

CoherentPaaS

CoherentPaaS aims at providing a full ACID coherent, scalable and efficient environment integrating NoSQL, SQL and CEP data management technologies, while allowing application developers to program using a simple and powerful programming model and query language.

AT A GLANCE

Project title:

Coherent and Rich PaaS with a Common Programming Model

Project coordinator:

Ricardo Jimenez-Peris, Universidad Politecnica de Madrid, SPAIN

Partners:

Universidad Politecnica de Madrid (ES), Neurocom SA (GR), INRIA (FR), Foundation for Research and Technology - Hellas (GR), Institute of Engineering Systems and Computers (PT), Sparsity (ES), MonetDB Solutions (NL), QuartetFS (UK), Portugal Telecom (PT), Institute of Communication and Computer Systems (GR), LeanXcale (ES)

Duration:

36 months (Oct. 2013-Sep. 2016)

Total cost:

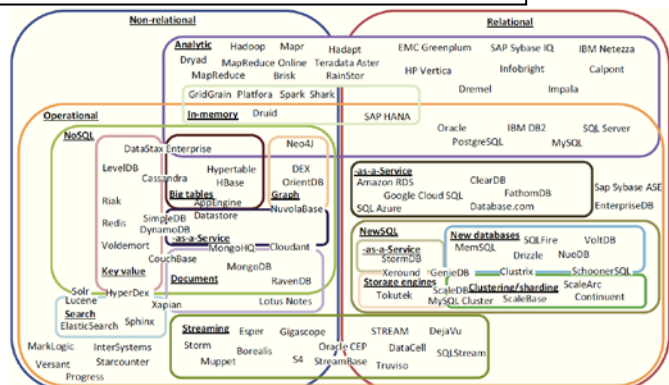
€6,54 M

Programme:

ICT-2013.1.2 Software Engineering, Services and Cloud Computing

Further information:

<http://coherentpaas.eu>



Why CoherentPaaS?

The increasing demand for efficiency has resulted in a *wide diversity of different cloud data stores*, each one specialized and optimal for specific applications, thus leading to a **“no one size fits all” situation**. The consequence refer to large *proliferation of APIs*, *lack of a common programming framework* and *lack of coherence* across different cloud data management technologies (in the cloud landscape ACID coherence has been totally lost). Performance and efficiency issues are of major importance for *Big Data* analytics where emphasis is put on large queries over big cloud data stores.

What makes CoherentPaaS unique?

CoherentPaaS provides a **rich PaaS with different “one size” data stores** optimized for particular tasks, data, and workloads. CoherentPaaS integrates

NoSQL data stores such as key-value data stores and graph databases, SQL data stores such as in-memory and column-oriented databases, hybrid systems such as SQL engines on top on key-value data stores, and complex event processing data management systems.

What is more, CoherentPaaS provides a **common query language, CloudMdsQL**, for all data stores. CloudMdsQL enables to combine the simplicity of SQL with the power of the native query languages of the underlying data stores, thus, enabling to fully exploit efficiently the capabilities of each cloud data store. CoherentPaaS enables the development of new cloud applications that exploit the performance and scalability of new cloud data management technology, while hiding the complexity of the underlying technology under a unified query language.

Additionally, a scalable transactional management system provides **holistic transactional coherence across data stores**. CoherentPaaS has designed and implemented within each data store the necessary functionality to enable the integration with the new transactional technology in a scalable way, providing a common transactional API to enable the holistic transactional management in order to orchestrate transactions across data management technologies transparently to the applications.

Value proposition

CoherentPaaS enables the development of cloud applications that exploit the

performance, coherence and scalability of an environment providing an integrated cloud data management technology; while hiding the complexity and diversity of the underlying technology under a simple and powerful programming model and query language. Thus, the added value of CoherentPaaS is summarized as follows:

- ✓ **Simplicity:** Programming combining the simplicity of SQL and the power of the native query languages.
- ✓ **Scalability:** Across all data management technologies.
- ✓ **Efficiency:** Avoiding copying and translating TBs of data from one data store to another (i.e. ETL jobs).
- ✓ **Coherence:** Transactional semantics across data stores.

Demonstrators

The CoherentPaaS outcomes are being validated through four different application scenarios: Cloud Machine-to-Machine for rating, invoicing and analysis of M2M events, and fraud identification and management; Real-Time Network Performance Analysis for collection and real-time analysis of information in a telco environment; Media Planning for analysis of social media, news or blogging data; and Bibliographic Search for exploiting information from various data sources.

Exploitation through a Startup

A startup, LeanXcale, has been created to exploit the core results of the project with personnel transferred from the coordinator (UPM) and a partner (INESC).

