

Big Data Matters

Why Financial Services Firms Need Platform MapReduce Now

In a recent Wall Street & Technology poll sponsored by Platform Computing and SAS Institute, both Sell-side and Buy-side firms of all sizes and types identified data volume and quality as key challenges in analytics. Agile analytics and powerful processing technology rank high on investment priorities over the upcoming year. Predictive analysis of both internal and external data result in better, proactive management of a wide range of issues from: credit and operational risk e.g. fraud and reputational risk; customer loyalty and profitability; etc.

Data growth – particularly of unstructured data e.g. email – poses a special challenge as the volume and diversity of data types outstrip the capabilities of older technologies such as relational databases. Financial Services firms are investigating next generation technologies for data analytics. One of the most promising technologies is the Apache Hadoop software and MapReduce framework for dealing with this “big data” problem. However current implementations of Hadoop lack the necessary enterprise robustness demanded by Financial Services firms for wide-scale deployment of MapReduce applications.

Introducing Platform MapReduce, the industry’s first enterprise-ready MapReduce solution with 100% Hadoop compatibility.

Available June 2011.

Platform MapReduce is an enterprise-class server platform designed to run all MapReduce applications in the most demanding production environments of Financial Services firms. Built on the Platform Resource Orchestrator, the

proven resource manager - deployed at 12 of Top 20 banks worldwide, Platform MapReduce is an enterprise-ready, enhanced version of the MapReduce framework including feature-to-feature compatibility with that included within the Apache Hadoop environment.

Platform MapReduce guarantees enterprise-class reliability, manageability, and scalability. Platform MapReduce solves the major runtime risks faced by organizations in deploying MapReduce applications on currently available Apache Hadoop implementations. Key enhancements include:

- High Availability and Disaster Recovery for guaranteed low mean-time to failure, and ease of management.
- Automated Resource Utilization and Management leading to optimized allocation of IT resources and higher return on infrastructure.
- Scalability to meet variable workload requirement
- Open Architecture to support heterogeneous development environments, applications and file systems beyond the Apache Hadoop subprojects such as Hadoop Distributed File System (HDFS)

Platform MapReduce includes (See Figure 1):

- Platform MapReduce Workload Manager that automates, distributes, and manages MapReduce workloads according to users’ Service Level Agreements (SLAs).
- Platform Resource Orchestrator that allocates and manages distributed pools of resources including clusters, servers, and cores.

- Application Programming Interfaces (APIs) that support:
 - Apache Hadoop subprojects including Hadoop Java MapReduce API, Pig, and Hive
 - Enhanced APIs including Python
 - Advanced file system and data access framework for connectivity to file systems and database architectures, thus eliminating the need to migrate existing data. Initial support includes HDFS, Appistry, GPFS, and Oracle.
- Rich set of management, troubleshooting, and reporting tools with a single GUI interface for managing multiple MapReduce applications across a shared set of resources. Full version compatibility supports rolling upgrades.

For more information, please visit www.platform.com or contact your local Platform Computing sales office.

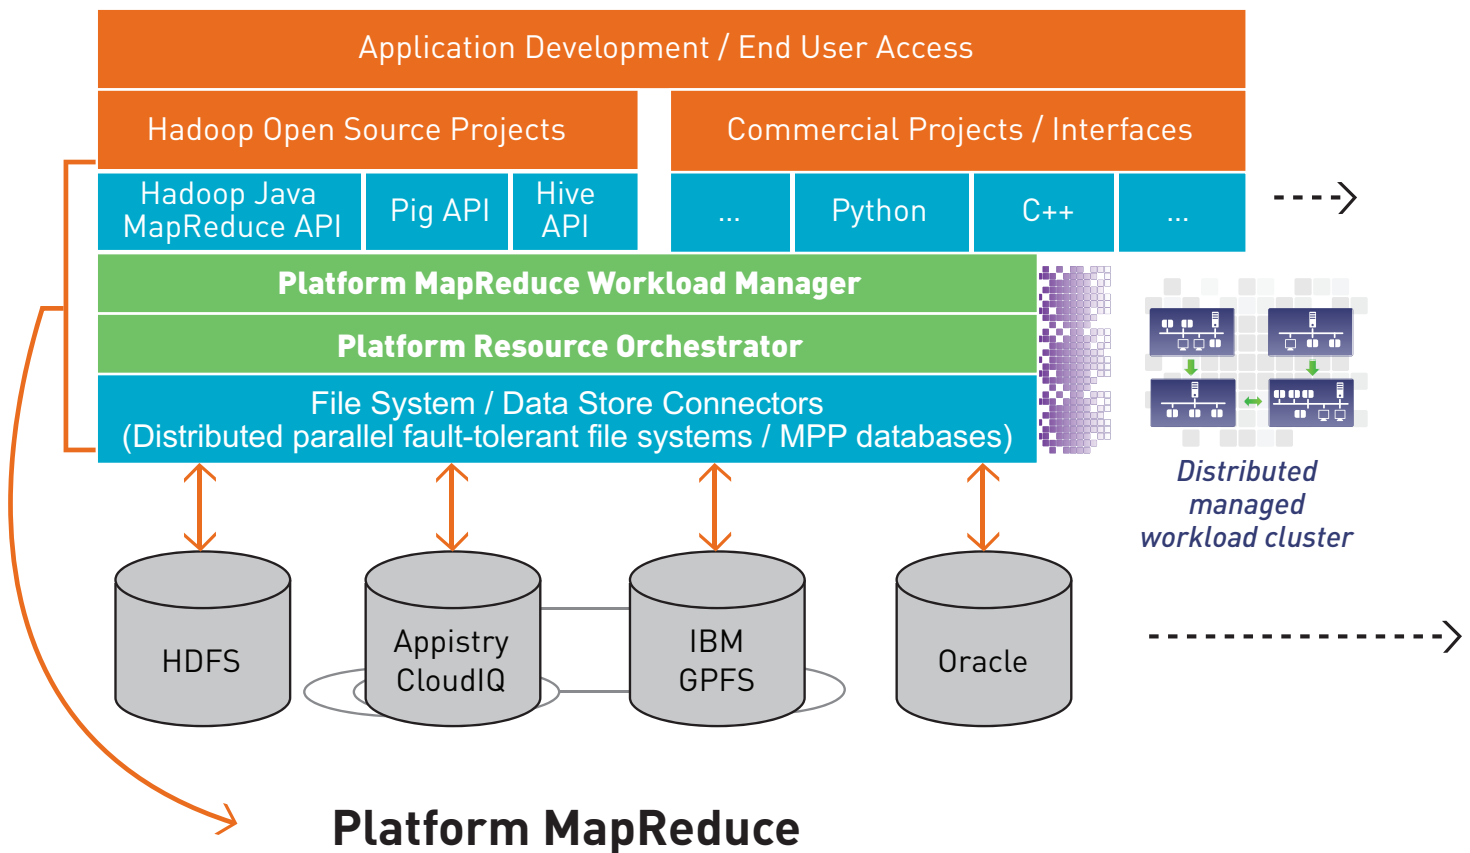


Figure 1. Platform MapReduce Architecture



Platform Computing is the leader in cluster, grid and cloud management software - serving more than 2,000 of the world's most demanding organizations for over 18 years. Our workload and resource management solutions deliver IT responsiveness and lower costs for enterprise and HPC applications. Platform has strategic relationships with Cray, Dell™, HP, IBM®, Intel®, Microsoft®, Red Hat® and SAS®. Visit www.platform.com.

World Headquarters
 Platform Computing Corporation
 3760 14th Avenue
 Markham, Ontario
 Canada L3R 3T7
 Tel: +1 905 948 8448
 Fax: +1 905 948 9975
 Toll-free Tel: 1 877 528 3676
info@platform.com

Sales - Headquarters
 Toll-free Tel: 1 877 710 4477
 Tel: +1 905 948 8448

North America
 New York: +1 212 888 6270
 San Jose: +1 408 392 4900

Europe
 Bramley: +44 (0) 1256 883756
 London: +44 (0) 20 3206 1470
 Paris: +33 (0) 1 41 10 09 20
 Düsseldorf: +49 2102 61039 0
info-europe@platform.com

Asia-Pacific
 Beijing: +86 10 82276000
 Xi'an: +86 029 87607400
asia@platform.com
 Tokyo: +81(0)3 6302 2901
info-japan@platform.com
 Singapore: +65 6307 6590
wliaw@platform.com