MODERNIZE PRICING AND RISK ANALYTICS WITH RED HAT

TECHNOLOGY OVERVIEW

58%

of business and technology executives surveyed at financial services institutions cited shortage of resources as their most significant challenge in improving their risk analytics.[1]

80%

of business and technology executives say that use of open source software is increasing or significantly increasing in their organization.[2]

BUSINESS BENEFITS

- Maintain regulatory compliance by supporting long-term data retention

- Reduce IT capital expenditures (CapEx) and operating expenses (OpEx)

- Change business rules without disrupting the system

- Remain competitive by more accurately predicting pricing and risk profiles

INTRODUCTION

In the financial services industry, market conditions change constantly and regulatory requirements limit your ability to respond in a timely manner. New technologies and the fast pace of innovation offer opportunities for growth, but can also introduce uncertainty. In this environment, it is extremely difficult to quickly and accurately calculate and predict pricing and risk profiles.

Risk analysts need to be able to rapidly change business rules, easily access data of any type from any source, make complex calculations in near-real time, and store large volumes of data. But, as an IT professional, you cannot deliver that level of service if you are struggling with legacy risk calculation platforms that hinder your ability to respond to business requests.

Many legacy systems are cobbled together using multiple point solutions that run on expensive proprietary vendor platforms. These systems are unable to scale to handle massive volumes of data. They cannot be easily updated without disrupting their operation and are usually hosted on dedicated infrastructure that does not scale up or down with the workload.

Red Hat can help you modernize your legacy risk calculation platforms—and say "yes" to business requests. Red Hat's software stack uses the latest open source software technologies, as well as developments in agile and DevOps, to provide risk analytics that are flexible, fast, easy to change, and cost-effective.

RED HAT OPEN SYSTEM ARCHITECTURE FOR RISK ANALYTICS

Red Hat combines the power of industrial-grade versions of the latest, proven open source software for cloud, mobility, analytics, virtualization, and containers to provide a flexible software architecture for risk analytics. With a Red Hat® solution, you can give your business users the ability to do much more on their own, like:

- Quickly change business rules.

- Reduce the complexity of accessing data from disparate data sources.

- Rapidly and accurately perform complex calculations on vast amounts of data.

- Securely store large amounts of data for compliance.

The architecture helps you use agile and DevOps approaches to deliver continuous integration and innovation to the business. You can also exploit the power of container and cloud technologies to:

- Increase the delivery speed of pricing and risk analytics applications and features.

- Improve scalability, flexibility, and hardware utilization rates by moving workloads seamlessly across on-premise, private cloud, and public cloud environments.

---

1   *Red Hat / TechValidate Survey "IT Trends in the Financial Services Industry," March 2017.*

2   *TechValidate. https://www.techvalidate.com/tvid/E4A-C3D-FA8, Survey of 85 Red Hat customers at financial services institutions conducted in April 2017.*

• Maintain the stringent levels of security required to comply with regulatory guidelines.

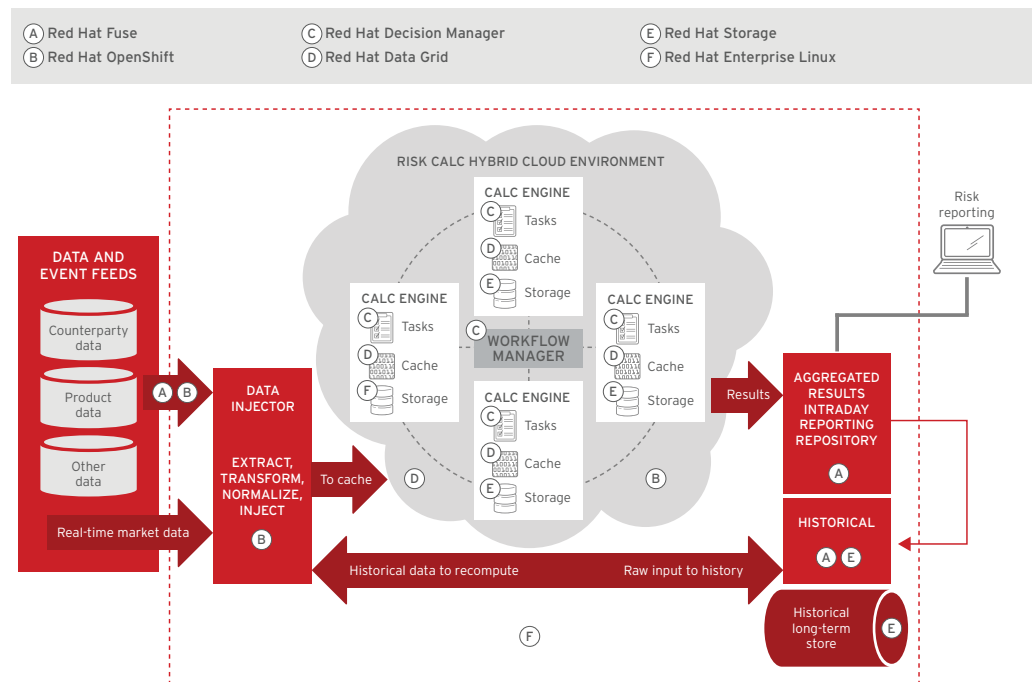• Reduce total cost of ownership (TCO) and avoid vendor lock-in by running on industry-standard hardware.

Figure 1. Pricing and risk analytics process flow

RED HAT APPROACH

The best way to understand the power of the Red Hat solution architecture for risk analytics is to examine the process flow—and where Red Hat components fit (see Figure 1). This document explains the functionality and value of the Red Hat products for each part of the process.

DATA AND EVENTS FEED

Business users need to access and analyze data from a variety of disparate sources. It is essential to be able to extract multiple types of structured, unstructured, static, and dynamic data from the data sources in both batch and real-time modes.

The solution uses Red Hat Fuse to provide this capability. It is a lightweight, pattern-based integration platform for noncritical path data loads. It lets you rapidly integrate the applications, data, and devices that populate and manage your data and event feeds. Its powerful set of integration and data transformation tools captures, aggregates, and routes the data to the data injector. Red Hat Fuse makes it easy for you to integrate new sources of data. Most importantly, it improves risk and pricing calculation accuracy by quickly providing business users with a greater breadth and depth of relevant data.

CASE STUDY

The compliance team at a large, multinational banking and financial services holding company embedded Red Hat JBoss® Data Grid in its global positions aggregation system. The system:

- Computes aggregate holdings across global businesses.

- Gives the firm the ability to comply with a multitude of geographical and juris-dictional regulatory requirements.

- Is implemented in embedded mode and the application is closely coupled, saving the high cost of network transfer for repetitive network calls to fetch remote data.

- Enables the compliance tech-nology team to build a state-of-the-art processing engine that uses synchronous and asynchronous processing modes and provides cache eviction features to remove and replace superseded data.

DATA INJECTOR

The role of the data injector is to reduce the complexity of ingesting disparate data feeds. It does this by normalizing multiple data formats, containerizing the data, and hydrating the distributed cache of the risk calculation hybrid cloud in its native format.

Red Hat Fuse is the data abstraction, transformation, and provisioning engine for this solution. It sits in front of multiple data sources and treats them as a single source, delivering the needed data, in the required form.

RISK CALCULATION HYBRID CLOUD

To process big data analytics in a timely, cost-effective manner, the risk analytics calculation engines execute calculations in a cluster, or clusters, of commodity hardware nodes. These nodes can be located on-premise or in a public cloud. Red Hat Decision Manager, Red Hat Data Grid, and Red Hat OpenShift® work together to create a high-performance hybrid cloud environment that is easy to deploy, easy to manage, and cost-effective to operate.

Red Hat Decision Manager helps you easily update market risk calculation platforms by moving the ownership of relevant sections of decision logic from developers to the business. This improves agility by allowing your business users to quickly react to changing market conditions on their own—without involving IT. Decision Manager executes the logic in real time, resulting in fewer failed transactions, improved flexibility, improved accuracy, and lower TCO.

Red Hat Data Grid is an in-memory data management solution that can be used as a distributed cache, NoSQL database, and event broker. In this solution, it powers the distributed caches in the calculation engines, as well as the aggregated results and reporting repository. Data Grid stores the required data close to, or in, the calculation engines—usually in memory. It also synchronizes copies of data across multiple servers to deliver continuous availability, information reliability, and linear scalability. All of these features eliminate data bottlenecks, freeing the calculation engines to execute at in-memory speed and deliver an incredibly fast user experience. In addition, Data Grid provides configurable ACID (atomicity, consistency, isolation, and durability) transaction support and integration with Apache Spark and Apache Cassandra cluster computing frameworks.

Data Grid lowers TCO because it operates in a distributed manner on commodity infrastructure and can scale up or down with ease. It also improves agility, throughput, and calculation speed. In addition, Data Grid strengthens compliance because data can be rapidly retrieved from in-memory stores and is more persistent due to the synchronized copies.

Red Hat OpenShift helps your developers quickly create new risk analytics applications. It manages container orchestration for development and operations across the data and events feeds, data injector, and risk calculation hybrid cloud. With it, you can provide developers with self-service pro-visioning and a suite of agile and DevOps tools that streamline development and let the develop-ers make changes without disrupting operations. When run on-premise, Red Hat OpenShift allows diverse workloads to be securely isolated on shared infrastructure, increasing utilization and lower-ing infrastructure TCO.

Running Spark[3] on Red Hat OpenShift mitigates the problem of managing large, multitenant Spark clusters by scheduling per-application Spark clusters with their dependent applications. This elimi-nates the need to design complicated scheduling policies to ensure quality of service (QOS) on a

---

3 *For more information see* **http://radanalytics.io/** *and* *https://github.com/radanalyticsio*

Red Hat risk analytics architecture lets your business take advantage of the power of open source technology to more accurately predict the risk profile of various combinations of assets while maintaining regulatory compliance.

large, multitenant cluster. Spark clusters that are managed within Red Hat OpenShift are isolated from clusters that are used in other projects, are capable of elastic scale-out and scale-down, and can be provisioned quickly.

Red Hat OpenShift also improves utilization rates and lowers TCO by giving you the ability to easily move workloads across on-premise and multiple public cloud service providers (CSPs) like Microsoft Azure, Amazon Web Services, and Google Cloud Platform. Using containers to package and deploy applications provides elastic scaling, higher density, and workload portability. This enables consistent, efficient, and highly scalable pricing and risk analysis across both public and private infrastructure—improving agility and reducing CapEx and OpEx.

AGGREGATED RESULTS AND REPORTING REPOSITORY

The role of the aggregated results and reporting repository is to cost-effectively aggregate long-term and intraday outputs from the risk calculation hybrid cloud. It also scales on industry-standard hardware without performance degradation.

Data Grid provides powerful data synchronization and persistence functionality combined with Red Hat Storage, it simplifies reporting and strengthens compliance.

Red Hat Storage provides storage for the risk calculation hybrid cloud, the aggregated results and reporting repository, and the historical and intraday store. Highly scalable and cost-effective, it offers software-defined file, block, and object storage that can span physical, virtual, container, and cloud environments, lowering storage TCO. It can quickly scale from hundreds of gigabytes to terabytes or petabytes, strengthening compliance by providing cost-effective storage for as much data as you need to store. Also, because it is software-defined, it reduces container storage OpEx by providing dynamic storage provisioning, persistence, elasticity, flexibility, and speed.

HISTORICAL AND INTRADAY STORE

To simplify reporting and strengthen compliance, the historical and intraday store houses the long-term data needed for reporting and provides a mechanism to recompute past risk calculations. In addition, it allows you to test updated algorithms before you run them on live data.

There are multiple options for implementing the historical and intraday store, depending on the type of data you need to store and your retention requirements. It is possible to use multiple implementation options simultaneously:

- **Distributed NoSQL.** This is one of the typical formats for storing large volumes of calculation results from each run. Data Grid can connect to a cache store, like Cassandra, using its distributed database architecture to provide a virtually unlimited, horizontally scalable, persistent store for calculation results.

- **Tick data.** A significant portion of financial data analytics is performed by sifting through large volumes of tick data. The higher the resolution of collected tick data, the larger the dataset and, therefore, the required storage capacity. The ability to look further back in time by storing more data with greater granularity helps financial institutions find more useful and actionable insights and meet regulatory requirements. Storage capacity for tick data can grow very quickly to terabytes or petabytes. Red Hat Storage offers a simple file storage platform that can scale easily and cost effectively by adding industry-standard hardware to the storage system.

- **Unstructured data.** Storage volumes often grow faster than the need for compute processing. In this case, it is optimal to separate storage capacity from the compute tier for greater cost savings and flexibility. Red Hat Storage lets large enterprises disaggregate big data compute analytics from big data storage. It gives you the flexibility to dynamically provision compute analytics separate from virtual machines and containers.

In addition, Data Grid and Red Hat Storage provide data management functionality that simplifies reporting and strengthens compliance.

The foundation of the entire solution is Red Hat Enterprise Linux®, the world's leading enterprise Linux platform. It delivers military-grade security, 99.999% uptime, and support for business-critical workloads. Red Hat Enterprise Linux is tried, tested, and trusted by more than 90% of Fortune Global 500 companies.[4]

RED HAT RISK ANALYTICS ARCHITECTURE BUSINESS BENEFITS

Market risk analysis is part of almost every decision in the financial sector. Timely and accurate risk analytics are essential to maintaining your competitive edge. Red Hat's risk analytics architecture gives you the power you need to navigate today's rapidly changing technology landscape and global financial markets.

GET MORE FROM INSIGHT FROM YOUR DATA

Red Hat's risk analytics architecture helps you turn big data into actionable information for faster, more accurate pricing, risk analysis, and forecasting by:

- Vastly simplifying user experience and reporting by masking the complexity of accessing and analyzing data from disparate sources.

- Eliminating data bottlenecks so that sophisticated pricing and risk algorithms can be executed on large amounts of data in near-real time.

- Providing greater flexibility of the data model, better integration with BI tools, and elastic scalability.

- Providing object, block, and file storage that can quickly scale from hundreds of gigabytes to petabytes of structured and unstructured data, to cost-effectively support your need to store any amount of any kind of data for analytics and compliance.

REDUCE I.T. INFRASTRUCTURE OPEX AND CAPEX

Red Hat's industrial-grade open source software stack is supported by Red Hat Engineering teams that work with one of the largest hardware and software interoperability certification ecosystems in the world. You can deploy Red Hat software across a wide range of industry-standard hardware and CSPs to scale up and down with ease—eliminating vendor lock-in and minimizing TCO. You can also seamlessly move workloads across internal and external infrastructure to reduce costs, implement disaster recovery, or take advantage of a CSP's innovation.

---

4 Red Hat client data and Fortune Global 500 list, 2016

Red Hat developer tools allow you to use the latest innovations in agile and DevOps approaches to achieve operational excellence. You can streamline and accelerate your development process and easily integrate systems and data, as well as automate manual, time-consuming processes.

QUICKLY CHANGE PRICING AND RISK CALCULATION ALGORITHMS

Agility is key to business success. With Red Hat's risk analytics architecture, you can give your business users the ability to modify business rules and immediately execute them in real time, without disrupting operation. Because the platform is open, the risk analytics architecture from Red Hat also gives you the ability to continuously take advantage of upcoming big data and analytics technologies (like machine learning and artificial intelligence) to provide your business users with new tools that let them quickly modify business rules to respond to changes in the marketplace.

CONCLUSION

In the financial services industry and today's global economy, market opportunities grow faster than IT budgets. Complex transactions, vast amounts of data, and the increasing demand for 24/7 response create challenges for even the most skilled IT teams. Open source technologies give you the low-latency, high-throughput infrastructure you need to help your business compete and thrive.

Standardizing on the risk analytics architecture from Red Hat lets your business more accurately predict risk profiles of various combinations of assets while maintaining regulatory compliance. The architecture helps you:

• Use emerging technologies and approaches, like agile and DevOps, to reduce OpEx and CapEx, as well as deliver continuous integration and innovation to the business.

• Vastly simplify the complexity of accessing and analyzing huge volumes of data from disparate sources.

• Give risk analysts the power to quickly define or modify sophisticated business rules on their own without disrupting their operations.

• Execute analytics in near-real time to keep up with changing market conditions and compliance regulations.

A trusted solution from Red Hat is affordable and reliable, providing the security, technology, and performance advantages financial institutions need to compete. Red Hat offers a comprehensive, curated, tested, supported portfolio of safe and stable open source infrastructure and application development solutions that integrate with a wide range of current and emerging technologies, so you can keep up with technology and stay ahead of your competition.

Learn more at redhat.com/finserv
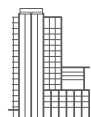
NORTH AMERICA
1 888 REDHAT1

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

ASIA PACIFIC
+65 6490 4200
apac@redhat.com

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