



Offer #2021-04074

Energy-efficient data stream processing

Contract type : Fixed-term contract

Level of qualifications required : Graduate degree or equivalent

Other valued qualifications : Master degree in distributed systems and/or cloud computing

Fonction : Temporary scientific engineer

Level of experience : Up to 3 years

About the research centre or Inria department

The Inria Rennes – Bretagne Atlantique Centre is one of Inria's eight centres and has more than thirty research teams. The Inria Center is a major and recognized player in the field of digital sciences. It is at the heart of a rich R&D and innovation ecosystem: highly innovative PMEs, large industrial groups, competitiveness clusters, research and higher education players, laboratories of excellence, technological research institute, etc.

Context

With the advent of Cloud computing, service-oriented software and distributed systems have become commonplace from end-user applications to the internal control plane of infrastructures (e.g., 5G, Cloud, Fog or Edge computing). However, the energetic cost of the Cloud is an important issue for the sustainability of the IT sector, and more broadly of our societies.

As an important Cloud provider in the IT market, the goal of OVHCloud in the years to come is to reduce and optimize its energy consumption. To achieve this goal many scientific challenges have to be studied. This is why a large project collaboration between Inria and OVHCloud has been built.

The selected candidate will have the opportunity to be co-supervised by Prof. Guillaume Pierre at Inria Rennes, Prof. Romain Rouvoy at Inria Lille, and by Paul Kerguillec and Hamouda Makhoulfi at OVHCloud.

Assignment

Data managed by companies are increasingly produced at the edge of the Internet, far from the data centers. One of the main reasons for this profound change is the increasing weight data from the Internet of Things (IoT). These data are generated as a data flow that many users want to be able to analyze in real time, in order to allow a reaction almost instantaneous to events detected in the data streams. Currently there are many frameworks for streaming data analysis, the most popular being Apache Storm and Apache Flink. The goal of this project is to study how such data stream processing systems can be made energy-efficient.

The recruited engineer will be in charge of adapting the existing data stream processing experimental platform to integrate energy measurement functionalities and support a broader range of workloads and application scenarios. The resulting system will be used to quantify the respective energy savings that may be expected from vertical and horizontal scaling techniques when the incoming data workload varies over time and/or when different costs for energy exist in different parts of the infrastructure.

This offer is advertised for an engineer position, with a duration of 12 months. If the work makes good progress there will be a possibility to extend it with a 3-years PhD student position.

[1] Hamidreza Arkian, Guillaume Pierre, Johan Tordsson and Erik Elmroth : An Experiment-Driven Performance Model of Stream Processing Operators in Fog Computing Environments. In Proc. ACM SAC, March 2020. <https://hal.inria.fr/hal-02394396>

[2] Hamidreza Arkian, Guillaume Pierre, Johan Tordsson and Erik Elmroth : Model-based Stream Processing Auto-scaling in Geo-Distributed Environments. In Proc. ICCCN, July 2021. <https://hal.inria.fr/hal-03206689>

[3] Guillaume Fieni, Romain Rouvoy et Lionel Seinturier : SmartWatts : Self-Calibrating Software-Defined

Power Meter for Containers. In Proc. CCGRID, 2020. <https://hal.inria.fr/hal-02470128>.

[4] Aurélien Havet, Rafael Pires, Pascal Felber, Marcelo Pasin, Romain Rouvoy et Valerio Schiavoni. SecureStreams : A Reactive Middleware Framework for Secure Data Stream Processing. In Proc. ACM DEBS, 2017. <https://hal.inria.fr/hal-01510699v2>.

Main activities

Main activities (5 maximum) :

- Integrate the experimental stream processing testbed in a cloud environment
- Integrate energy measurements in the platform
- Extend the platform to support a broader range of workloads and application scenarios
- Perform experiments to quantify the respective energy savings from vertical and horizontal scaling techniques

Skills

- A master degree in distributed systems and/or Cloud computing.
- Excellent programming skills in Linux environments.
- Excellent communication and writing skills.
- Good command of English.
- Knowledge of the following technologies is not mandatory but will be considered as a plus:
 - Cloud resource scheduling
 - Distributed container systems: Kubernetes, Docker Swarm.
 - Single-board computers such as Raspberry PI
 - Python and shell scripting
 - Revision control systems: git, svn.
 - Linux distributions: Debian, Ubuntu.

Benefits package

- Subsidized meals
- Partial reimbursement of public transport costs
- Leave: 7 weeks of annual leave + 10 extra days off due to RTT (statutory reduction in working hours)
- Possibility of teleworking and flexible organization of working hours
- Professional equipment available (videoconferencing, loan of computer equipment, etc.)
- Social, cultural and sports events and activities
- Access to vocational training
- Social security coverage

Remuneration

Monthly gross salary from 2562 euros according to diploma and experience.

General Information

- **Theme/Domain** : Distributed Systems and middleware System & Networks (BAP E)
- **Town/city** : Rennes
- **Inria Center** : [Centre Inria de l'Université de Rennes](#)
- **Starting date** : 2021-11-01
- **Duration of contract** : 1 year
- **Deadline to apply** : 2021-11-15

Contacts

- **Inria Team** : [MYRIADS](#)
- **Recruiter** :
Pierre Guillaume / Guillaume.Pierre@irisa.fr

About Inria

Inria is the French national research institute dedicated to digital science and technology. It employs 2,600 people. Its 200 agile project teams, generally run jointly with academic partners, include more than 3,500 scientists and engineers working to meet the challenges of digital technology, often at the interface with other disciplines. The Institute also employs numerous talents in over forty different professions. 900 research support staff contribute to the preparation and development of scientific and entrepreneurial projects that have a worldwide impact.

The keys to success

This offer is advertised for an engineer position, with a duration of 12 months. If the work makes good progress there will be a possibility to extend it with a 3-years PhD student position.

Warning : you must enter your e-mail address in order to save your application to Inria. Applications must be submitted online on the Inria website. Processing of applications sent from other channels is not guaranteed.

Instruction to apply

Please submit online : your resume, cover letter and letters of recommendation eventually

For more information, please contact guillaume.pierre@irisa.fr

Defence Security :

This position is likely to be situated in a restricted area (ZRR), as defined in Decree No. 2011-1425 relating to the protection of national scientific and technical potential (PPST). Authorisation to enter an area is granted by the director of the unit, following a favourable Ministerial decision, as defined in the decree of 3 July 2012 relating to the PPST. An unfavourable Ministerial decision in respect of a position situated in a ZRR would result in the cancellation of the appointment.

Recruitment Policy :

As part of its diversity policy, all Inria positions are accessible to people with disabilities.