

## Ignis: an efficient and scalable multi-language Big Data framework

**Título** Ignis: an efficient and scalable multi-language Big Data framework

**Autores** César Piñeiro, Rodrigo Martínez-Castaño and Juan C. Pichel

**Tipo** Artículo de revista

**Fonte**  [Future Generation Computer Systems](#), ELSEVIER, Vol. 105, pp. 705-716 , 2020.

**Rank**  [Provisionally ranked Q1 in Hardware and Architecture by SJR 2018](#)

**ISSN** 0167-739X

**DOI** [10.1016/j.future.2019.12.052](https://doi.org/10.1016/j.future.2019.12.052)

**Abstract** Most of the relevant Big Data processing frameworks (e.g., Apache Hadoop, Apache Spark) only support JVM (Java Virtual Machine) languages by default. In order to support non-JVM languages, subprocesses are created and connected to the framework using system pipes. With this technique, the impossibility of managing the data at thread level arises together with an important loss in the performance. To address this problem we introduce Ignis, a new Big Data framework that benefits from an elegant way to create multi-language executors managed through an RPC system. As a consequence, the new system is able to execute natively applications implemented using non-JVM languages. In addition, Ignis allows users to combine in the same application the benefits of implementing each computational task in the best suited programming language without additional overhead. The system runs completely inside Docker containers, isolating the execution environment from the physical machine. A comparison with Apache Spark shows the advantages of our proposal in terms of performance and scalability.

**Palabras clave** Big Data, Multi-language, Performance, Scalability, Container

## LIGAZÓNS

 [Versión da editorial](#)

## DESCARGAS

 [Referencia BibTex](#)

## PROXECTOS DE INVESTIGACIÓN

eRISK: Tecnoloxías para a predición temperá de signos relacionados con trastornos psicolóxicos

## PROGRAMAS CIENTÍFICOS

Computación avanzada

Tecnoloxías da Linguaxe Natural