

# Granules: A Lightweight Runtime for Cloud Computing

Shrideep Pallickara  
Indiana University

Scientists in several domains are grappling with vast volumes of data. The proliferation of networked sensors, observational equipments, and measuring instruments has led to a sustained increase in data volumes. Most scientists lack the machinery to process and manage the massive datasets that are being produced in their domains. Cloud computing offers the prospect of meeting these challenges.

In this talk, I will describe Granules, which is a lightweight streaming-based cloud runtime. Granules simplifies the development of resource-intensive scientific applications by insulating them from the complexities inherent in these large-scale settings.

Granules orchestrates the execution of an application's computations concurrently on multiple machines to achieve significant execution speedups. Additionally, at a given computational resource, Granules interleaves the execution of multiple computations for optimal resource utilizations.

The dominant programming model in cloud computing settings is Map-Reduce. In Granules, I have incorporated support for the Map-Reduce framework. In addition to traditional static datasets such as files, Granules also incorporates support for dynamic, streaming datasets. I will describe this work, and also the support for variants of the basic Map-Reduce framework that make Granules particularly suitable for scientific applications.

The Granules project utilizes NaradaBrokering for all its content dissemination needs. I will touch upon some of the key innovations in NaradaBrokering and how they relate to my research activities in Granules.

## **Bio**

Dr. Shrideep Pallickara is a Researcher at the Community Grids Labs at Indiana University. He received his PhD in Computer Engineering from the Department of Electrical Engineering & Computer Science at Syracuse University. He is the architect of two projects in the area of large-scale systems: Granules and NaradaBrokering. His research has been funded through two grants from the United States National Science Foundation, and two grants from the United Kingdom's Open Middleware Infrastructure Institute.