Title: GraalVM and MLE: How to embed a multilingual virtual machine into the Oracle Database

Abstract

In this talk, we present GraalVM, a universal virtual machine that allows just-in-time compilation and execution of applications written in a variety of different programming languages. We introduce some basic concepts such as partial evaluation and just-in-time compilation, before we show how GraalVM uses this to turn a language interpreter into a self-optimizing compiler with speculative optimizations. We will show how such a speculative just-in-time compiler can significantly improve the performance of typical data processing tasks like encoding and decoding JSON data. Finally, we demonstrate how a database system can benefit from such techniques in the context of the Oracle Database Multilingual Engine - the embedding of GraalVM in the Oracle Database.

Main Presenter: Lucas Braun <lucas.braun@oracle.com>
Bio:
Dr. Lucas Braun received his doctoral degree from the Computer Science Department (Systems Group) of ETH Zurich in the beginning of 2017 and joined Oracle Labs Zurich right afterwards. At Oracle Labs, Lucas is part of the team that develops the Oracle Database Multilingual Engine (MLE). This engine allows to use modern languages and tools to write and deploy database code and has been released as part of Oracle Database 21c in December 2020. MLE is enabled by GraalVM, a universal virtual machine for running applications written in JavaScript, Python, Ruby, R, JVM-based languages like Java, Scala, Kotlin, and LLVM-based languages such as C and C++. During his doctoral research, Lucas studied the performance and confidentiality of cloud databases. Lucas also has a Computer Science teaching degree and has given various guest lectures at the applied universities of Northwestern Switzerland (FHNW) and Zurich (ZHAW).

Intro Presenter: Vasileios Trigonakis <vasileios.trigonakis@oracle.com>
Bio:
Vasileios Trigonakis is a researcher at Oracle Labs Zurich, where he is leading the efforts to develop PGX Distributed, a fast and scalable distributed graph-processing engine. Apart from graph processing, his main research interests include concurrent, parallel, and distributed runtimes and programming. Vasileios received his PhD (titled "Towards Scalable Synchronization on Multi-Cores") from EPFL Switzerland in 2016. His dissertation was awarded the 2017 EuroSys Roger Needham PhD award for the best thesis on Software Systems in Europe.
During his PhD, Vasileios joined Oracle Labs Cambridge for a three-month internship, working on scalable runtimes for multi-core processors.