

ADAPTIV FRTB RISK AGGREGATION

Empowering the Financial World
FISGLOBAL.COM

FIS

Adaptiv FRTB Risk Aggregation: A look from the IT Perspective

The new rules under the FRTB Internal Models Approach (IMA) bring with them an explosion of data that will challenge traditional IT architectures and technology. These new rules are significantly different from the current capital rules and prompt major changes to infrastructure and business models.

Where simulation-based measures are employed, risk managers and regulators are demanding richer drilldown and aggregation, as well as attribution, explanation, and an auditable correction process.

Additionally, the non-simulation Standardized Approach (SA) measures require clever and intricate aggregation. End-users are demanding full transparency into results, full audit and reproducibility, dynamic what-if risk analysis and real-time correction updates for current and past as-of-dates. The trend is clear – a complete picture of risk must include powerful aggregation analytics.

Separation of price valuation and aggregation

The burden of introducing and vetting new models is driving many banks to choose an approach to FRTB where they re-use P&L vectors and sensitivities generated in existing pricing or risk systems to feed into a risk aggregation system that calculates the FRTB capital numbers. Many banks' risk infrastructures now separate the process of simulating and producing sensitivities and P&L vectors from the aggregation and post aggregation analytics of VaR and FRTB calculations. This approach can also help pass the FRTB P&L attribution test when the engine used to generate official P&L is used to generate the vectors.



Aggregation Requirements

This approach of making the market risk system act as an aggregator of external valuation services avoids some of the challenges of FRTB but creates new challenges.

- ✔ Firstly, the aggregation solution must be able to calculate the FRTB metrics quickly, using fast engines for post aggregation analytics and provide all the drill through and measures a risk user would expect using commoditized BI tools.
- ✔ The cube schema's dimensionality needs to cover the large number of measures needed for FRTB with ability to see them holistically within the same report to actively perform result examination and what-if on multiple measures at once. Different measures will have different dimensionality; it is critical the architecture manages these together in one cube schema.
- ✔ The data velocity performance must be highly responsive to end-users in processing and loading into memory, and measure navigation and drill through. It is increasingly important to find data errors on today's and past date results. The architecture must allow for fast incremental intraday corrections of erroneous data in which users can isolate the initial loads from corrections for analysis.
- ✔ The solution must be scalable and cost-justified, as there will be billions of P&Ls (scenario PV changes) being passed to the solution. The faster the aggregation can take place, the larger the time window allowed for analysis and reporting.
- ✔ The solution must be accessible to a large number of concurrent end-users at all times, handling and streaming data updates in real time as users query and navigate.
- ✔ The FRTB solution needs to deliver value through performing interactive what-if investigations, analyzing and resolving data quality issues, active risk management via advanced what-if.
- ✔ The solution must be easy to deploy within the existing bank's risk infrastructure.

- ✔ The solution must be easy to use – users will need to access information fast and flexibly for analysis and management reporting.
- ✔ The solution must be extendible, adaptable and open to allow for changes as regulation emerges and changes over time.

In-memory technology

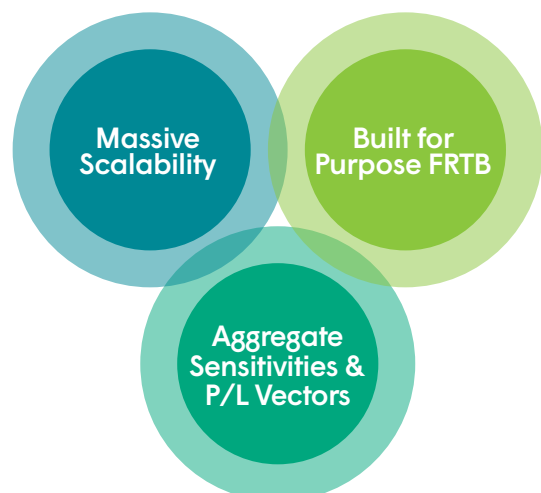
Adaptiv FRTB Risk Aggregation provides Business Intelligence (BI) and in-memory technology built for purpose to manage the high data volume (billions of P&L scenarios) and risk aggregation analytics of FRTB calculations. The management of data allows for real-time incremental correction updates and incremental trading activity as end-users are connected simultaneously analyzing results.

The scalable in-memory cube technology used in the solution addresses the unique challenges of risk calculations. Its object-oriented design gives it massive performance wins as the objects are vectors as opposed to scalars. The benefits are numerous:

- Optimizes how non-additive measures such as VaR, FRTB IMA Expected Shortfall, are structured in memory
- Allows for the aggregation objects for example, to be the set of SBA sensitivities and curvature data for the Standardized Approach, the IMA P&L vectors for IMA-ES calculation.
- Allows for highly customized post aggregation logic, i.e., optimally support the IMA Default Risk Charge (DRC) object to only import the sparse non-zero scenario values with their associated scenario sequence number, reducing the DRC vector (100K scenario values) by over 90% on average.

Aggregation of the required FRTB data will reach new high volumes while at the same time end-users demand dis-aggregation to position level measures.

The burden of pricing model validation drives banks to reuse pricing services already vetted against their P&L.



Not the traditional OLAP

In traditional BI, solutions are limited to aggregating scalar values. When aggregation occurs by vector, the indexed value is the entire vector and not the individual scenario P&L, resulting in a massive increase in performance velocity in load time and savings in memory usage. For example, in traditional BI, each scenario P&L had to be indexed by trade and scenario number. Here, it only needs to be indexed once for the entire vector object. The overhead is same for any vector size. The only change is the time to populate the vector in memory. The built-for-risk schema supports defining any object and any logic on the objects for a true object oriented OLAP. For example, we deliver support for IMA DRC by allowing hydration of only the sparse non-zero scenario values of a 100,000 scenario P&L vector.

The result is that end-users get a rich drilldown and dis-aggregation of risk metrics, with extremely fast query and aggregation-times based on a schema optimal for risk navigation. In addition, this is -orders of magnitude faster than traditional relational or OLAP database technologies.

Traditional OLAP technology was intended for solutions with infrequent data loads. Unfortunately, the risk production process is complex and initial daily loads from feed applications are never right the first time. The solution must allow for frequent intraday data updates due to missing data feeds, and the myriad of intraday corrections requiring recalculations at the source. The in-memory technology has optimized performance times to ensure results can be streamed in real-time, while ensuring users have a consistent view over all measures.

In-memory analytics, should optimally deliver in one in-memory schema a single cube for all measures. This allows what-if and views of all measures to be available for side-by-side analysis and comparison. For example, having SA measures and IMA measures in one cube for consistent updates.

Scalability

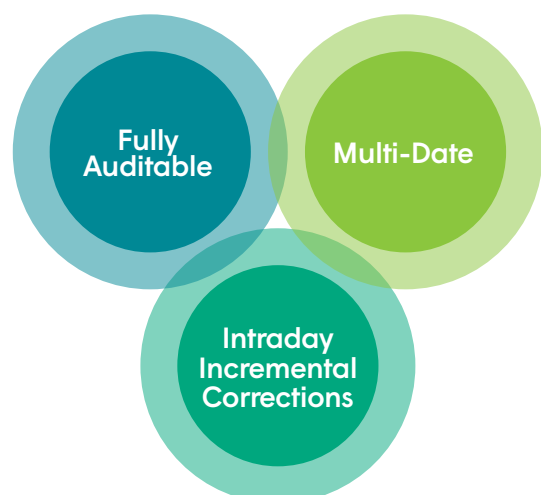
The in-memory technology is vectorized and multi-threaded. The solution can scale, allowing increases in the number of cores and RAM as required. Increase cores as the number of users increase, reduces querying speed and aggregation performance.

The solution can scale up and scale down to fit the client's business requirements and deployment. For very large clients, the solution supports multi-server cube distribution based on, providing an experience as if the user has connected to one uniform cube.

In-memory Hydration

Hydration speeds into memory are fast since instream data is first written to a detail pre-aggregation data store. This allows for:

- Multiple date hydration and multiple offline cube hydration from the single pre-aggregation data store
- Local user hydration of cubes per users for isolated querying and investigation
- Allows for point-in-time hydration for reporting and downstream report extracts
- Support for correcting today and past as-of-dates
- Full audit and reproducibility as hydration does not require reprocessing from the provided input files. Hydration occurs from the pre-aggregation data store, which reflects the imported validated and enriched data
- Simple in-memory cube recovery and back up. In the event of failure, re-hydration occurs from the detail pre-aggregation data store ensuring the initial imported data plus all the corrections applied to the initial data is intact and sequenced.



Extensibility

The solution's in memory analytics are extensible, as new calculations are introduced, aggregation plug-in calculations are added or changed to reflect custom aggregation logic. The solution comes with built out- of-the-box functionality for all FRTB calculations under SA (SBA, DRC, Residual Risk), IMA (ES, NMRF, DRC, P&L attribution, back testing), Aggregate Capital Charge and SA-CVA (once final rules are published) with continued functional solution support as regulations are modified.


Clients can also add their own measures, simple measures are supported using the schema, highly complex measures can be added and work in an identical way to the existing measures.

What-if Analysis

What-if analysis is supported in the in-memory design as an incremental overlay session. Once the user dismisses the session it is gone. The user builds up what-if instructions which can be saved and then replayed again or shared with other users. The what-if instructions include:

- Ability to scale any measure, or group of measures at trade or any aggregation node
- Remove a trade or an aggregation point (e.g., remove a desk)
- Flip a desk from SA to IMA eligibility or IMA to SA
- Move a trade or a set of trades from one aggregation point to another
- Import a new trade portfolio not currently in the production set from an external folder and insert into the portfolio structure

When what-if instructions are applied, recalculation occurs immediately; the measures in view are dynamically updated; the user sees the before and after results of the what-if.



WHAT-IF OVERLAYS

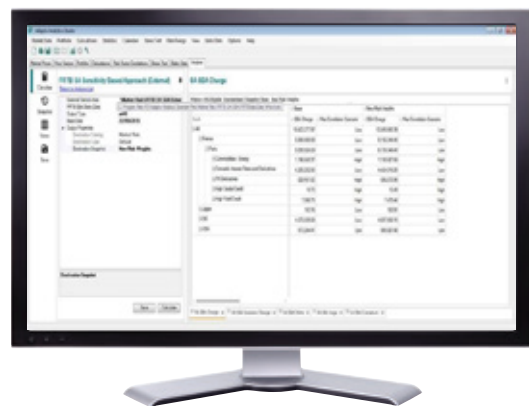
Sandbox Analysis

The solution comes with the ability to recreate production results from the feed inputs. The desktop version loads the provided input files and static data into a user stateless in-memory sandbox for independent analysis.

The desktop version hydrates the same in-memory schema with the identical in-memory analytics as in production. The only difference is that the offline in-memory schema is stateless (temporary) and memory is released once the desktop application is closed.

The sandbox analysis allows for non-incremental type what-if analysis, where a user can change an input or a calculation parameter setting that requires a full recalculation of the FRTB analysis. For example, when changing risk weights or risk buckets for the SA SBA sensitivity aggregation, the user can see the new snapshot against the initial and the resulting differences.

The power of the in-memory technology allows for the local user to have the exact same aggregation deployment as the enterprise deployment, limited only by the amount of memory on his or her desktop.



QUANTITATIVE ANALYST SANDBOX



End-user Experience

There are multiple methods for end-user connectivity. The in-memory connectivity uses industry standard BI tools based on XMLA, ODBO and MDX.

The solution provides an HTML5 based, built-for-purpose Web UI that allows users to see the status of production processes and production alerts, and provides methods to define and link analysis views for navigation of results.

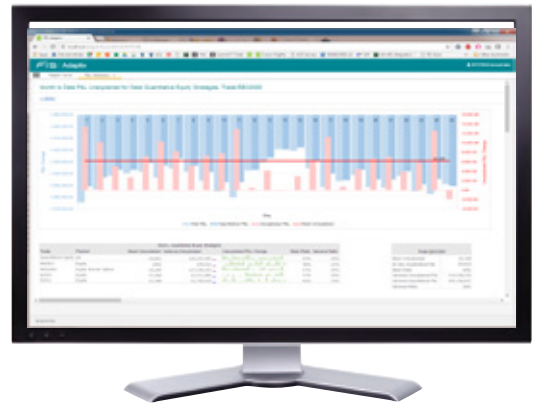
Result views can be linked so that as a user drills through and clicks on a result the linked view is filtered to the associated trade set.

All views are user-designed. The cube schema is accessible in design mode to end-users. The user can choose to utilize the Web UI that is provided with the solution (with powerful what-if analysis in place, and default views for each FRTB analysis).


FIS Risk Reporting allows users to create dashboards in Excel and publish them to the web. PDF Reports can also be produced and distributed on a schedule. The system also natively integrates with Excel pivot tables so users can access and interact with golden source results directly from Excel. Or they can use other XMLA compliant BI tools.



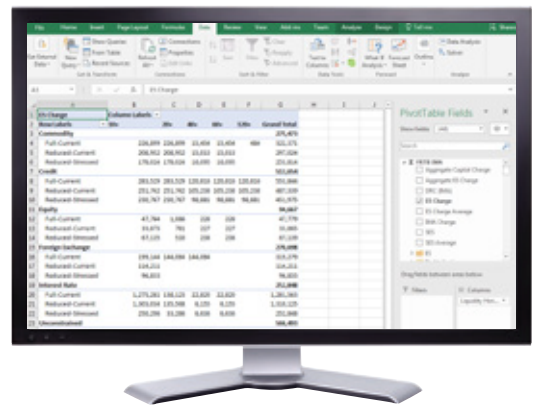
DASHBOARDS



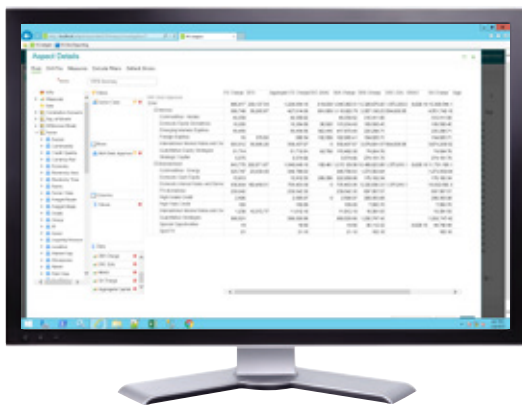
INTERACTIVE REPORTING



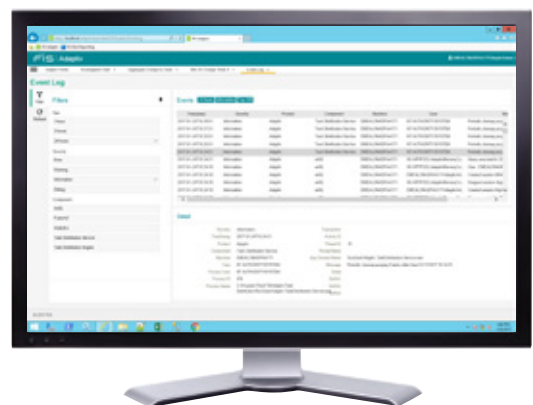
RISK INVESTIGATIONS



EXCEL CONNECTIVITY



VIEW DESIGNER



EVENT PROCESSING LOGS

Performance

The solution emphasizes performance in the following three areas:

- The time it takes to import and optionally join data for enrichment tagging and then store optimally in the pre-aggregation data store for hydration
- The speed of hydration of data from the pre-aggregation data store into the in-memory structures, where data is in memory at the lowest level
- Query time performance

For example, in just a few minutes, a real-world enterprise portfolio of 250,000 position lines with each position producing on average 16 P&L vectors of 250 P&L values (i.e., 1 billion facts) for its FRTB scenario set, factor class and liquidity horizon combination is imported and enriched.

Hydration of the enterprise portfolio into memory takes 1 minute. The organizational portfolio has roughly 15 levels with about 500 leaf portfolios and 50 trade attributes.


On this size enterprise portfolio, most queries are completed with a sub-second response time. An example of this is a top-of-house query of the FRTB ES calculation measure by desk.

Multiple date hydration into memory is just a multiplier of the above. For example, five days of hydration of this size portfolio will take approximately five minutes.

The hardware kit is representative of 16 Virtualized cores on 4 X Intel E7-4970 @ 2.4 GHz, 64GB RAM and 4 x 200GB SATA SSD HDD in Raid 0 at an approximate kit cost of 20,000USD.

Conclusion

Adaptiv FRTB Risk Aggregation combines a high-end technology solution with risk-specific intelligence. It is a built-for-purpose, out-of-box solution, enabled by optimized in-memory technology that allows FIS to deliver rich risk aggregation solutions that target requirements such as FRTB. In this ever-changing regulatory environment, there is a demand for auditable and transparent data to aggregate and disaggregate as needed. This requires in-memory scalable solutions optimizing memory consumption, lowering hardware costs and maximizing user experience. The FRTB Risk solution is available as a standalone aggregation engine or embedded in the Adaptiv Risk solution.



About FIS

FIS is a global leader in financial services technology, with a focus on retail and institutional banking, payments, asset and wealth management, risk and compliance, consulting and outsourcing solutions. Through the depth and breadth of our solutions portfolio, global capabilities and domain expertise, FIS serves more than 20,000 clients in over 130 countries. Headquartered in Jacksonville, Florida, FIS employs more than 55,000 people worldwide and holds leadership positions in payment processing, financial software and banking solutions. Providing software, services and outsourcing of the technology that empowers the financial world, FIS is a Fortune 500 company and is a member of Standard & Poor's 500® Index. For more information about FIS, visit www.fisglobal.com



www.fisglobal.com



twitter.com/fisglobal



getinfo@fisglobal.com



linkedin.com/company/fisglobal