IBM Analytics
Thought Leadership White Paper

Banking breakthroughs

Big-data-powered, next-generation asset liability and liquidity risk management

Contents

2 Investing in business insights

4 More informed asset liability management

5 Implementing with expertise

7 Insights accessible on demand

8 Plumbing an efficient big-data-enabled architecture

10 Fortune favors the bold

11 About the authors

Investing in business insights

Under the profitability burden of new banking regulations, a growing number of banks are betting on big data technology to drive future growth. Over the last year, several banks have advanced big data experiments into production-ready platforms, and IDC predicts that by 2019 banking industry spend on big data and analytics will exceed USD 187 billion.[1] Investments in big data are expected to help banks test new business strategies, reduce the operational costs of business analytics and increase revenues with new client-driven solutions.

For banks, growing profitability requires competing against their peers and the increasing number of FinTech firms. Smaller FinTech rivals are seeking to penetrate financial markets by building new offerings on the latest technology, including big data. Established banks that once regarded open source big data approaches like Apache Hadoop as experimental, now recognize that the technology is undeniably mainstream. Banks using more traditional data management solutions see the potential for big data to dramatically reduce the operating costs of analytics by efficiently distributing both data and analytical models across clusters of low-cost servers, and aggregate results at high speeds. Many banks are already improving their competitiveness with insights from big data projects, and see the expansion of these operations as the beginning of a long-term digital transformation.

Among banks already using big data approaches, several have reported a backlog of business applications waiting for the necessary resources to deploy big data. In assessing the priority of big data projects, teams are looking for applications that will use a high volume of “clean” and readily identifiable data. They also want applications that will yield a short-term return on investment and lay the foundation for the bank’s broader transformation towards an increasingly digital operating model.

These are reasons why big data project teams within banks are beginning to apply big data analytics for smarter asset liability management (ALM) and liquidity risk management. Proper ALM and liquidity risk management is critical to the success of any bank, which covers a wide scope of activities. These activities range from the day-to-day management of interest rate risk and liquidity risk exposures, to enterprise-wide balance-sheet optimization, which analyzes the impact of various risk factors across a range of scenarios. Effective and coordinated management of a bank's entire balance sheet must bring together teams from both finance and risk management. These teams must be equipped with the analytics needed to develop more risk-aware business plans for increasing profitability and successfully weathering future financial crises.

Figure 1: Asset liability management at the intersection of finance and risk management departments.

Dramatic client success was achieved by applying big data to increase the granularity and interactivity of the analysis, and reduce the runtime		
	IBM Algo One + traditional data management	**IBM Algo One + big data technology**
Data granularity	**350 thousand pooled positions** Processed to reduce granularity from six million source positions	**6 million source positions**
Interactivity with reporting results	**Limited drill down** Liquidity risk reporting generated as part of the overall batch run is limited, where drill down into the underlying data requires other interfaces	**Full drill down to source** Reporting generated supports drill down, and slice and dice—all the way down to interpreting IBM® Algo One® with IBM RiskWatch® inputs, outputs and log files
"What-if" analysis	**Full rerun required** Adding new "what-if" stress scenarios requires a rerun of the batch process	**Incremental updates** New "what-if" stress scenarios can be computed by the user adding more scenarios to the table
Runtime	**2 hours**	**20 minutes**

Table 1: Actual client results from a big data implementation

More informed asset liability management

For banks, undertaking a more rigorous approach to balance-sheet risk analysis demands the scale of big data approaches that can reliably distribute the high analytical workloads involved. Versus big data, applying traditional analytical approaches on millions of complex cash flows across many scenarios would be too costly and slow to effectively support decision making. Until now, banks have managed to cope by reducing the data granularity from many millions of source positions into several hundred thousand pooled positions.

This compromise has enabled firms, using traditional approaches, to execute the analysis by way of overnight batch runs. Unfortunately, this pooling approach not only impacts accuracy, it also limits validation and troubleshooting, because users are unable to drill into the underlying data that was pooled.

To deliver more informed and flexible analysis, IBM worked with the center of excellence (CoE) team at a large Canadian bank to implement a big data approach for liquidity stress testing. The team analyzed over 15 times more data in less than one-fifth the previous runtime.

Liquidity stress tests measure a bank's ability to meet its payment obligations on time at reasonable costs under strained conditions. By applying big data approaches, banks can fully re-evaluate liquidity risk with the flexibility of running "what-if" scenarios through incremental updates. This approach aids informed funding decisions, which are especially important during times of high market volatility.

"These results have determined that big data can achieve the scale we need for liquidity stress testing without revising our existing risk models," says the CoE team leader at the bank.

Because liquidity crises tend to occur suddenly, a bank is expected to continuously manage liquidity, analyzing possible impacts on the bank's cash flow profile, profitability and solvency. The timing of cash inflows and outflows, both off- and on-balance, and the evolution of unencumbered liquid assets that can be sold or repoed to generate cash, must be considered in the liquidity stress test.

The team leader at the bank explains: "The flexibility of our big data approach to liquidity stress-testing enables us to easily change our assumptions about the realization of cash inflows and outflows and the extent to which assets can be liquidated or used as collateral — it has allowed us to be much more specific and confident in the reporting we present to committee."

Implementing with expertise

Center of excellence teams at the bank include the most experienced technology professionals and a select group were chosen to lead the big data project for ALM and liquidity risk management. "Most of us started our careers in financial services as programmers building solutions for banks, either working within banks or for software vendors focused on the banking industry. We've gained our financial engineering knowledge over years of experience in detailed project work." This team was confident in their ability to lead the big data project for liquidity stress testing, and knew they could get results faster by partnering with trained professionals who have years of experience in implementing big data technology.

"IBM was the natural choice; we've built strong relationships with many of their consultants who know our systems just as well as we do. We wanted a partner who could help execute the big data implementation for this project, and who could continue to offer support as we make progress towards the bank's long-term vision for big data. The success of this project is a big step towards infrastructure reforms that will enable us to do more with less, and significantly reduce IT infrastructure costs for the bank."

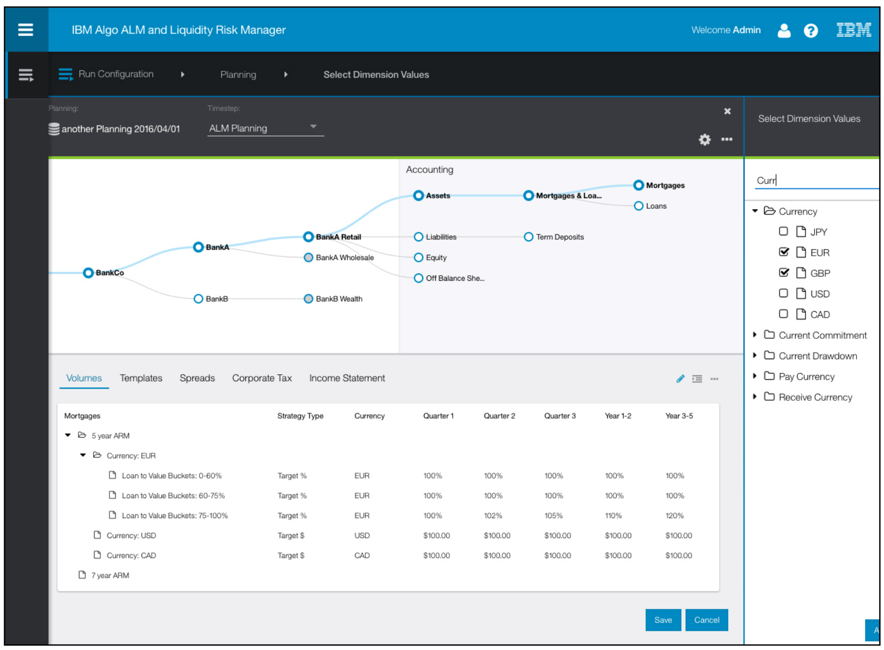

Figure 2: Interface for asset liability management and liquidity risk management

With more granular approaches, banks can reliably manage the risks arising from asset and liability mismatches and test risk-aware strategies for increasing profitability through alternative asset and liability structures, generating higher spreads.

Efficiently modeling alternative structures for the asset and liability mix requires interfaces where end-users are in full control of the setup parameters for the analytics.

Figure 3 is one example of an interface for configuring the product templates, hierarchy structures, reporting time steps and modelling approaches used in the analysis.

Powered by a big data infrastructure that can quickly output more granular results, the application provides decision makers with a comprehensive and intuitive way to execute direct control over the analysis. In fact, it's expected to revolutionize current practices within the banking industry.

Insights accessible on demand

Until recently, a lack of computing power, isolated data and analytical approaches forced the bank to do business and then, compute its risks. The CoE team leader thinks they can do better: "Long term, we believe that further consolidating more of the bank's data into big data platforms will enable us to run more holistic and on-demand analysis across departments." Many banks have spent years on projects aimed at consolidating information pools into a big data lake. These consolidations have included unravelling and redirecting data feeds from front-office trading data, middle-office risk and operations data, and back-office financial data.

By moving towards an integrated big data architecture, with a clean pool of data to draw from, decision makers across departments can be confident that consistent data will drive trustworthy results. With the performance gains of big data analytics, decision makers can run on-demand analytics to assess the risk of virtually any major transaction or change in strategy before taking any action. This analysis supports more risk-aware decision making at multiple levels.

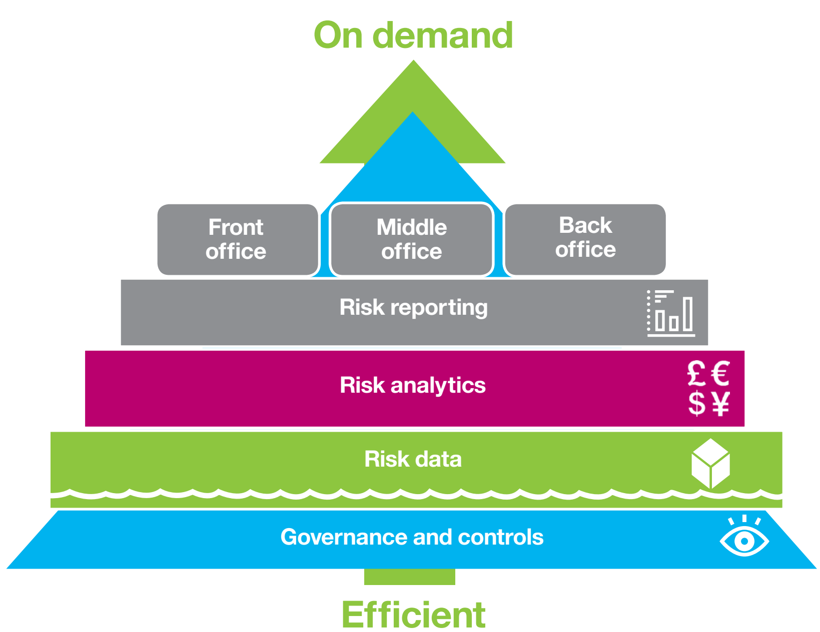

Figure 3: Big-data-driven gains in analytical efficiency can yield on-demand results to support key business decisions across the bank.

Plumbing an efficient big-data-enabled architecture

An enterprise-wide infrastructure comprises an enormous system of data and analytics, which breaks down into relatively few technologies for performing specific tasks. The continually evolving ecosystem of big data technologies is designed to be mixed and matched in ways that simplify and streamline the variety of processes involved in converting raw data into actionable insights.

Figure 5 provides an overview of the key components in a big data infrastructure that's capable of delivering analytical insights to multiple departments.

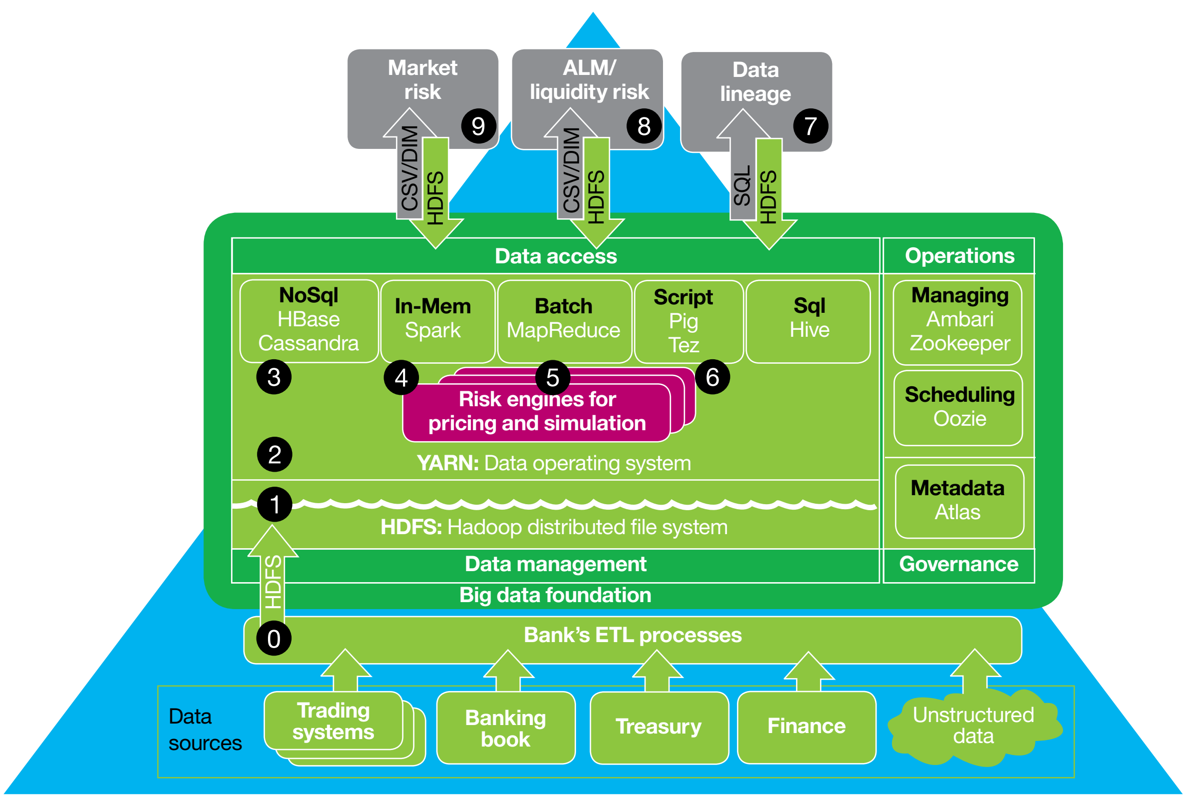

Figure 4: The key components and data flows of a big data architecture for risk analytics. See the next page for an explanation of these components.

Key components and data flows in figure 4

0. ETL processes pull through ETL sources (for example, market data and reference data) and internal sources (for example, transactions across the trading book and banking book).
1. A Hadoop Distributed File System (HDFS) stores all data required for analysis, including input data, output data and log files.
2. Yet Another Resource Negotiator (YARN) is the big data cluster management technology that co-ordinates and optimizes interactions between input data stored in the HDFS, the product mapping handled by Cassandra and Spark, the Hive aggregation engine, and the data partition with the risk engines for pricing and simulation.
3. The ETL source data stored on the HDFS, containing financial instrument definitions and other input data, are transformed and partitioned into risk-engine-ready inputs by the product mapping built on Cassandra and Spark, with Cassandra and Spark delivering higher performance over HBase.
4. Input data is partitioned into simulation-complete sessions for the distributed risk engines. A YARN-distributed cache collects and deploys static, market and scenario data, and simulation settings to multiple risk engine nodes.
5. Risk engine simulations are run as MapReduce jobs across all available nodes with work scheduled by YARN. YARN performs retries and achieves data affinity, optimizing data processing and ensuring risk engine simulations are run against the data available on the same node. The risk engines output simulation results into the HDFS.
6. Hive "schema-on-read" features interact with all risk engine inputs, outputs and log files stored in the HDFS without ever changing, duplicating or moving the data. Avoiding data transportation results in significant performance gains, and enables a full account of the data lineage and no loss of data granularity available for risk reporting. Hive provides a SQL interface to query data stored in the HDFS.
7. Third-party business intelligence (BI) tools query intermediate and final results stored on the HDFS and databases through the SQL interface provided by Hive. These queries can also support troubleshooting activities or an assortment of other applications appropriate for generic reporting tools.
8. Interfaces for asset liability management and liquidity risk enable the efficient modeling of alternative structures for the asset and liability mix. End-users are in full control of the setup parameters for the analytics, configuring the product templates, hierarchy structures, reporting time steps and modelling approaches used in the analysis.
9. A big data foundation can support risk analytics for multiple risk types, and integration with multiple user interfaces (UI) built for specific purposes. Multiday market risk analysis, for example, can read simulated and dimension data directly from the HDFS, supporting complete FRTB and market-risk reporting with drill-down and sandboxing capabilities.

Fortune favors the bold

Royal Bank of Scotland[2] and Citibank[3] have talked openly about their ongoing transformation towards a big data architecture, and many other banks are increasingly releasing job postings looking for big data experience and expertise.

For banks just getting started with big data, IBM recommends initiating pilot projects that can:

- Deliver a measureable short-term return on investment
- Follow best-practices already established through implementations at other banks
- Be expanded into full-scale, organization-wide big data programs that support more use cases

With leading banks already applying big data solutions towards ALM and liquidity risk management, IBM recommends banks prioritize this application for their next big data project. Successful ALM and liquidity risk projects help improve strategies for coordinated management of a bank's entire balance sheet across the disciplines of finance and risk management, and directly help improve the short-term profitability of the bank.

Over the long-term, investments in big data and other business analytics technology that provide end-users with more access and control can be reapplied towards supporting new client-driven solutions. These technologies enable banks to leapfrog the recent advances in market penetration made by emerging FinTech firms. FinTech firms have benefitted from taking the fast-mover advantage on new technologies, but are a long way from building a client base that rivals the scale of established banks. Banks that can most efficiently implement technologies to better understand and serve their clients will continue to dominate the market.

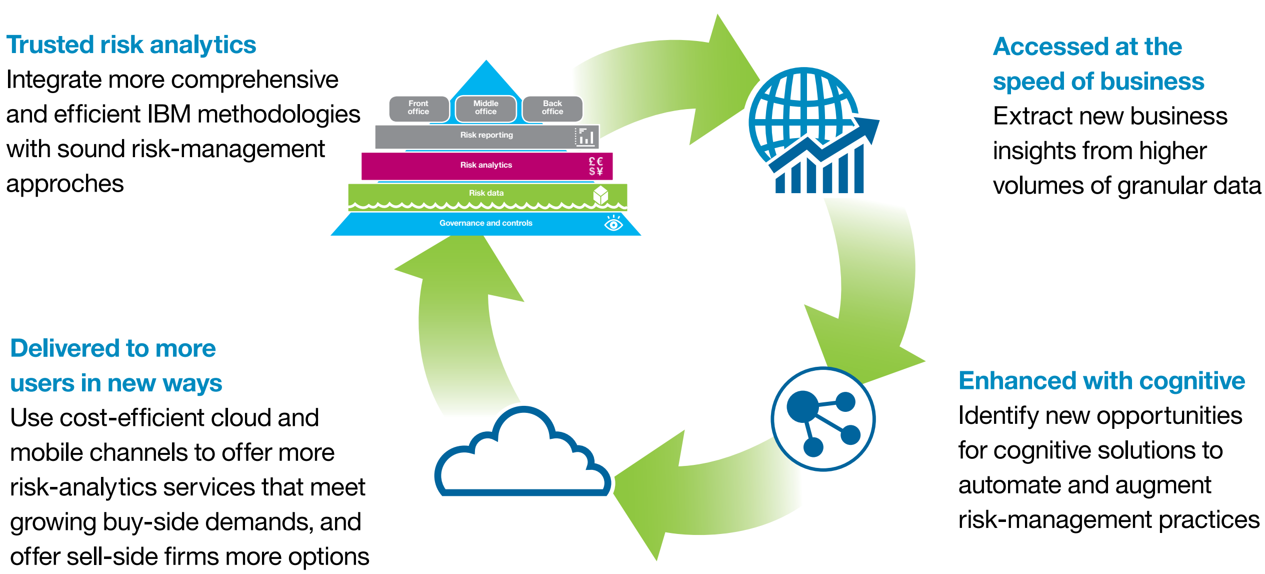

Figure 6: The positive feedback loop for firms that invest in delivering risk-based insights to more users in new ways

About the authors

Fraser Schad is Senior Strategist in Financial Risk Product Management for IBM. He works with IBM banking clients and technology partners, sharing insights on how regulatory changes are impacting risk management practices, and market demand for advances in software and technology. As an Offering Manager for big data applications in banking, Fraser draws on over ten years of experience in leadership roles across product management and marketing within the financial services sector. He can be reached via Twitter @fraserschad and Fraser.Schad@ca.ibm.com

Steven Good is the Global Offering Manager for the IBM Risk Analytics Asset Liability Management, Funds Transfer Pricing and Liquidity Risk solutions. He has over 10 years of experience in leading change management programs within the banking industry, aimed at improving balance sheet management operations. Prior to joining IBM, Steve was responsible for ALM and liquidity risk management at Lloyds Banking Group, where he developed a depth of practical experience and recognized how emerging regulations are creating opportunities for progressing best practices with new technologies. Steve hold degrees in Physics and Financial Mathematics. He can be reached at sgood@uk.ibm.com

Peter Timofejew is Lead Architect in the Research and Innovations team of IBM Watson Financial Services, directing the development of fresh approaches to risk analytics. He is committed to taking on the most challenging client projects and advancing best practices with cutting-edge approaches, such as big data technologies. His experience includes 14 years at TD Securities, designing systems for market and counterparty credit risk analytics, and credit derivatives trading. Peter was part of the team that converted the equity trading floor of the Toronto Stock Exchange to be fully electronic, authoring the messaging protocol STAMP for electronic exchange traded equities and derivatives. He can be reached at Peter.Timofejew@ca.ibm.com

Benjamin Hui, Chartered Financial Analyst (CFA), is the Offering Manager for banking solutions developed within the Research and Innovations team of IBM Watson® Financial Services. With over 20 years of experience in the banking industry, he leads collaborations across technical and functional teams to refine promising new innovations into production-ready offerings. Benjamin has global experience with blue chip names in international banking as both a practitioner and independent consultant—which includes the design and testing of financial models for corporate treasury hedging strategies. He can be reached at Benjamin.Hui@ca.ibm.com

IBM Corporation
Route 100
Somers, NY 10589

Produced in the United States of America
June 2017

1 http://www.informationweek.com/big-data/big-data-analytics/big-data-analytics-sales-will-reach-$187-billion-by-2019/d/d-id/1325631

2 http://www.forbes.com/sites/bernardmarr/2016/04/13/the-wonderful-big-data-strategy-at-royal-bank-of-scotland/ - 3cf0705b4ae1

3 https://www.forbes.com/sites/bernardmarr/2016/09/09/big-data-in-banking-how-citibank-delivers-real-business-benefits-with-their-data-first-approach/#3ff4b17d47e0

ASW12435-USEN-00