CoherentPaaS

Coherent and Rich PaaS with a Common Programming Model

ICT FP7-611068

X-Ray Subsystem Implementation (final version)

D8.5

(Prototype)

March, 2016

Document Information

Scheduled delivery	30.03.2016
Actual delivery	30.03.2016
Version	1.0
Responsible Partner	INESC

Dissemination Level:

RE Restricted to a group specified by the consortium (including the Commission)

Nature:

Prototype

Revision History

Date	Editor	Status	Version	Changes
27.08.2015	J. Pereira	Draft	0.1	Initial revision.
15.09.2015	J. Pereira	Draft	0.2	Corrections after internal revision.
20.03.2016	J. Pereira	Draft	0.3	Revision for A&V and software 1.0.
30.03.2016	J. Pereira	Final	1.0	Final version.

Contributors

P. Guimarães (INESC TEC), J. Pereira (INESC TEC)

Internal Reviewers

Sotiris Stamokostas (ICCS/NTUA), Ying Zhang (MonetDB Solutions)

Acknowledgements

Research partially funded by EC 7th Framework Programme FP7/2007-2013 under grant agreement n° 611068.

More information

Additional information and public deliverables of CoherentPaaS can be found at: http://coherentpaas.eu

1. Executive Summary

The integration of multiple database technologies, including both SQL and NoSQL, allows using the best data management tools or each aspect of a complex data-centred application and is increasingly sought in practice. Unfortunately, this makes it difficult for database developers and administrators to obtain a clear view of the resulting composite data processing paths, as they combine operations chosen by different query optimisers, implemented by different software packages, and partitioned across distributed systems.

In CoherentPaaS this challenge is addressed with the X-Ray Subsystem, a framework for monitoring and analysis of distributed and heterogeneous data processing systems. First, it provides a way to add monitoring code to applications and data stores running on the Java platform. By using bytecode instrumentation, it does not rely on the availability of the source code and can be applied conditionally to avoid overhead in production systems. In addition, it provides interfaces for integration of C/C++ applications. Second, it provides mechanisms for tracking the interaction of multiple threads, on synchronisation primitives, and of distributed processes communicating with sockets.

The analysis and visualization components have the ability to track and display requests across distribution and heterogeneity boundaries in the CoherentPaaS system. They offer different views of a data processing operation, showing both the composition and structure of elementary operations, the resources used for those operations, and their performance.

Experimental evaluation using the Apache Derby query engine and the industry standard on-line transaction processing benchmark TPC-C, show that X-Ray does in fact show data processing requests across multiple processes, and that results in minor overhead compared to the original setting.