

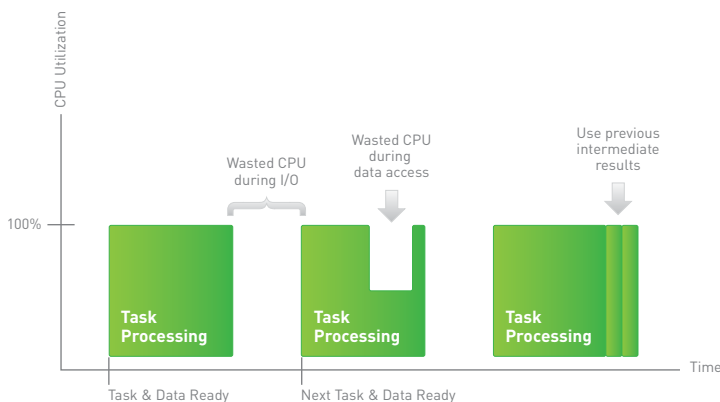
Solution Overview

Platform Symphony Data Affinity

A breakthrough in data access bottlenecks

Have you ever wished that you could:

- Do more accurate risk assessment and still meet regulatory reporting deadlines
- Leverage the power and cost savings of grid for front office applications
- Co-locate your data and compute grids efficiently and effectively
- Reduce the amount of time spent moving data around the grid – by up to 4x



Without Platform Symphony Data Affinity there is a potential latency as each task waits for the data to be transported to the network and loaded into memory.

What is Data Affinity for Platform Symphony?

As processors increase in performance and grid farms scale to 10s of 1,000s of nodes, it becomes impossible to keep up with throughput requirements by simply scaling up NAS architectures. Platform Symphony Data Affinity is a breakthrough in solving the latency bottlenecks caused by moving large data sets around on a compute grid. This add-on product allows Platform Symphony to intelligently schedule application tasks and significantly improve performance by taking into account data locality in order to place workload in the most optimal and cost efficient way.

Platform Symphony Data Affinity ensures that financial services organizations can:

- Move from scale-up of expensive hardware to scale out of inexpensive off the shelf servers
- Cut task run times by up to 80%
- Improve grid utilization by on-boarding new applications that previously were difficult to run in a shared environment
- Reduce reliance on expensive network storage hardware

Already own a data caching solution or parallel file system?

While Data Affinity for Platform Symphony does not require data caching or parallel file system, it is integrated with leading distributed cache systems to seamlessly tie into your company's data management and optimization initiatives:

- Oracle Coherence
- GemFire
- GigaSpaces XAP
- HP IBRIX Fusion parallel file system

As you distribute and cache data across the nodes in a grid, Platform Symphony Data Affinity intelligently places workload to ensure optimal performance with minimal data-staging wait-times.

Where to deploy Platform Symphony Data Affinity

Middle office

Platform Symphony Data Affinity virtually eliminates the time it takes to load large (giga- and terabyte) data volumes required by applications such as counterparty risk valuations. By placing tasks on nodes that already contain the required data you can accelerate time to results by up to 80%. This effectively scales the application, allowing you to run more calculations in the same amount of time.

- Enables better, more accurate risk assessment
- Relieves pressure on meeting regulatory reporting deadlines
- Eliminates the need to buy expensive network storage equipment by replacing it with cost-effective storage on the grid

Front office

Many real-time, data-intensive applications (e.g. bond pricing) are still confined to expensive SMP servers since the time it would take to load the required data onto a grid far exceeds the 2-3 second calculations required. Platform Symphony Data Affinity eliminates the latency bottleneck that currently prevents on-boarding these applications to the grid.


- Front office applications can leverage the cost-reduction and linear scalability benefits of grid by running on commodity hardware with no increase in latency
- Improves the overall utilization of the grid within the organization by sharing resources among more applications
- Improves application performance and scalability by efficiently reusing intermediate results

Location aware data services

Platform Symphony Data Affinity enables consolidation of multiple layers of server, network, and storage infrastructure in order to increase the performance and manageability of data-intensive applications; such as historical pricing analysis. While the data required for these applications is stored in a way that is typically easy to find, accessing the files creates a flurry of network activity that has a less than optimal effect on overall performance.

Platform Symphony Data Affinity solves the data bottleneck problem by intelligently placing workload on systems that are physically close to the required data – the benefits:

- Network traffic between nodes in the cluster is reduced or eliminated
- Faster file access increases overall application performance (some studies show improvements of 2x to 4x)
- Consolidation of previously separate layers of server, network and storage infrastructure lead to savings in capital and operating costs



Without Platform Symphony Data Affinity, 16% of tasks are sent where the data is already stored. Platform Symphony Data Affinity increases this to 90+% of tasks.



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 17 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 Inc.
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: +1 646 290 5070
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