White
Paper

IBM Algo Workspace Analyzer

Can your risk engine keep pace?

Watson Financial Services

IBM

Contents

Winner of the RegTech Awards 2018, A-Team Group

"Best Vendor Solution to Address an FRTB Requirement"

High-level summary: The risk management landscape is shifting

Rapidly emerging business requirements are shifting the landscape for risk management systems. Financial firms strive to both improve business oversight and comply with increasing regulatory demands. At the same time, data generated by financial institutions is growing in volume and complexity. As a result, traditional risk solutions are becoming obsolete. Firms are re-thinking how to apply new technology to meet these mounting challenges.

What used to be called risk reporting has morphed into a broader and deeper set of requirements. Static reports were once the norm, but today business users need to perform dynamic analyses that span the enterprise. Transparency is required from high-level views down to the most granular detail. Users need to easily navigate data generated by multiple systems and stored in multiple locations. Finally, the risk measures themselves and the associated aggregation logic are rapidly becoming more complex and require more computational power, which in turn stresses existing infrastructures.

These challenges are exacerbated by the growing data volumes and pressure to report on the most granular data and then to aggregate it across any possible dimension. Performance is the critical metric in evaluating success. Decision-makers can't wait days or even hours for risk reports, regardless of complexity. IBM® Algo® Workspace Analyzer offers an evolutionary approach with established performance benefits to meet these challenges.

Limits of existing architectures and tools

Traditional risk solutions are often comprised of two major components: pricing/simulation and aggregation/reporting. Pricing/simulation components continue to be enhanced as new security types are introduced. However, aggregation/reporting components are evolving even faster. In general, today's risk aggregations are intricate because results are more than simply linear combinations of raw inputs. In addition, the use of simulation-based risk measures means that for larger institutions, daily raw data volume can be at terabyte scale.

One challenge facing legacy aggregation/reporting components is that institutions working with larger data sets often want to see a consolidated view of the firm, while the data often resides in different groups and departments. To solve this problem, firms often build data warehouses for staging the data. They expend resources building and maintaining the extraction, transformation and loading (ETL) logic required to bring in data from disparate sources across the enterprise. In addition to the effort required to maintain all of the ETL logic, it is common practice to load pre-aggregated data. This makes it difficult to run reports on the most granular data and limits the flexibility to add new reports that rely on data not available in the warehouse.

For large data volumes, most conventional aggregation/reporting components rely on various pre-aggregation techniques. Approaches such as online analytical processing (OLAP) are a devil's bargain. In return for simpler technical implementation, users are faced with inflexible solutions that are difficult to adapt to new requirements. These so-called "solutions" cannot answer unforeseen questions that require aggregations across dimensions that are inconsistent with the pre-aggregation.

There is a growing business need for risk management systems to perform new analyses rather than generate new reports on existing data. This requires solutions that can re-compute results under different scenarios on demand. Legacy systems struggle to meet modern risk management demands where the aggregation and reporting layer must perform second-stage analytics. Consider an institution running portfolio stress tests. In the past, it was sufficient to save portfolio stress test results over several years for historical reporting. Today, institutions want to back-test new portfolios over the stress tests. Yet, it is not possible to run back-tests on new portfolios if the system only saved results aggregated at the portfolio level. To run these back-tests, systems must save results at the more granular instrument level. Furthermore, to back-test a new scenario is even more complicated. Before the final report can be created, the new scenario must be generated, the simulation run and the results must be aggregated. Additional data is required for each of these steps. Managing increasing amounts of data and supporting reporting requirements that trigger new calculations (for example, run new scenarios), are big challenges for existing tools.

Requirements for next-generation aggregation tools

From a technical perspective, computational and data challenges need to be overcome to achieve a solution that is powerful and flexible enough to meet both current needs and future requirements. To deliver solutions that are not limited to a fixed set of pre-aggregated data, aggregation technology must harness significant amounts of hardware. This helps deliver solutions that can aggregate and filter any aspect of data and report at any level ranging from the enterprise "top-of-house" level down to the most granular levels of detail. Furthermore, solutions need to be both scalable and cost-effective. To achieve these goals, they must be completely elastic and compatible with the most generic hardware. Finally, the solutions need to be equally at home when deployed in the data center or in the cloud.

Modern risk aggregation solutions need to function easily in the complex and sometimes ugly world of production data. They must be able to source data from many existing sources, rather than require a transformation or movement of data to an intermediate location. In addition, aggregation solutions need to serve multiple functions in the broader risk architecture. The solutions need to work as both front-end analytical workbenches and back-end engines for other downstream processes.

Algo Workspace Analyzer delivers a proven, evolutionary approach

Algo Workspace Analyzer uses different methods for risk aggregation than traditional tools. It helps organizations consolidate risk data and analytics from diverse systems more efficiently.

Algo Workspace Analyzer presents dynamic reporting to decision makers, improving the quality and comprehensiveness of risk reporting. All data is loaded in memory and spread across multiple servers, supporting a massively parallel architecture. These features provide updated benefits that address modern requirements.

Data flexibility — Data and analytics obscured within the firm's existing infrastructure can easily be accessed from different locations such as files, databases, trading systems, messaging systems, big data environments and more. Data adapters can easily be written, allowing data to be read in any format and from any system. Data is retrieved from the various "golden sources," eliminating the need for intermediate data warehouses. Data permissions control what users can access, supporting flexible operating models. For example, a user might be able to see aggregated data for the entire firm, but only be able to drill into and see lower-level data for their department.

High performance with scalability — Calculations are parallelized across multiple servers and multiple cores using commodity hardware. New cores, compute nodes and compute services can be added dynamically to scale for peak workloads. The ability to utilize a pool of commodity and potentially heterogeneous hardware provides the ability to dynamically flex the hardware footprint up and down. These features help reduce the total cost of ownership (TCO).

Extending analytics — Calculations that roll up and aggregate risk data are more easily implemented as new Java classes. There is a large number of pre- and post-hooks that permit calculations to be highly customized. The Java classes can be loaded either statically or dynamically to efficiently build the complex logic required by both new financial regulations and internal management reporting. Simple expressions can be written in a scripting language that is cross-compiled into Java at execution time, allowing business users without Java skills to add analytics. In addition to implementing new embedded calculations, it is also possible to trigger calculations to be run in other systems, retrieve additional data, or perform other complex operations.

Visualization — Interfaces are open and easy to customize, allowing users to build reporting dashboards with the ability to interactively drill through data elements to the underlying analysis and data, helping users understand the key risk factors and contributing sources. Users define and save reports by selecting data elements, dimensions and corresponding visualization elements.

Options include tables, bar graphs, pie graphs, histograms, scatter plots and more. Reports are interactive, providing high-speed slice-and-dice of instrument and portfolio analytics, so that changing a parameter (for example, the portfolio to view), will cause the report to be redisplayed. Reports can also be linked so users can easily navigate from a parent report to see more details in a child report. The linkage is dynamic so that changes in the parent report are reflected in the child report.

Dependence graphs — All the data and calculations are linked dynamically. Changes to any input data or formula will cause all corresponding outputs to automatically re-compute. In turn, all reports on users' screens will update to reflect these changes. This feature serves two purposes:

- The system can be set up to be real time and data can be brought into the system to reflect business and market activity. Trades, market data or changes to any other activity data can be continually fed into the system, and all reports will update instantly.
- Sandboxes can be created to support what-if scenarios that can be quickly run and analyzed. Changes can be made to any aspect of the data, and all outputs in that sandbox will be updated instantly. Examples could include simple changes to single input values, large sets of proposed trades or even edits to the corporate hierarchy.

Collaborative environment and sandboxing — All users share a common copy of in-memory data, so that all reports are consistent by default. Individual users can share views and reports but also build individual reports. Users can create a sandbox that allows them to change both data and calculations, which can be used to investigate data anomalies, compare multiple versions of a calculation, pose scenarios and other tasks. Changes to the user's sandbox are private and do not impact the work of other users. For collaborative work, other users can be invited to access a sandbox. When changes are made in a sandbox, the dependence graph is used to determine what needs to be recalculated, leading to improved performance.

Case studies exhibit exceptional performance

Performance is a key metric for evaluating the success of a modern risk aggregation engine. Legacy aggregation engines make trade-offs that limit reporting flexibility to achieve acceptable performance. In contrast, Algo Workspace Analyzer provides an architecture designed to meet today's reporting requirements with exceptional performance.

Some of the main factors that influence performance are type of hardware, hardware utilization, number of current users and types of reports being requested. In the three case studies below, the hardware is dedicated for a single user. The reports are all cold full aggregation, meaning no data or results are cached. Subsequent reports using cached data would be significantly faster.

Client A —Market risk example

Algo Workspace Analyzer provides a fully interactive interface that allows for dynamic reports with day-over-day comparisons. In an example that provides market risk exposures across a broad range of asset classes and portfolios, value at risk (VaR) and other risk analytics are computed using Monte Carlo simulation results for different confidence intervals and across multiple time horizons. Daily VaR measures are also computed and compared over a series of consecutive periods.

Data set size

	Value
Positions/day	200,000
Risk measures (VaR)	14
Scenarios /measures	500
Nodes in hierarchy	100
Number of days	10
Total number of P&L values	14 billion

Calculations performed

Description
Top-of-house aggregation for 14 scenarios sets on 1, 2, 4, and 8 dates
Test results show a cold full re-aggregation and reporting scenario (no caching)
Times are in seconds

IBM SoftLayer® environment accessed via internet

	Value
Bare metal servers	E5-2690
Number of hosts	7
Cores/host	12 (84 total)
I/O channels/host	3
Memory/host	64GB
Core speed	2.6GHz
Interconnect	10Gb Ethernet
Client connect speed (via internet)	50Mb

Results — Scalability across cores

Test	7	14	21
1 day Top-of-house VaR/14 measures	4	2.5	1.5
1 day by node VaR/14 measures	4	3	2
10 days Top-of-House/14 measures	12	7	5.5
10 days by node (sec)/1 measure	4	4	5

Results — Scalability across dates (Time in seconds)

Test	Time
1 date	1.5
2 dates concurrent	1.5
4 dates concurrent	2
8 dates concurrent	4

Client B — Fundamental Review of the Trading Book example

Algo Workspace Analyzer computes required analytics to help banks comply with Fundamental Review of the Trading Book (FRTB) reporting on minimum capital requirements for market risk. The FRTB reporting requirements include:

- Mandatory reporting on standardized approach (SA) results, where the standardized market capital charge needs to be computed at the desk level.
- Internal model approach (IMA) that includes internal model capital charge to capture tail risk and capital adequacy during periods of significant financial market stress and incorporation of the risk of market illiquidity.

This example focuses on data loading and in-memory caching for a medium-sized bank. It computes the aggregate capital charge as well as the internal capital charge on the approved desks and the standardized capital charge on the unapproved desks.

Data set size

Parameter	Value
Positions	400,000
Instruments	400,000
Risk measures (VaR)	Not applicable
Scenarios/measures for IMA	22,500
Sensitivities for SA	Not applicable
Nodes in hierarchy	100
Number of days	2

Calculations performed

Description
FRTB for both IMA and SA including expected shortfall for the prescribed risk groups and liquidity horizons, deltas and curvatures
Reports aggregate more than 270,000 positions across 55 desks from 63 different profit and loss distributions for full/current, reduced/current and reduced/stress simulations
Test results show a cold full re-aggregation and reporting scenario (no caching)
Times are in seconds

IBM SoftLayer® environment accessed via internet

Parameters	Value
Virtual machines	E5-2683
Number of hosts	3
Cores/host	32 (96 total)
I/O channels/host	2
Memory	640GB
Core speed	2.0GHz
Interconnect	1000 Mbps
Client connect speed (via internet)	50Mb

Results

Test	Time
Total data retrieval	6.4 minutes
Retrieve portfolio and instrument files	<1 minute
Retrieve IMA simulations, including default risk charge (DRC) simulation	5 minutes
Retrieve SA sensitivities	<1 minute
All key reports returned	5 – 20 seconds
Base Report Generation Top-of-house, aggregated Total Capital Charge including: IMA aggregate capital (63 simulations x approximately 300 scenarios) IMA DRC (100,000 scenarios) SA (5.0 million sensitivities)	5.6 seconds

Client C — FRTB example

Algo Workspace Analyzer computes all the required analytics to help banks comply with Fundamental Review of the Trading Book (FRTB) reporting on minimum capital requirements for market risk. The FRTB reporting requirements include:

- Mandatory reporting on standardized approach (SA) results, where the standardized market capital charge needs to be computed at the desk level.
- Internal model approach (IMA) that includes internal model capital charge to capture tail risk and capital adequacy during periods of significant financial market stress and incorporation of the risk of market illiquidity.

While example B illustrates performance on data loading and computation of the Aggregate Capital Charge, this example highlights the Algo Workspace Analyzer capabilities of drill-through and computation dependency. The aggregate capital charge as well as the underlying internal capital and standardized capital charges are broken down into their components. Each desk is further sliced into more granular groupings for in-depth drill-through analysis.

Conclusions

Rapidly emerging business requirements are rendering traditional risk solutions obsolete. Both regulatory reporting requirements and internal management reporting requirements require solutions that can do more than just generate reports. Solutions must re-compute results on-demand and must allow users to navigate from top-level aggregated numbers down to the most granular level of detail. In addition, the data involved has grown larger and more complex. Solutions must be able to roll up results from the most granular level of data and not rely on pre-aggregated results. In light of these requirements, performance becomes the key metric when evaluating the success of a modern risk solution. Algo Workspace Analyzer is a modern architecture designed to deliver the performance users require while addressing a broad set of risk aggregation demands for financial institutions, covering three pillars: user experience, analytics and technology.

Data set size

	Value
Positions	650,000
Instruments	400,000
Scenarios/measures for IMA	22590
Sensitivities for SA	6,000,000
Nodes in hierarchy	100
Number of days	3

Calculations performed

Description
FRTB for both IMA and SA were calculated using different levels of aggregation
Test results show a cold full re-aggregation and reporting scenario (no caching)
Times are in seconds
Aggregation was completed at the following levels: – 2 levels: Desk->Book – 4 levels: Desk->Book->Region->Currency – 7 levels: Desk->Book->Region->Currency->Industry Sector->Obligor->Position

IBM SoftLayer® environment accessed via internet

Parameters	Value
Virtual machines	E5-2683
Number of hosts	4
Cores/host	56 (224 total)
I/O channels/host	2
Memory	247GB (988 GB)
Core speed	2.0GHz
Interconnect	1000 Mbps
Client connect speed (via internet)	50Mb

Results

Test	2	4	7
IMA Desk 1	25	25	50
IMA Desk 2	60	55	110
SA Desk 3	75	40	60
SA Desk 4	65	80	65

The analytic landscape in risk aggregation is constantly evolving. Algo Workspace Analyzer delivers features designed to satisfy the most demanding reporting and analysis needs. It empowers users to ask novel questions, perform wide- ranging what-ifs, have complete transparency at all levels of the organization and react in real-time. All these tasks can be done across the enterprise. Furthermore, Algo allows clients to build on pre-existing risk calculations or build their own from scratch. In either case, complex aggregation regimes can be quickly built for the most bespoke need.

Algo Workspace Analyzer technology drives the performance and flexibility. The flexible approach to data modeling means that many inputs can come from multiple sources across the institution. The in-memory architecture allows for data to be distributed across multiple servers, delivers established performance and supports transparency to the most granular level. Finally, Algo Workspace Analyzer is designed to be quickly and seamlessly deployable, while placing minimum demand on existing infrastructure.

About IBM Watson Financial Services

IBM works with organizations across the financial services industry to use IBM Cloud™, cognitive, big data, RegTech and blockchain technology to address their business challenges. IBM Watson® Financial Services merges the cognitive capabilities of Watson and the expertise of Promontory Financial Group to help risk and compliance professionals make better informed decisions to manage risk and compliance processes. These processes range from regulatory change management to specific compliance processes, such as anti-money laundering, know your customer, conduct surveillance and stress testing.

For more information

For specifics about IBM Algo Workspace Analyzer, please visit http://ibm.biz/AlgoWorkspaceAnalyzer

To learn more about IBM RegTech solutions, please contact your IBM representative or IBM Business Partner, or visit the following website: ibm.com/regtech

IBM Corporation
IBM Watson Financial Services
New Orchard Road
Armonk, NY 10504

Produced in the United States
of America, October 2018

71015671USEN-03

Watson Financial Services