Research internship (Master 2)

Management of data movement for resiliency across the Edge-Cloud Continuum

Guillaume Rosinosky, Daniel Balouek
Inria, IMT Atlantique, France

2025

Keywords *:* Edge-to-Cloud continuum ; urgent computing ; Cloud Computing ; Data movement

1 Introduction

With the advent of distributed infrastructures, the Cloud computing paradigm is progressively moving towards a full continuum from IoT devices and sensors to the centralized Cloud, with Edge (edge of the network) and Fog computing (core network) in between [1]. Simultaneously, distributed applications also evolve. Urgent computing tackles services that require time-critical decisions that improve quality of life, monitor civil infrastructures, respond to natural disasters and extreme events, and accelerate science (e.g., autonomous cars, disaster response, precision medicine, etc.). These services are typically sensitive to latency and response time [2].

In this internship, we consider to increase the resiliency of IoT-to-Cloud scenarios by managing the data associated with services. First, these services are often being under the constraints to move data from the edge of the network to the Cloud [3]. Second, these services and applications require system support to program reactions that occur at runtime, especially when the target infrastructure capacities and capabilities are unknown during the design [12].

This internship focuses on data movement support for the management on databases and event stores distributed across the Edge-Cloud Computing continuum. The location (which physical server) and the execution parameters (which software configuration) have significant effect on the performance of applications.

2 Expected work

Data movement or *live data migration* is a way for providers to optimize the usage of their computing infrastructure and to permit their users to enjoy low-latency in geo-distributed contexts without downtime. Numerous approaches have considered migrations at the virtual machine level [7], the container level [11], or the tenant level [14]. However these approaches' level of granularities are often too high as the whole data of a given module or tenant is moved amongst devices while the target user data can be a subset of the hosted data. To address these problems, main SaaS providers often use instead live migrations at the shard level [6, 4, 9, 5]. This shard migration approach is also supported by large scale proprietary database providers such as MongoDB [10], Google Spanner [5].

In this work, we want the candidate to study the effects of data placement and data migration on latency, as the migrations provoke overheads during the operations [13]. In this context, the candidate

will be in charge of evaluating the effects of placement and live data migrations on a software stack composed of the IoT platform SiteWhere and geo-distributed IoT devices. The SiteWhere platform's goal is to facilitate the ingestion, storage, processing and integration of IoT device data at a large scale [8]. Its architecture is microservice-based and makes usage of the Kubernetes container orchestrator and the Istio service mesh. Its data layer modules include the Postgresql relational database, the InfluxDB time series database, and the Kafka data event store, those being the target of the live data migrations.

The main challenge consists in evaluating the different outcomes of the system based on data placement, live data migrations, and deducing management policies based on their impact.

The objectives of this internship are :

— a state-of-the-art to assess the concepts associated to the computing continuum and data-driven analytics

— experiment live-migrations techniques for the target data layer modules. These techniques can imply the development of scripts or the usage of synchronization tools such as MirrorMaker for Kafka

— evaluating the model on a real-platform using a realistic micro-service application

In this modeling and evaluation tasks, consideration of metrics relative to latency, quality of service, and throughput during the lifecycle of an urgent application is of particular importance.

We expect the successful candidate to create repeatable processes and artifacts that will be used at scale to develop and evolve edge computing designs. Experiments and validation will occur on FIT IoT-Lab [1] and Grid'5000 [2], the biggest share network dedicated to research in Computer Science.

Note that if satisfactory, the successful candidate will probably have an opportunity to start a Ph.D. thesis after the internship.

3 Skills

The following skills are expected from the successful candidate :

— a student in the last year of a Master's degree in Computer Science (or in the last year of an engineering school with a computer science option) ;

— knowledge of the Python programming language and also preferably low-level languages such as Rust, C or C++

— knowledge of the Kubernetes platform

— knowledge or interest in learning the usage of event stores and databases such as Kafka, InfluxDB, Postgres

— Basics in probability and statistic

— a good level of English to contribute to the writing of a research paper ;

— an ability to collaborate and communicate ;

— curiosity and an appetite for learning new things.

4 Additional information

Advisors

— Guillaume Rosinosky, Inria, IMT Atlantique guillaume.rosinosky@inria.fr

— Daniel Balouek-Thomert, Inria, IMT Atlantique daniel.balouek@inria.fr

1. https://www.iot-lab.info/
2. https://www.grid5000.fr/w/Grid5000:Home

Duration 6 months

Salary legal amount of 3,90€ / hour, full time

Location IMT Atlantique, équipe Inria Stack, laboratoire LS2N à Nantes

Références

[1] Daniel Balouek-Thomert, Eduard Gibert Renart, Ali Reza Zamani, Anthony Simonet, and Manish Parashar. Towards a computing continuum : Enabling edge-to-cloud integration for data-driven workflows. *The International Journal of High Performance Computing Applications*, 33(6) :1159–1174, 2019.

[2] Daniel Balouek-Thomert, Ivan Rodero, and Manish Parashar. Harnessing the computing continuum for urgent science. *SIGMETRICS Perform. Eval. Rev.*, 48(2) :41–46, November 2020. ISSN 0163-5999. doi:10.1145/3439602.3439618. URL https://doi.org/10.1145/3439602.3439618.

[3] Kevin Fauvel, Daniel Balouek-Thomert, Diego Melgar, Pedro Silva, Anthony Simonet, Gabriel Antoniu, Alexandru Costan, Véronique Masson, Manish Parashar, Ivan Rodero, and Alexandre Termier. A distributed multi-sensor machine learning approach to earthquake early warning. *Proceedings of the AAAI Conference on Artificial Intelligence*, 34(01) :403–411, Apr. 2020. doi:10.1609/aaai.v34i01.5376. URL https://ojs.aaai.org/index.php/AAAI/article/view/5376.

[4] Garrett Fidalgo. Herding elephants : lessons learned from sharding Postgres at Notion, October 2021. URL https://www.notion.so/blog/sharding-postgres-at-notion.

[5] Google. Minimal downtime migrations. URL https://googlecloudplatform.github.io/spanner-migration-tool/minimal.

[6] Sangmin Lee, Zhenhua Guo, Omer Sunercan, Jun Ying, Thawan Kooburat, Suryadeep Biswal, Jun Chen, Kun Huang, Yatpang Cheung, Yiding Zhou, Kaushik Veeraraghavan, Biren Damani, Pol Mauri Ruiz, Vikas Mehta, and Chunqiang Tang. Shard Manager : A Generic Shard Management Framework for Geo-distributed Applications. In *Proceedings of the ACM SIGOPS 28th Symposium on Operating Systems Principles*, pages 553–569, Virtual Event Germany, October 2021. ACM. ISBN 978-1-4503-8709-5. doi:10.1145/3477132.3483546. URL https://dl.acm.org/doi/10.1145/3477132.3483546.

[7] Haikun Liu, Cheng-Zhong Xu, Hai Jin, Jiayu Gong, and Xiaofei Liao. Performance and energy modeling for live migration of virtual machines. In *Proceedings of the 20th international symposium on High performance distributed computing*, HPDC '11, pages 171–182, New York, NY, USA, June 2011. Association for Computing Machinery. ISBN 978-1-4503-0552-5. doi:10.1145/1996130.1996154. URL https://dl.acm.org/doi/10.1145/1996130.1996154.

[8] SiteWhere LLC. SiteWhere, 2021. URL https://github.com/sitewhere.

[9] Paarth Madan. Shard Balancing : Moving Shops Confidently with Zero-Downtime at Terabyte-scale, September 2021. URL https://shopify.engineering/mysql-database-shard-balancing-terabyte-scale.

[10] MongoDB. Migrate Ranges in a Sharded Cluster - MongoDB Manual v8.0. URL https://www.mongodb.com/docs/current/tutorial/migrate-chunks-in-sharded-cluster/.

[11] Shripad Nadgowda, Sahil Suneja, Nilton Bila, and Canturk Isci. Voyager : Complete Container State Migration. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, pages 2137–2142, June 2017. doi:10.1109/ICDCS.2017.91. URL https://ieeexplore.ieee.org/abstract/document/7980161. ISSN : 1063-6927.

[12] Eduard Gibert Renart, Daniel Balouek-Thomert, and Manish Parashar. An edge-based framework for enabling data-driven pipelines for iot systems. In *2019 IEEE International Parallel and Distributed Processing Symposium Workshops (IPDPSW)*, pages 885–894, 2019. doi:10.1109/IPDPSW.2019.00146.

[13] Guillaume Rosinosky, Chahrazed Labba, Vincenzo Ferme, Samir Youcef, François Charoy, and Cesare Pautasso. Evaluating Multi-tenant Live Migrations Effects on Performance. In Hervé Panetto, Christophe Debruyne, Henderik A. Proper, Claudio Agostino Ardagna, Dumitru Roman, and Robert Meersman, editors, *On the Move to Meaningful Internet Systems. OTM 2018 Conferences*, pages 61–77, Cham, 2018. Springer International Publishing. ISBN 978-3-030-02610-3. doi:10.1007/978-3-030-02610-3_4.

[14] Guillaume Rosinosky, Samir Youcef, François Charoy, and Etienne Rivière. A Methodology for Tenant Migration in Legacy Shared-Table Multi-tenant Applications. In Miguel Matos and Fabíola Greve, editors, *Distributed Applications and Interoperable Systems*, pages 3–20, Cham, 2021. Springer International Publishing. ISBN 978-3-030-78198-9. doi:10.1007/978-3-030-78198-9_1.