



Risk in the Cloud: I want it all ... and I want it now!

Why now and how to get it right
the first time

Cloud-based risk processing has become increasingly attractive for banks. The immediate appeal is not having to constantly add “on-premises” compute and storage, instead using on-demand infrastructure. Beyond “going faster for less,” the strategic advantage of knowing your exposure before anyone else can provide a significant competitive advantage to your traders and risk managers.



Abstract

Cloud-based risk processing has become increasingly attractive for banks. The immediate appeal is not having to constantly add on-premises compute and storage, instead using on-demand infrastructure to achieve better business outcomes. Beyond “going faster for less,” moving risk to the cloud also gives you an opportunity re-architect risk, pricing and profit-and-loss solutions. In addition, applying analytics and Artificial Intelligence (AI) to risk data sets in the cloud can provide a competitive advantage.

This white paper, intended for risk managers and risk IT professionals at investment banks, describes the advantages of risk processing in the cloud and outlines a five-step journey.

What FRTB Means to Risk Computation

Europe’s Fundamental Review of the Trading Book (FRTB), the new regulatory framework for market risk capital, will go live in 2022. FRTB is among the most significant changes since the introduction of model-based risk capital regulations. It increases the variety and volume of data included in computations, compelling banks to significantly increase storage and computation resources. To understand the impact, consider that just one change, the new expected shortfall-based capital measure, will require approximately 50 times more computational power.

How On-Premises Grids Fall Short for Risk Computation

Investment banks have invested in sophisticated risk computation systems, often with thousands of cores and massive storage. Hardware demands continue to grow as regulatory bodies require more risk measures on every trade. Already, some businesses need almost a full day to compute their overall exposure, particularly for credit valuation adjustments and other complex calculations. Meeting increasing computational and data throughput requirements with an on-premises data center is not cost-effective—and sometimes not even feasible.

Advantages of Risk in the Cloud

The benefits of moving risk computation to the cloud include:

- **Cost savings.** This is the most obvious benefit. In the cloud, you pay only for the resources you actually use. It’s not uncommon for servers to be idle 70-80% of the time, so usage-based billing can significantly reduce costs.
- **Competitive advantage of solving the problems faster.** The more servers you assign to the task, the faster you can understand your risk exposure. With a cloud deployment, you can scale from, say, 200 to 2000 cores (and with a little effort, even 200,000 cores) with a few clicks. The strategic advantage of knowing your exposure before anyone else—or pricing your trades with a full re-valuation—is among the most compelling reasons for banks to look at the cloud.
- **A chance to re-architect risk and pricing environments for greater efficiency.** Most enterprise and front-office risk and pricing systems use commercial grid management and

distribution middleware to orchestrate and dispatch parallel workloads, collect results, and monitor and report on usage and billing. Cloud-based risk processing enables you to re-think grid distribution approaches by taking advantage of the cloud’s auto-scaling and open-source middleware for parallel computing and high-performance computing (HPC). Commercial grid middleware is now available in cloud-aware variants that integrate with Apache Spark (the open-source cluster computing framework), big data and other technologies.

- **Innovation.** The combination of the cloud, big data analytics, AI, open source and HPC is ushering in a new generation of innovative risk infrastructure and solutions. So it’s not just about cheaper and faster, it’s also about doing things in ways previously unimaginable – just imagine Quantum computers and multi-dimensional directed acyclic graphs (DAG) in the cloud.
- **Early access to faster computation technology.** For banks and other enterprises, keeping up with the frantic pace of evolution in non-traditional compute technologies and products is onerous and expensive. Cloud providers, in contrast, invest continually in R&D for innovations ranging from graphics processing units to field programmable gate arrays (FPGA) to Quantum cores, including them in their normal offerings. Quantitative research and risk groups are already using non-traditional compute technologies—such as General Purpose Graphical Processing Units (GP-GPUs) and FPGAs to dramatically accelerate critical computations and support highly parallel and vectorized calculations.

Risk in the Cloud: Why Now?

Risk departments have not led the charge to the cloud, slowed in part by security concerns. In particular, banks traditionally have been wary of moving sensitive trading data, customer transaction and pricing and risk models to external facilities or clouds. Today, external clouds have robust security features equaling or exceeding enterprise levels. At the end of the day, running a data center is not and should not be a bank's core competency. Additionally, today's networks are robust enough to provide the last-mile connectivity to the bank that's needed to transfer big data loads. As a result, financial institutions are beginning to move customer data and proprietary models to the cloud.

Getting the Cloud Right: Five Incremental Steps

Moving risk to the cloud requires a mix of expertise: HPC, multidimensional data analytics, big data and AI and deep learning. The approach we recommend is to draw from the best existing solutions and then extend them with advanced engineering, raising the bar in buy-side risk and pricing computation.

1. Like for Like: Grid in the Cloud

This is a not-so-obvious first step for organizations sprinting to experience the benefits of the cloud. It's particularly appealing for banks that are relatively new to the cloud and grappling with some of the hygiene factors associated with security and governance. Moving the grid to cloud gives you a first-hand feel for cloud challenges and an understanding of the integration points between your on-premises and cloud environments.

This step involves lifting and shifting the grid as-is, continuing to use existing scheduling solutions such as Symphony, Data Synapse and the like. Depending on which cloud you use, the orchestration technologies may

already be cloud-ready. While this exercise is simple, it has a clear return on investment, especially if your grid is under-utilized.

2. Grid in the Cloud with Native Orchestration

Instead of using the expensive grid software you used when the grid was on-premises, use the cloud provider's scheduling, orchestration and auto-scaling capabilities. This step takes you closer to re-architecting your risk grid in the cloud.

3. Look Beyond the Grid to the Full Lifecycle

Risk platforms are seldom isolated. The overall solution spans multiple source systems, transaction systems, grids and analytical architectures. Trades and market data are usually pulled into risk workloads, offloaded to the grid, computed and then aggregated in a warehouse for analytical insights. Cloud migration gives you an opportunity to look beyond just the grid. Candidates for reengineering include:

- Market data management and analytics. Every risk-computation execution requires a market data snapshot. If the market data doesn't change all the time, housing it in the cloud will significantly reduce traffic volume between the data center and the cloud.
- Warehousing and analytics. Moving this to the cloud is possible only if users can access cloud-based output from their interface. This may be difficult to achieve quickly if trading desktops need to be updated.

4. Re-architect and Re-engineer from the Ground Up

Steps 1-3 are incremental changes; this one is radical. It's about taking a completely new look at the way you compute, store and consume risk data.

Most cloud providers offer technologies to build DAG risk-computation architectures—à la Goldman Sachs' SecDB and JP Morgan's Athena—that are scalable and flexible. The DAG architecture supports real-time computation of risk with co-located trade and market data (Figure 1). It also allows the "in-situ" aggregation of data. While

Lessons from BCBS239

BCBS239, the regulatory framework for banking data quality, lays down eleven principles for ensuring robust infrastructure, risk analysis and reporting. Discussions about BCBS239 often focus on technical and data management issues to the exclusion of the stated purpose, which is for banks to realize "gains in efficiency, reduced probability of losses and enhanced strategic decision-making, and ultimately increased profitability." This objective is very relevant in the context of cloud-based risk processing and FRTB.

Many firms have invested heavily in data management and governance tools and processes to support the BCBS239 principles, especially with respect to Completeness, Accuracy and Timeliness, the so-called CAT standards. Any cloud migration must be BCBS239-compliant. It should conform to the standards, solutions and tools supporting BCBS239 in the institution. This process will likely also entail extension of the BCBS239 tools and processes, taking into account cloud security, bandwidth and usage.

Today's pipeline architectures focus on traditional data warehouse designs, DAG architectures support object-oriented designs that significantly reduce computation and data movement over the network.

Organizations that have already built on-premises DAG-based platforms can go directly to this step and start leveraging the power of the cloud to turbocharge their DAG engines. Organizations that still have pipeline-based architectures need to make a judgment call about whether to jump straight to this step or first take steps 1-3. The decision depends on where you are today, your internal capabilities and how crucial the cloud is to your overall strategy.




Figure 1 Reference Architecture: Risk in the Cloud Using Directed Acyclic Graphs (DAG)

Challenges to Consider as Part of Re-engineering

- Limited revaluation capabilities of existing Value at Risk (VaR) systems, which rely on multi-stage pipeline architectures
- Inconsistent and unintegrated trade, market data and reference data sources
- Inefficient and flawed dependency tracking for risk factors, pricing models, instruments, securities and books by date and contexts
- Compliance with internal and external regulations, including FRTB, as requirements increasing computational demand

5. Leverage AI in the Cloud

This is the icing on the cake. (The cake is a mature risk-computation platform in the cloud.) Popular applications of AI for risk management include:

- Identifying and remediating model and risk defects (Figure 2)
- Improving capital efficiency
 - Optimizing FRTB trading desk structuring to achieve efficient capital
 - Reducing non-modellable risk factor capital impact
- Choosing the appropriate internal

risk models to replace existing models using the standard approach to computation

- Business planning and capital strategy

Applying AI in the cloud gives you the advantage of accelerated computation, scale on demand and ability to use the cloud service provider's high-performance AI and machine learning tools.

The Hybrid-Cloud Route

Let's Be Real: Not Everything Will Move to the Cloud—Yet

As mentioned earlier, risk is not an isolated function. It touches numerous stakeholders, systems and process throughout the bank, and not all of these systems are ready to sit in the cloud. Reasons to keep some systems in the cloud include data privacy, jurisdiction concerns, economics and latency requirements. Therefore, design the cloud-based risk architecture with the hybrid approach in mind: some elements in the cloud, some on-premises. Considerations in hybrid-cloud design include operational costs, data latency and placing the right components on the right side of the firewall.

Conclusion

Moving enterprise risk processing to the cloud is a major milestone in the

history of risk. Taking full advantage of the opportunity requires a fresh look at the challenges—and also openness to adopting innovations like AI and new compute technologies.

As the saying goes, good hardware (or even the cloud) cannot overcome a badly designed platform. Make sure to work with professionals who understand banking and risk as well as the ins and outs of enterprise architecture, the hand-offs with the cloud and what an optimal solution looks like in its entirety.

Hitachi Consulting is actively working with some of the world's largest banks to help them embrace the cloud in a pragmatic way. We deliver incremental visibility and value at every point in their journey. Our engagements range from back-office high-volume regulatory reporting data lakes to front-office, near-real-time trading risk analytics solution with sub-second latency. Hitachi Consulting also provides significant capabilities in enterprise architecture and cloud technology and is a strategic partner for Microsoft, Oracle and Google Cloud solutions.

For more information please contact learnmore@hitachiconsulting.com

Suranjan Som

Vice President, Financial Services, Hitachi Consulting

Suranjan has been working in the risk and regulation space with some of the largest investment banks in world in a consulting capacity and hands-on experience of architecting some of these solutions himself. He has a broad appreciation of the financial services industry and leads the industry vertical at Hitachi Consulting.

John Barclay

Independent Financial Service Industry SME

John is a very well known figure in the risk space having worked at Goldman Sachs, JP Morgan, Credit Suisse and Deutsche Bank. He is a proponent of DAG based risk computation and has assisted in Hitachi's Cloud based risk computation strategy.