a couple of network hops. This requires one to distribute machines in a very large number of geographical locations so computation capacity is always available in immediate proximity of any end user. In this presentation I will discuss the application scenarios where fog computing is or isn't useful, and the architectural challenges one needs to face when designing the next-generation fog computing architectures. Finally, I will introduce the FogGuru European project which aims to train the next generation of fog computing specialists.