

Intern with Oracle Labs MLE!

The Multilingual Engine (MLE) team at Oracle Labs has open internship positions about extending MLE with new exciting functionalities.

Oracle

Oracle, a global provider of enterprise cloud computing, is empowering businesses of all sizes on their journey of digital transformation. Oracle Cloud provides leading-edge capabilities in software as a service, platform as a service, infrastructure as a service, and data as a service.

Oracle's application suites, platforms, and infrastructure leverage both the latest technologies and emerging ones – including artificial intelligence, machine learning, blockchain, and Internet of Things – in ways that create business differentiation and advantage for customers. Continued technological advances are always on the horizon.

Oracle Labs

Oracle Labs is the advanced research and development arm of Oracle. We focus on the development of technologies that keep Oracle at the forefront of the computer industry. Oracle Labs researchers look for novel approaches and methodologies, often taking on projects with high risk or uncertainty, or that are difficult to tackle within a product-development organization. Oracle Labs research is focused on real-world outcomes: our researchers aim to develop technologies that will someday play a significant role in the evolution of technology and society. For example, chip multithreading and the Java programming language grew out of work done in Oracle Labs.

Multilingual Engine (MLE)

Stored procedures, user-defined functions and triggers - in general server-side procedural logic - complement SQL processing and are an important part of many database applications, in particular in enterprise environments. Multilingual Engine (MLE) [1] is a feature of Oracle® Database that enables the use of modern programming languages like JavaScript for server-side procedural logic. MLE is based on GraalVM [2], a high-performance embeddable virtual machine developed at Oracle Labs that supports a multitude of programming languages. On the basis of GraalVM, MLE builds a platform for efficient, scalable program execution in Oracle® Database that is tightly integrated with processing of relational and non-relational data. Multilingual Engine has been released as a feature of Oracle® Database and Autonomous Database and initially allows the execution of JavaScript directly within the database.

The Multilingual Engine team at Oracle Labs works on further developing this platform to provide innovative features for database applications. Internships in the Multilingual Engine team at Oracle Labs offer the opportunity to work with state-of-the-art database and virtual machine technology, and have direct impact on the world's premier database system.

[1] https://docs.oracle.com/en/database/oracle/oracle-database/21/arpls/dbms_mle.html

[2] <https://www.graalvm.org/>

Internship Details

The goal of this project is to extend MLE and research new possibilities. We offer various topics depending on the skills and the interests of the candidate (topics are not limited to the ones below; see also the "Related Topics" sub-section below):

HTTP fetch in MLE

Using MLE, developers can implement all or part of their application in the database using JavaScript. The goal of this internship is to extend the MLE JavaScript platform to provide an efficient and secure implementation of the JavaScript Fetch API that allows developers to fetch resources across the network. This will not only give JavaScript developers a convenient, familiar way to make HTTP callouts and access REST endpoints, but would also allow them to use existing libraries that are built on top of the Fetch API.

WebAssembly for In-Database Applications

Combined with JavaScript, WebAssembly allows the safe, sandboxed execution of languages like Rust and C/C++. Next to the Web platform, WebAssembly is a promising technology for example for edge computing platforms. In this internship, we will explore WebAssembly as a means to enable languages like Rust and AssemblyScript for application code in database systems. We will extend the Multilingual Engine with WebAssembly support, develop SDKs for WebAssembly-based programs to access database functionality and explore performance aspects.

Co-Optimizing JavaScript and SQL Queries

Execution of procedural logic in the database manages to avoid expensive network round-trips for transaction processing workloads. Further optimization potential lies in the interaction between an in-database language runtime on the one side, and the SQL engine and data layer on the other side. Elimination of context switches and data copies, as well as holistic compilation to machine code promise further speed-ups. Concretely, this internship will investigate methods to accelerate the combined execution of JavaScript data processing logic and simple SQL statements.

This internship requires a solid understanding of some aspects of database systems implementation (in particular SQL execution layer).

Location Optimization for Data-Driven Microservices

Enterprise applications structured into microservices typically feature services that interact heavily with database systems as the persistent data store. These services can benefit from being located close to the database system in order to minimize network interaction. Deployment to FaaS platforms offers the opportunity to optimize the placement of services, for example in close proximity to the database. In this project, we will explore techniques to automatically

“During my internship I could dive deep into the inner workings of a novel language implementation framework and a JIT compiler, among many other things. The experience gave me more confidence in my abilities as a programmer and showed me where I still have room to grow. People at the company were responsive to my questions and were eager to help. Overall, I had a great time at the Labs.”

Aleksei Kashuba

EPFL student, now employed by Oracle Labs, 6-months intern

optimize the placement of data-driven microservices based on communication patterns.

Just-in-Time Compilation of Regular Expressions in Oracle® Database

In database systems, regular expressions can occur both in SQL queries, but also in JSON-based document store workloads. In high-throughput performance-sensitive database applications, efficient evaluation of regular expressions is critical. This internship will investigate the use of GraalVM's regular expression engine for SQL processing and other database components. GraalVM provides a high-performance regular expression engine that leverages JIT compilation.

Required Skills

The successful candidate is expected to complete the internship using a wide and diverse set of skills.

- Basic understanding of databases and SQL (having completed advanced database implementation courses is a plus)
- Experience with C and Java programming. Experience with JavaScript is a plus

For more information about the internship, please contact Martin Keppner (martin.keppner@oracle.com) or Lucas Braun (lucas.braun@oracle.com)

Related Topics

Our group at Oracle Labs further offers various internship topics in the following areas:

- Automated Machine Learning with Explainability (AutoMLx)
- Automating OCA Verification of GitHub Pull Requests
- BPF Linux Schedulers
- Extending a Distributed Graph Engine (Oracle Labs PGX)
- Extending a Web-Based Enterprise Data Science Platform
- Graph Machine Learning at Oracle
- Graph Support in the Oracle Database
- Machine Learning and Data Analysis Techniques for Domain Global Graphs
- Machine Learning for Optimizing Oracle Database Performance
- Machine Learning Processing in DB Systems

If you are interested in more than one of these areas, it is sufficient to apply once. In our interview process, we are going to take all your areas of interest into account.