

WHITEPAPER

OPEN MODERN DATA ARCHITECTURE FOR FINANCIAL SERVICES RISK MANAGEMENT

A top-tier global bank's end-of-day risk analysis jobs didn't complete in time for the next start of trading day. To solve the problem, they are using Red Hat® JBoss® Data Grid running across 80+ commodity servers (and growing). Now, in addition to the analysis jobs completing in time with room for growth, the scalability and cost efficiency opens the door to new more comprehensive analytics. The Hortonworks Data Platform stores both the results and unstructured data in new types of analysis.

facebook.com/redhatinc

@redhatnews

linkedin.com/company/red-hatredhat.com

EXECUTIVE SUMMARY

Financial service firms are operating under significantly increased regulatory requirements such as Basel III, in the wake of the 2008 financial crisis. The risk management systems that each firm operates must respond not only new reporting requirements but also handle ever-growing amounts of data to perform more comprehensive analysis. Existing systems that aren't designed to scale for today's requirements can't finish reporting in time for start of trading. And, many of these systems are inflexible and expensive to operate.

Red Hat and Hortonworks offer a joint solution that provides an open, modern data architecture for financial services risk management. The solution is designed to scale-out using small commodity systems. An in-memory data grid, stores real-time data with very low-latency, while the Hortonworks Data Platform uses Apache Hadoop to efficiently process and store huge volumes of data that support advanced analytics capabilities. The key to scalability is that each additional system provides increased memory and storage, as well as more capacity for computation.

Flexibility is the other main advantage of the joint solution. The data is easily accessible to batch jobs, real-time processing, or interactive research. Standard interfaces are provided that allow you to use a variety of off-the-shelf business software. Hadoop is having such a large impact on enterprise big data that almost every business intelligence software vendor is adding support.

The strategic alliance between Hortonworks and Red Hat has a number of benefits, including engineering and integrated customer support. Both companies are aligned on the same open source business model, which focuses on making innovations from open communities enterprise ready through tested, certified software releases, and world-class support.

RISK MANAGEMENT SYSTEM CHALLENGES

As capital and liquidity reserve requirements have increased, the need to know exactly how much capital needs to be reserved, based on current exposures, is critical. Unnecessarily tying up excess capital can keep the firm from taking advantage of business and market opportunities.

The risk management systems that each firm operates must respond to new reporting requirements and also handle ever-growing amounts of data to perform more comprehensive analysis of credit, counter-party, and geopolitical risk. However, existing systems that aren't designed to meet today's requirements can't finish reporting in time for start of business or trading, which can lead to uninformed decisions. This is compounded by increasing requirements for intra-day reporting, as well as a shorter window for overnight batch processing required by global trading and electronic exchanges.

Other problems with the aging in-house solutions include:

- Data stored across many silos throughout the firm using multiple technologies that all require different methods for obtaining access. Instead of focusing on analysis and reporting, valuable time is wasted trying to figure out how to reliably obtain the necessary data.
- Proprietary solutions built using high performance computers or grid computing clusters that are inflexible and consume large portions of the available technology budget without meeting evolving challenges. Since these systems often don't use any standard interfaces, off-the-shelf tools can't be used or require custom development.
- Existing systems typically lacking the security and controls necessary to keep up with compliance and data security requirements.

Hortonworks and Red Hat offer products that can address a number of these challenges. The two companies recognize a solution based on both of their product offerings can address the broad range of challenges faced by financial services firms.

A MODERN DATA ARCHITECTURE FOR RISK MANAGEMENT

By combining Hortonworks and Red Hat solutions, you can create a modern data architecture for risk management. The foundation of the architecture is a centralized repository capable of aggregating all of the data necessary for analysis of all the exposures including positions, counter-parties, and liabilities at a firm-wide level. The joint products provide a repository that is capable of efficiently storing both frequently-updated, real-time data, as well as long-term historical data.

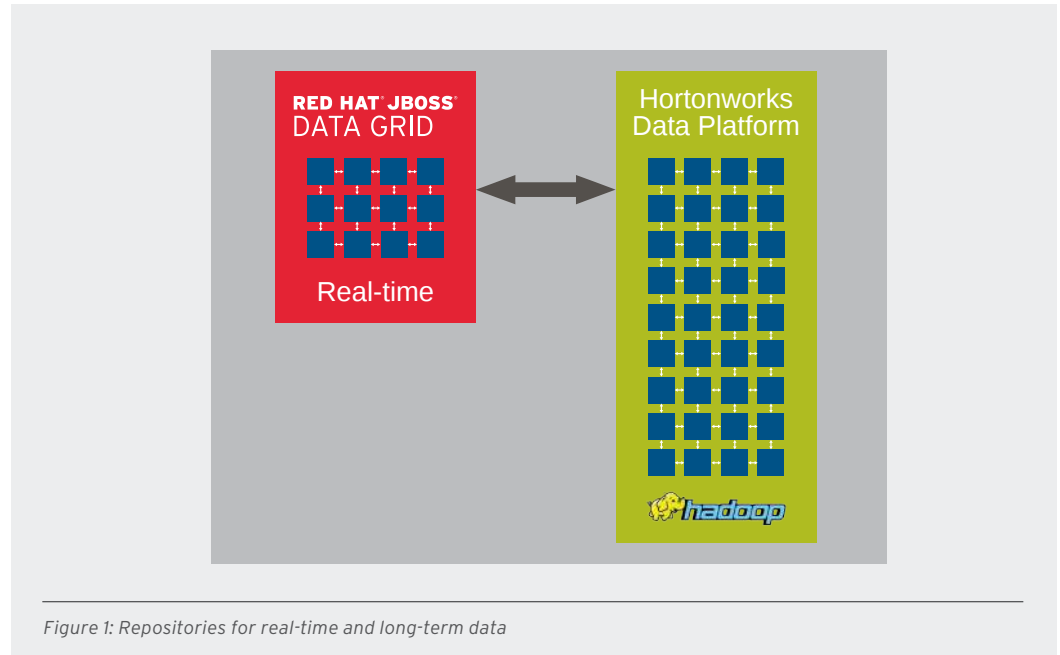
The architecture of the joint repository provides both performance and accessibility to data. This allows you to build the capability to execute liquidity risk analysis on an intra-day or multi-day basis, while providing long-term data retention. The data is easily accessible to batch oriented, interactive query, or real-time analysis and reporting tools.

The architecture is also capable of triggering business workflows and rule-based decision making without waiting for overnight batch processing to occur, which increases business efficiency.

Additionally, the architecture can easily scale not only to support ever-increasing volumes of data, but—more importantly—to enable increasingly complex and comprehensive risk analysis and reporting that are in response to new global business opportunities and their inherent exposures. Finally, the joint solution provides comprehensive data security controls to meet today's compliance requirements.

BOTH REAL-TIME INTRADAY AND LONG-TERM HISTORICAL DATA

Analyzing liquidity risk requires access to both real-time, frequently-updated pricing data as well as many sources of longer-term information like reference data, historical pricing, and results of past valuations. The two types of data have very different requirements for processing and storage. In the past, systems that attempted to address both types of requirements with a single data store generally either failed to meet performance requirements or were prohibitively expensive. To address this, the modern data architecture for risk management includes both a repository that is well suited to real-time data with frequent updates, as well as a repository that is specifically designed to keep accumulating large volumes of data as shown in figure 1.

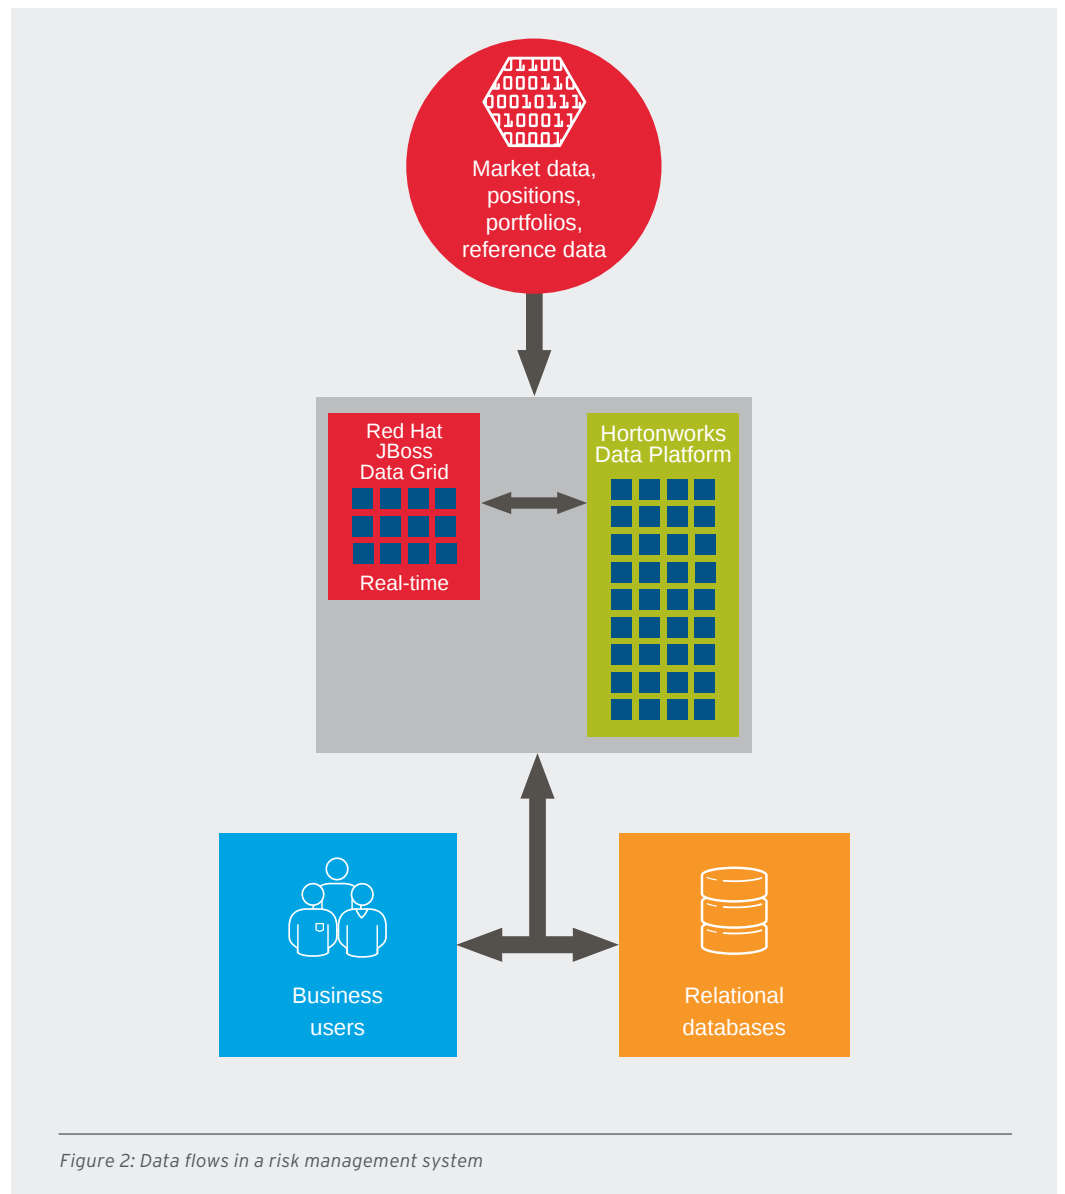


Frequently updated, real-time data is stored in the Red Hat JBoss Data Grid, which provides fast in-memory storage with low-latency access across an array of commodity systems. Longer-term data is stored in the Hortonworks Data Platform, which is designed to efficiently store massive amounts of data across an array of commodity systems using Apache Hadoop.

The Hortonworks Data Platform, along with Red Hat JBoss Middleware, provides a wealth of flexibility for accessing data that is stored in the two repositories. In addition, data that resides in other systems throughout the firm, can be made available through the same access capabilities either by loading it or federating it, which allows access without the need to copy or move all of the data. The modern data architecture for financial services makes it easy for batch or real-time applications to get the data needed. It also allows analysts to perform interactive ad-hoc research. Many standard data interfaces are available, which lets you use off-the-shelf software without having to develop custom bridges.

UNDERSTANDING THE DATA FLOWS

The data flows for a typical implementation of the open, modern data architecture for risk management are shown in figure 2.



1. Real-time data, like pricing feeds or trade execution data, flows directly into Red Hat JBoss Data Grid, which keeps the data in-memory for low-latency.
2. Historical and batch-oriented data is loaded into the Hortonworks Data Platform, which provides tools for transforming data and the capability for storing raw data.
3. Analytic calculations can be performed directly on Red Hat JBoss Data Grid. The results of the analysis can be routed to the Hortonworks Data Platform to make it easily available to other applications like batch-oriented analytics, or interactive query and visualization tools.
4. Results can also be stored in traditional relational database systems that can be accessed by existing applications in other parts of the firm.
5. Analysts can leverage business intelligence (BI) tools to perform interactive research on the data stored in the Hortonworks Data Platform. Other sources like relational databases can be seamlessly incorporated using Red Hat JBoss Data Virtualization.

ADVANTAGES

The modern data architecture for financial services provides a number of advantages over other approaches, including:

- Scalability and cost savings. The incremental cost for adding capacity is low because the solutions uses farms of commodity systems. Both the Red Hat JBoss Data Grid and Hortonworks Data Platform utilize arrays of commodity systems that can be scaled cost-effectively by adding more systems as needed. Due to economies of scale, these commodity systems have the best price-to-performance ratio. Each additional system increases the amount of disk, memory, and computational power, which increases the capacity for analysis as well as storage.
- New opportunities for analysis, modeling, and simulation. New efficiencies for storage and processing make it possible to process and store amounts of data that are cost prohibitive with traditional IT enterprise systems. The efficiencies allow analysis of longer-term trends, more scenarios can be run in less time, or additional factors can be added from unstructured data sources like consumer sentiment and social media. The value of some data only becomes apparent when high volumes can be analyzed. Exploiting these efficiencies can enable firms to not only better manage risk, but also find opportunities for competitive advantage.
- The ability to turn data into a strategic asset. Much of the data stored within financial services firms' numerous silos isn't fully exploited for business benefit. Hortonworks Data Platform and Red Hat JBoss Data Virtualization can integrate existing data sources—including data warehouses, SQL/NoSQL databases, enterprise applications, flat and XML files—into business-friendly, reusable virtual data models with unified views. The result allows analysts to concentrate on performing valuable research instead of trying to obtain and transform multiple data sources into something usable.
- Use of business intelligence tools and business rule management systems with your existing data and systems. The combined Hortonworks Data Platform and Red Hat JBoss Middleware solution provides standard interfaces to your data and systems so you can use off-the-shelf business software without extensive custom development. Because of the impact Apache Hadoop has on big data, almost all of the major business software vendors support Apache Hadoop, which is the core of the Hortonworks Data Platform. Hadoop YARN is the architectural center of enterprise Hadoop, allowing interactive SQL, real-time streaming, and batch processing across the data stored in the Hortonworks Data Platform.

100% Apache Hadoop

Hortonworks Data Platform (HDP) is powered by 100% open source Apache Hadoop. HDP provides all of the Apache Hadoop related projects necessary to integrate Hadoop alongside your existing systems as part of a modern data architecture. It ships with efficient data management and versatile data access methods, along with the capabilities enterprises require for widespread adoption: data governance, data integration, security, and operations.

- Streamlined processing. Business rule management software can react to real-time data. Workflows and automated decision making can be triggered by real-time events using the Red Hat JBoss Data Grid and Red Hat JBoss Business Rule Management System. This provides many opportunities to be more responsive and streamline processing without waiting for over-night batch processing.
- Meeting compliance and data security requirements for controlling and monitoring access to data. The Hortonworks Data Platform and Red Hat JBoss Middleware provide comprehensive data security capabilities that allow fine-grained control over data access.

The modern data architecture for financial services provides a number of advantages over other approaches, including: You can deploy the open, modern data architecture for risk management with confidence using proven, enterprise-grade infrastructure from Red Hat, like Red Hat Enterprise Linux® and Red Hat Storage. Flexible deployment options are available using Red Hat Enterprise Linux OpenStack® Platform, which allows you to enable your applications for the cloud.

WHY HORTONWORKS DATA PLATFORM

Apache Hadoop is an open source technology created from the experience of web-scale consumer companies like Yahoo, Facebook, and others, who were among the first to confront the need to store and process massive quantities of digital data. Key technology powerhouses like Microsoft, SAP, Teradata, Yahoo!, Facebook, Twitter, LinkedIn, and many others are continually contributing to enhance the open source platform, each bringing their unique capabilities and use cases. As a result, the innovation of enterprise Hadoop continues to outpace all proprietary efforts.

The Hortonworks Data Platform provides supported, enterprise-ready Apache Hadoop. The most current community innovation is incorporated into the Hortonworks Data Platform and is tested on the most mature Hadoop test suite running on thousands of nodes. The Hortonworks Data Platform is developed and supported by engineers with the deepest and broadest knowledge of Apache Hadoop. All of the Apache Hadoop related projects necessary to integrate Hadoop with your enterprise data systems are provided with the Hortonworks Data Platform.

Hortonworks was founded in 2011 by 24 engineers from the original Yahoo! Hadoop development and operations team, Hortonworks has amassed more Hadoop experience under one roof than any other organization.

WHY RED HAT JBOSS MIDDLEWARE

Red Hat JBoss Middleware integrates software from various open source communities, including the JBoss community, Apache Software Foundation, and Eclipse foundation into robust, fully tested, integrated platforms available via subscriptions that include support and long-term maintenance.

Red Hat JBoss Middleware helps you evolve your middleware infrastructure by providing the tools needed to rapidly build connected systems. Red Hat JBoss middleware transforms your costly, hard-to-manage application infrastructure into one that is truly dynamic across multiple environments.



WHITEPAPER OPEN MODERN DATA ARCHITECTURE FOR FINANCIAL SERVICES RISK MANAGEMENT

HORTONWORKS AND RED HAT TOGETHER

Hortonworks and Red Hat work together to provide integrated support to make it easy to resolve interoperability issues. Certification of the Hortonworks Data Platform with Red Hat Enterprise Linux ensures the solution works optimally. And integration of the Hortonworks Data Platform with Red Hat JBoss Middleware enables millions of existing JBoss developers to quickly build applications that take advantage of Apache Hadoop, giving you a greater choice of off the shelf software.

Both companies are aligned on the same open source business model, which focuses on innovation built in strong, open communities. The benefit to you is that you can take advantage of rapid innovation with software that is tested to ensure the highest levels of stability and reliability.

CONCLUSION

The lower costs achieved by using commodity systems, open-source software, and standard interfaces makes it economically feasible to store, process, and analyze more data than ever before. These efficiencies drive the modern data architecture for financial services risk management.

To find out more, contact your Red Hat or Hortonworks sales team to schedule a demonstration and discussion of how a modern data architecture for financial services can improve your risk management systems to better meet not only today's challenges, but also to open new opportunities in risk analysis and reporting.

ABOUT RED HAT

Red Hat is the world's leading provider of open source solutions, using a community-powered approach to provide reliable and high-performing cloud, virtualization, storage, Linux, and middleware technologies. Red Hat also offers award-winning support, training, and consulting services. Red Hat is an S&P company with more than 80 offices spanning the globe, empowering its customers' businesses.

ABOUT HORTONWORKS

Hortonworks develops, distributes and supports the only 100% open source Apache Hadoop data platform. Our team comprises the largest contingent of builders and architects within the Hadoop ecosystem who represent and lead the broader enterprise requirements within these communities. Hortonworks Data Platform deeply integrates with existing IT investments upon which enterprises can build and deploy Hadoop-based applications. Hortonworks has deep relationships with the key strategic data center partners that enable our customers to unlock the broadest opportunities from Hadoop. For more information, visit www.hortonworks.com.



facebook.com/redhatinc
[@redhatnews](https://twitter.com/redhatnews)
linkedin.com/company/red-hat

redhat.com
#INC0207199_V1_0115_KVM

NORTH AMERICA
1 888 REDHAT1

**EUROPE, MIDDLE EAST,
AND AFRICA**
00800 7334 2835
europa@redhat.com

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

LATIN AMERICA
+54 11 4329 7300
info-latam@redhat.com

Copyright © 2015 Red Hat, Inc. Red Hat, Red Hat Enterprise Linux, the Shadowman logo, and JBoss are trademarks of Red Hat, Inc., registered in the U.S. and other countries. Linux® is the registered trademark of Linus Torvalds in the U.S. and other countries. The OpenStack® Word Mark and OpenStack Logo are either registered trademarks / service marks or trademarks / service marks of the OpenStack Foundation, in the United States and other countries and are used with the OpenStack Foundation's permission. We are not affiliated with, endorsed or sponsored by the OpenStack Foundation or the OpenStack community.