IBM Platform Computing Solutions: Banking

Create a high utility, high-performance environment for deeper analytics

Highlights

- Accelerate fraud analytics with high-performance big data solutions
- Enable deeper customer analytics that improve profitability and retention
- Cost-effectively perform analysis on demand for faster time to results

Banks today are facing intense pressure on multiple fronts. Regulators are becoming more prescriptive, requiring more detailed and frequent disclosure of information and undertaking more proactive risk management. Along with growing competition, more sophisticated financial crime and more complex customer demands, banks require more innovative business models and analytics to meet these challenges and to deliver growth and performance.

IBM solutions such as IBM® Platform Computing™ cluster, grid and high-performance computing (HPC) cloud management software and the IBM General Parallel File System (GPFS™) have helped banks cost-effectively accelerate data analysis of all types for improved business decisions on a timely basis.

Mining data for improved risk management and customer insights

Billions of dollars per year continue to be "laundered" or processed in the form of loan applications, card services and other transaction types. For instance, fraudulent Card Not Present (CNP) sales and e-payments can cost banks as much as 7 percent of sales transactions—if not more. Often, petabytes of data have been collected but not analyzed by the bank due to lack of technical resources, interest or subject-matter knowledge.

In many cases, the bank's first and primary focus is on the analysis of structured data; analysis of unstructured data such as images or correspondence likely remains a secondary focus or is not done.

Banks need a faster and more cost-effective way to expand their data analysis to include unstructured data for improving fraud analytics, whether perpetrated by customers, employees or partners. Big data technologies such as Hadoop MapReduce enable analysis of petabytes of unstructured data cost-effectively. Such technologies can help banks better anticipate a customer's most likely actions, determine "next best actions" and create more targeted offers and services.

Accelerating results for compute- and data-intensive applications

Firms developing big data analytics can choose the IBM InfoSphere® BigInsights™ big data solution. Adaptive MapReduce (AMR), based on IBM Platform™ Symphony - Advanced Edition (AE) grid management software, is included in InfoSphere BigInsights. With Adaptive MapReduce enabled, InfoSphere BigInsights applications can deliver up to four times faster time to results according to a STAC benchmark.[1] For firms developing their own big data solutions, Platform Symphony AE accelerated select Hadoop MapReduce workloads by seven times according to another STAC benchmark.[2]

In addition to big data analytics, Platform Symphony AE software supports traditional compute-intensive analytics such as credit risk analysis from IBM Algorithmics®, Murex and Calypso. Organizations already using InfoSphere BigInsights can add Platform Symphony AE to create a single platform for faster traditional risk and big data analytics (see Figure 1). Banks that use Platform Symphony AE can significantly accelerate time to results for compute- and data-intensive applications of all types for better decision making.

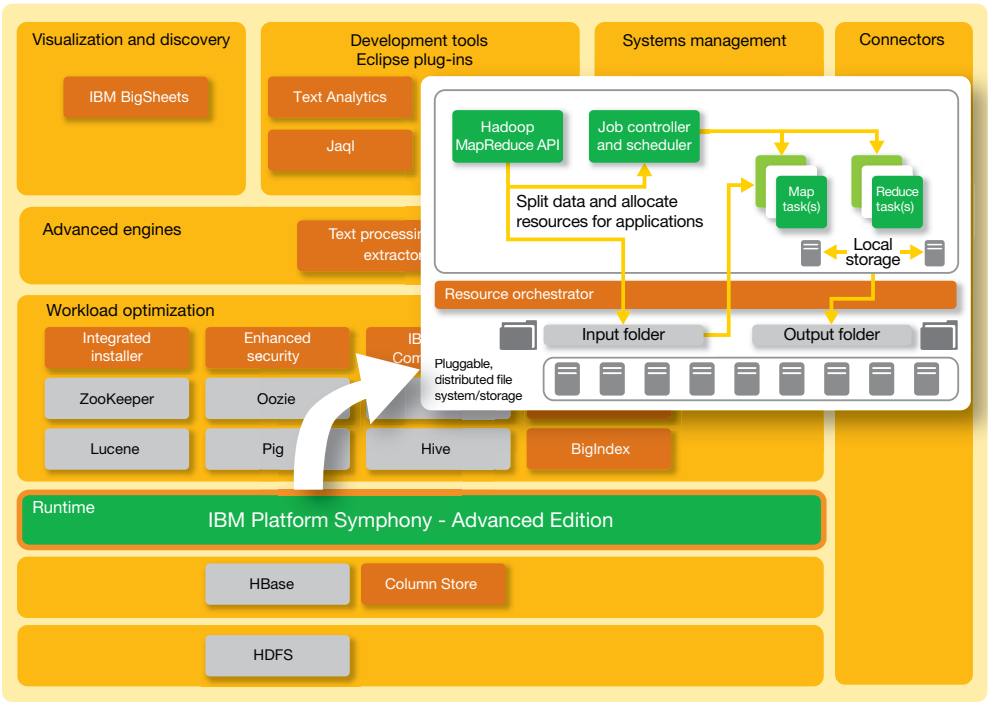

Figure 1. IBM Platform Symphony AE accelerates both compute and data-intensive analytics.

Eliminating file system bottlenecks

When running data-intensive applications such as big data analytics, file servers and their network connections can become overwhelmed, while I/O bottlenecks can lead to project delays. Alleviating the problem by adding more file servers creates additional costs and IT headaches.

The solution:

- GPFS removes data-related bottlenecks by providing parallel access to data, while retaining investments in existing network file servers through its support of Active File Management (AFM).
- If a file server is in high demand, GPFS AFM can cache the data and deliver it through multiple network connections.
- By eliminating hot spots, this approach helps accelerate data access, computation and time to market.

For clients running workloads where data is distributed rather than localized—for example, Hadoop MapReduce—GPFS supports a file-placement option (FPO) to maximize performance by ensuring that processing takes place right on the nodes where the data resides.

Greater efficiency and cost savings

Platform Symphony grid management software helps organizations gain more utility from existing hardware resources. Platform Symphony allows companies to pool compute resources into a common shared environment for greater scalability and performance. Provisioning time is also reduced. Rather than wait for days, weeks or even months for IT to provide IT capacity, users gain additional capacity on-demand as needed. By creating a shared environment, firms have deferred millions of dollars in hardware purchases by improving utilization rates of existing IT resources.

Platform Computing Cloud Service

Platform Symphony is also available as a Platform as a Service (PaaS). With Platform Computing Cloud Services, banks now have the option of cost-effective subscription-based service for traditional risk analytics.

Application Ready for Algorithmics

IBM has a new addition to its Platform Computing portfolio that can help financial services clients deploy, run and manage big data and analytics environments with increased efficiency and agility. IBM Application Ready Solution for Algorithmics provides an expertly designed, tightly integrated and performance-optimized architecture for IBM Algo One® solutions. IBM Application Ready Solutions are developed in partnership with leading ISVs and can accelerate production readiness of applications by up to 75 percent. This solution, based on IBM servers running Algo One integrated with Platform Symphony software, is designed to optimize risk analytics performance and reduce total cost of ownership, while enabling users to become productive more quickly.

Why IBM?

Most banks are continually evolving their systems to support new analytics. Platform Computing high-performance grid and data management software is designed to support heterogeneous systems and wide range of analytics from credit risk to fraud analytics. Platform Computing software allows banks to build a common compute and data platform that is multipurpose, thereby providing greater utility and cost efficiency. By combining Platform Computing software with powerful big data analytics systems such as InfoSphere BigInsights, end users can gain faster time to results without needing to focus on infrastructure management. Solutions such as Platform Computing Cloud Services and Application Ready for Algorithmics accelerate time to value with lower implementation time and cost.

For more information

To learn more about IBM InfoSphere BigInsights and IBM Platform Computing solutions, please contact your IBM representative or IBM Business Partner, or visit the following website:
ibm.com/platformcomputing/industries/financialmarkets.html

Additionally, IBM Global Financing can help you acquire the IT solutions that your business needs in the most cost-effective and strategic way possible. We'll partner with credit-qualified clients to customize an IT financing solution to suit your business goals, enable effective cash management, and improve your total cost of ownership. IBM Global Financing is your smartest choice to fund critical IT investments and propel your business forward. For more information, visit: **ibm.com**/financing

IBM Corporation
Systems and Technology Group
Route 100
Somers, NY 10589

Produced in the United States of America
June 2014

Actual available storage capacity may be reported for both uncompressed and compressed data and will vary and may be less than stated.

[1] For more information about these results, see the white paper, "Analyzing big data at the speed of business." **ibm.com**/common/ssi/cgi-bin/ssialias?infotype=SA&subtype=WH&htmlfid=DCW03042USEN

[2] For an audited STAC Report commissioned by IBM, visit: **ibm.com**/systems/technicalcomputing/platformcomputing/products/symphony/highperfhadoop.html

Please Recycle

DCS03061-USEN-00