



Offer #2024-08214

Engineer: Data placement and retrieval in erasure coded IPFS

Renewable contract : Yes

Level of qualifications required : Graduate degree or equivalent

Fonction : Temporary scientific engineer

About the research centre or Inria department

The Inria Centre at Rennes University is one of Inria's eight centres and has more than thirty research teams. The Inria Centre 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

Financial and working environment.

This engineer position will be in the context of Alvearium challenge (<https://project.inria.fr/alvearium/>), a project between HIVE, Coast, Magellan and Wide Inria teams. The engineer will be recruited and hosted at the Inria Center at Rennes University and will be visiting Inria Nancy-Grand Est and the Hive offices in Cannes. The work will be carried out within the MAGELLAN team in collaboration with other partners.

The position is for one year, with the possibility of an extension to 18 months.

About Hive:

Hive intends to play the role of a next generation cloud provider in the context of Web 3.0. Hive aims to exploit the unused capacity of computers to offer the general public a greener and more sovereign alternative to the existing clouds where the true power lies in the hands of the users. It relies both on distributed peer-to-peer networks, on the encryption of end-to-end data and on blockchain technology.

Assignment

Context:

Erasure coding (EC) has been increasingly used in storage systems as an alternative for replication to provide high data availability at relatively lower storage and power costs. For example, EC has been deployed in data analytic systems [1, 2, 3, 4] and in-memory storage systems on cached (hot) data [5]. EC can be an ideal candidate for large scale peer-to-peer storage systems (exploits parallel read and write of data, involves large number of nodes in storing and repairing data). However, unlike previous efforts where EC is mainly used for achieved data in P2P system [6, 7], performing EC on the critical path of data access (which is the case in this project) in large scale P2P storage system (to store hot and frequently accessed data) poses many research challenges on how to ensure high data availability and meet data and node dynamicity, and on how to provide cost-effective and heterogeneity-aware data repair.

Objectives:

The goal of this engineer position is to implement existing Erasure coding families in Go and to develop new data placement and retrieval stratifies for erasure coded IPFS [8]. The engineer will work closely with a PhD student on this topic.

[1] H. Chen, H. Zhang, M. Dong, Z. Wang, Y. Xia, H. Guan, and B. Zang. "Efficient and available in- memory KV-store with hybrid erasure coding and replication". In: ACM Transactions on Storage (TOS) 13.3 (2017), pp. 1–30. doi: 10.1145/3129900.

[2] J. Darrous and S. Ibrahim. "Understanding the performance of erasure codes in hadoop distributed file system". In: CHEOPS@EuroSys 2022: Proceedings of the Workshop on Challenges and Opportunities of Efficient and Performant Storage Systems, Rennes, France, 5 April 2022. Ed. by M. Kuhn, K. Duwe, J. Acquaviva, K. Chasapis, and J. Boukhobza. ACM, 2022, pp. 24–32. doi: 10.1145/3503646.3524296.

[3] J. Darrou, S. Ibrahim, and C. Peñrez. "Is it Time to Revisit Erasure Coding in Data-Intensive Clusters?" In: 27th IEEE International Symposium on Modeling, Analysis, and Simulation of Computer and Telecommunication Systems, MASCOTS 2019, Rennes, France, October 21-25, 2019. IEEE Computer Society, 2019, pp. 165–178. doi: 10.1109/MASCOTS.2019.00026.

[4] Z. Zhang, A. Deshpande, X. Ma, E. Thereska, and D. Narayanan. Does erasure coding have a role to play in my data center? Tech. rep. MSR-TR-2010-52. May 2010.

[5] K. V. Rashmi, M. Chowdhury, J. Kosaian, I. Stoica, and K. Ramchandran. "EC-Cache: Load-Balanced, Low-Latency Cluster Caching with Online Erasure Coding". In: 12th USENIX Symposium on Operating Systems Design and Implementation, OSDI 2016, Savannah, GA, USA, November 2-4, 2016. Ed. by K. Keeton and T. Roscoe. USENIX Association, 2016, pp. 401–417.

[6] J. Kubiatoiwicz, D. Bindel, Y. Chen, S. Czerwinski, P. Eaton, D. Geels, R. Gummadi, S. Rhea, H. Weatherspoon, W. Weimer, C. Wells, and B. Zhao. "OceanStore: An Architecture for Global-Scale Persistent Storage". In: SIGPLAN Not. 35.11 (Nov. 2000), pp. 190–201. issn: 0362-1340. doi: 10.1145/356989.357007.

[7] R. Rodrigues and B. Liskov. "High Availability in DHTs: Erasure Coding vs. Replication". In: Peer-to-Peer Systems IV, 4th International Workshop, IPTPS 2005, Ithaca, NY, USA, February 24-25, 2005, Revised Selected Papers. Ed. by M. Castro and R. van Renesse. Vol. 3640. Lecture Notes in Computer Science. Springer, 2005, pp. 226–239. doi: 10.1007/11558989_21.

[8] J. Benet. "IPFS - Content Addressed, Versioned, P2P File System". In: CoRR abs/1407.3561 (2014). doi: 10.48550/arXiv.1407.3561. arXiv: 1407.3561.

Main activities

- Implement several erasure coding families in Go and integrate them within IPFS
- Study new algorithms for data placement and retrieval
- Prototype key algorithms
- Run experiments and Evaluation of results
- Reporting, disseminating and presenting results

Skills

- Engineering and/or Master 2 degree in Computer science
- Good knowledge in distributed systems
- Software development skills: Python and Go
- Knowledge on storage and file systems
- Good collaborative and networking skills
- Excellent written and oral communication in English

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 exceptional leave (sick children, moving home, etc.)
- Possibility of teleworking (after 6 months of employment) 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

The engineer will receive a salary of €2,695 gross per month, based on their profile.

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** : 2025-01-01
- **Duration of contract** : 12 months
- **Deadline to apply** : 2024-12-28

Contacts

- **Inria Team** : [MAGELLAN](#)
- **Recruiter** : Ibrahim Shadi / Shadi.Ibrahim@inria.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.

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.

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.