WHITE PAPER

Mastering Market Risk Capital

Navigating the Fundamental Review of the Trading Book (FRTB)

As global regulators enforce the Fundamental Review of the Trading Book (FRTB), financial institutions must navigate a complex new landscape for market risk. This analysis focuses on the central strategic decision: the trade-offs between the straightforward Standardized Approach (SA) and the demanding Internal Models Approach (IMA). We dissect the key implementation hurdles—from desk-level governance, data integrity, and model validation to managing divergent jurisdictional rules—and outline how modern, scriptable risk calculation engines (platforms that allow users to write and execute custom risk calculations using programming interfaces) provide a robust path to not only meet compliance but also enhance risk management capabilities for the future.

MathWorks®

Executive Summary

The Fundamental Review of the Trading Book (FRTB) is reshaping how banks quantify and manage market risk. It replaces Value at Risk with Expected Shortfall for internal models and sets rigorous validation standards, demanding granular, desk-level analysis and recalibrated capital allocations.

Implementation timelines vary across jurisdictions and continue to evolve. Some supervisors in Asia-Pacific are moving earlier, while Europe and the U.K. are sequencing the Standardized Approach first and Internal Models thereafter. In the U.S., agencies are still finalizing the rule text and effective timing. These staggered schedules reduce "big-bang" risk but require banks to plan for parallel go-lives and rolling adjustments as rules are finalized.

Financial institutions must decide between a revised Standardized Approach (SA) and a more complex Internal Models Approach (IMA). The SA calculates capital using prescribed formulas and serves as the mandatory fallback for the IMA. By contrast, the IMA promises greater risk sensitivity but requires daily backtesting, Profit & Loss (P&L) attribution tests, and approval at the desk level—a hurdle that only a few banks are prepared to tackle.

Data quality, especially for long horizons and illiquid products, is critical, as errors in the tail of loss distributions can bias Expected Shortfall calculations. This paper discusses the strategic decisions facing banks, highlights challenges around data, validation, and desk-level governance, and outlines how a flexible, scriptable risk engine can help firms meet these obligations while building a more insightful risk function.

Introduction: The New Landscape of Market Risk

The 2008 financial crisis exposed weaknesses in capital rules. Regulators responded with the Basel III package and, within it, FRTB. The new framework clarifies the boundary between the trading and banking books and requires banks to hold capital that better reflects the risks of their trading positions.

Under the Standardized Approach, banks compute three components: a sensitivity-based charge, a default risk charge, and a residual risk add-on. These components capture changes in market factors, jump-to-default risk, and risks not accounted for in standard models, respectively.

The IMA uses Expected Shortfall at a 97.5% confidence level and requires regulatory approval at the trading desk level. By demanding more stringent validation and more granular data, FRTB aims to make capital requirements more comparable across jurisdictions and reduce the risk of regulatory arbitrage. However, differences in national timelines risk fragmenting the market. Banks operating globally must plan for these staggered deadlines and the possibility of further delays.

Methodology Choices: Standardized vs Internal Models

Choosing between the SA and the IMA is a strategic decision with implications for capital, resources, and governance.

The **Standardized Approach (SA)** is formulaic and generally simpler to implement, though potentially more conservative in its capital calculations. Its key characteristics include:

- It does not require prior regulatory approval to use.
- It calculates capital using prescribed risk weights, correlations, and aggregation formulas set by the rule text.
- The capital charge has three components: **Sensitivity-Based Measure** (SBM) — including delta, vega, and curvature; **Default Risk Charge** (DRC); and **Residual Risk Add-On** (RRAO).
- It operates at the **trading-desk level**, enabling transparent attribution by risk class/bucket and consistent firm-wide roll-ups.
- It is the **mandatory default and fallback** for every desk; even IMA-approved desks must maintain an SA calculation as a floor and contingency.
- Scriptable SA engines should **ingest or emit ISDA CRIF** (Common Risk Interchange Format) for interoperability and testing, while formal supervisory reports remain **jurisdiction-specific** (CRIF is an industry interchange format, not a regulatory report).

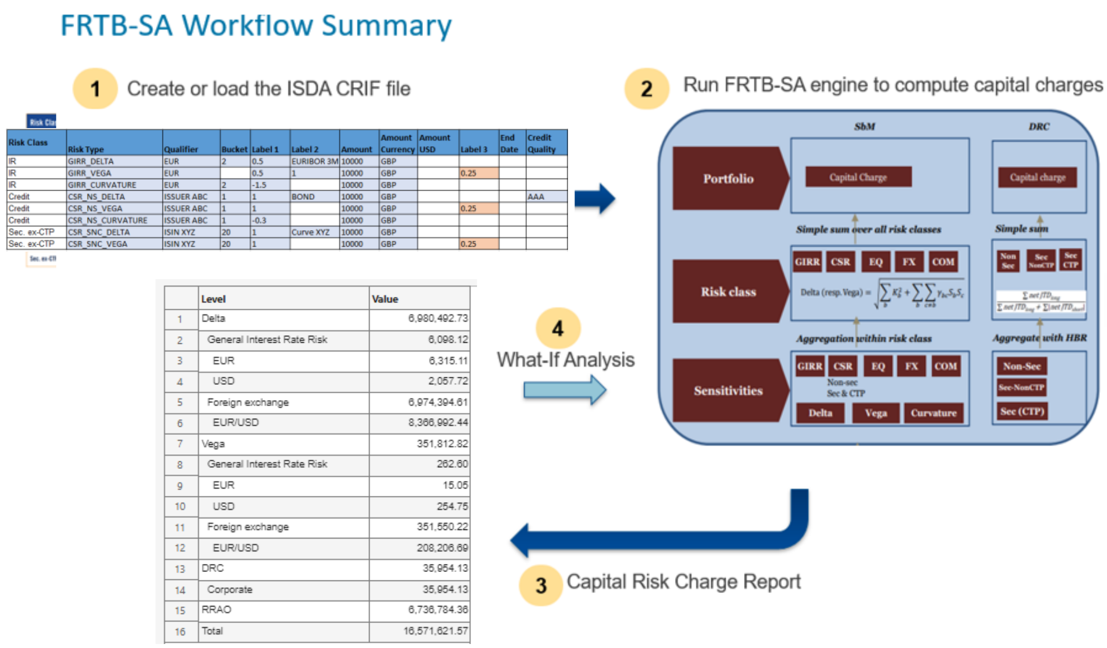

***Figure 1:** FRTB-SA workflow in MATLAB—from ingesting ISDA CRIF to capital computation, reporting, and what-if analysis. See also SA examples for [Basel] and [CRR].*

The **Internal Models Approach (IMA)** offers greater risk sensitivity and the potential for lower capital but comes with significant operational and validation overhead. Key requirements include:

- **Modelling Expertise:** Banks must demonstrate a sophisticated ability to model risk.
- **Rigorous Testing:** Firms must pass several stringent tests, including risk factor eligibility tests to ensure data quality, daily backtesting, and P&L attribution tests to prove the model's accuracy.
- **Desk-Level Approval:** Approval is granted on a desk-by-desk basis, meaning risk must be effectively measured and modelled where it is taken.

Given its demanding nature, only a handful of European banks are preparing to adopt the IMA; most institutions plan to start with the SA and consider moving to the IMA later.

The Three IMA Prerequisite Tests

- **P&L Attribution (PLA) Test:** This test checks if the bank's risk model "thinks" like its front-office pricing model. It works by comparing the daily profit and loss (P&L) calculated by the risk management model against the P&L from the trading systems. If the two P&L streams are not statistically similar, it proves the risk model is missing key factors and doesn't accurately reflect how the desk's portfolio behaves.
- **Backtesting:** This test determines if yesterday's risk forecast was good enough for today's reality. It compares the model's risk forecast (specifically, its Value-at-Risk or VaR) to the actual profit or loss the desk experienced on the following day. If actual losses breach the VaR forecast more often than statistically permitted, the model is considered unreliable for predicting potential losses.
- **Risk Factor Eligibility Test (RFET):** It vets each risk factor to confirm enough recent, 'real' price observations from eligible transactions or committed quotes; factors that fail are Non-Modellable Risk Factors (NMRFs) and attract a punitive capital charge.

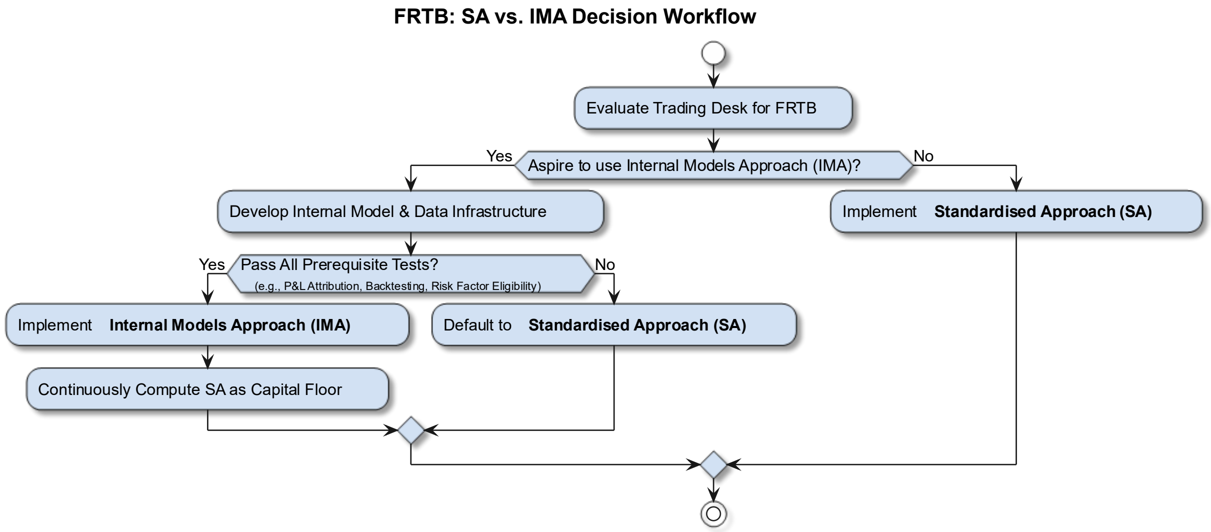

Figure 2: The Strategic Decision Path. *This workflow illustrates the choice between the Standardized Approach (SA) and the Internal Models Approach (IMA). The IMA path requires desk-level approval and ongoing tests (e.g., risk-factor modellability, P&L attribution, backtesting); SA remains the default/fallback.*

Successfully navigating this strategic path hinges on a robust foundation of high-quality data and precise risk measurement.

Data Quality and Risk Measurement

Accurate capital calculations depend on high-quality data. The SA requires detailed sensitivities for each risk factor, while the IMA needs historical time series for market prices, volatilities, and correlations. For both the SA and IMA frameworks, the ISDA (International Swaps and Derivatives Association) FRTB-SA CRIF (file, based on industry-standard XML or CSV formats, is widely used for portfolio interchange, validation, and unit testing across SA/IMA workflows.

For the IMA, banks must go beyond static sensitivities and construct time series of real price observations to prove risk factor modellability. Data gaps in the long tail of distributions are especially problematic because the Expected Shortfall metric averages losses in the tail. Small errors or outliers can lead to overstated or understated capital. Institutions, therefore, need robust processes to identify and cleanse data, perform outlier analysis, and document assumptions. They must also understand how different national versions of the FRTB treat data, such as varying probability-of-default floors or correlation structures and ensure consistency when aggregating results.

Performance: Meeting Daily Compute Demands

FRTB workloads—SA sensitivities/DRC/RRAO and IMA ES with daily VaR backtesting and P&L attribution—are compute-intensive at the desk level. Design for throughput, determinism, and auditability from day one.

- **Parallelize end-to-end**: Vectorize calculations; distribute by desk/portfolio/date across cores, nodes, and GPUs; treat scenario/path loops as the parallel dimension.
- **Burst to cloud, safely**: Containerize engines; autoscale for peaks; enforce cost/time guardrails; keep market/position data close to compute; encrypt and log access.
- **Reuse what's expensive**: Cache scenario shocks and RNG seeds; persist sensitivities/Greeks; do incremental recomputes for changed trades/risk factors; version inputs/outputs for reproducibility.
- **Orchestrate the daily loop**: Build a DAG from ingest → calc → tests (VaR backtest; P&L attribution) → reports; emit immutable artifacts with provenance; alert on SLA misses.
- **Accelerate exotic pricing**: Where nested MC/PDEs dominate, consider surrogate models (e.g., **physics-informed neural networks**) or AAD to speed Greeks—gate with model-risk controls (hold-out tests, error budgets) and document limits.

See it in MATLAB: scriptable SA/IMA workflows map to parallel pools, batch jobs, and GPU-accelerated pricing, on-premises or in the cloud

Managing data at scale calls for a disciplined architecture. Banks should establish central data repositories that collect market and position data, enforce quality rules, and supply both SA and IMA calculations. These repositories must support granular tagging of positions so that capital can be attributed to the right desk. Modern hardware—multi-core CPUs and GPUs—as well as cloud resources, can accelerate large simulations. The computational load is significant: default risk charges may require hundreds of thousands of Monte Carlo paths per desk, and backtests must run daily. A robust pipeline is also essential for the reproducibility and audit trails that regulators expect.

Validation and Governance

The IMA's promise of lower capital hinges on rigorous validation. While the following validation requirements are specific to the IMA, firms using the SA also benefit from robust model governance and data quality frameworks, though with different testing requirements. Banks must perform daily backtesting of VaR using Actual P&L (APL) and Hypothetical P&L (HPL) and run P&L Attribution (PLA) tests comparing Risk-Theoretical P&L (RTPL) from the risk model with Front-Office HPL. Only desks that meet thresholds on these tests and demonstrate sufficient real-price data qualify for IMA capital. Persistent failures against regulatory thresholds trigger remediation and may lead to reversion to SA. Regulators like the Prudential Regulation Authority see the P&L attribution test as a scientific way to assess whether models capture all material risks.

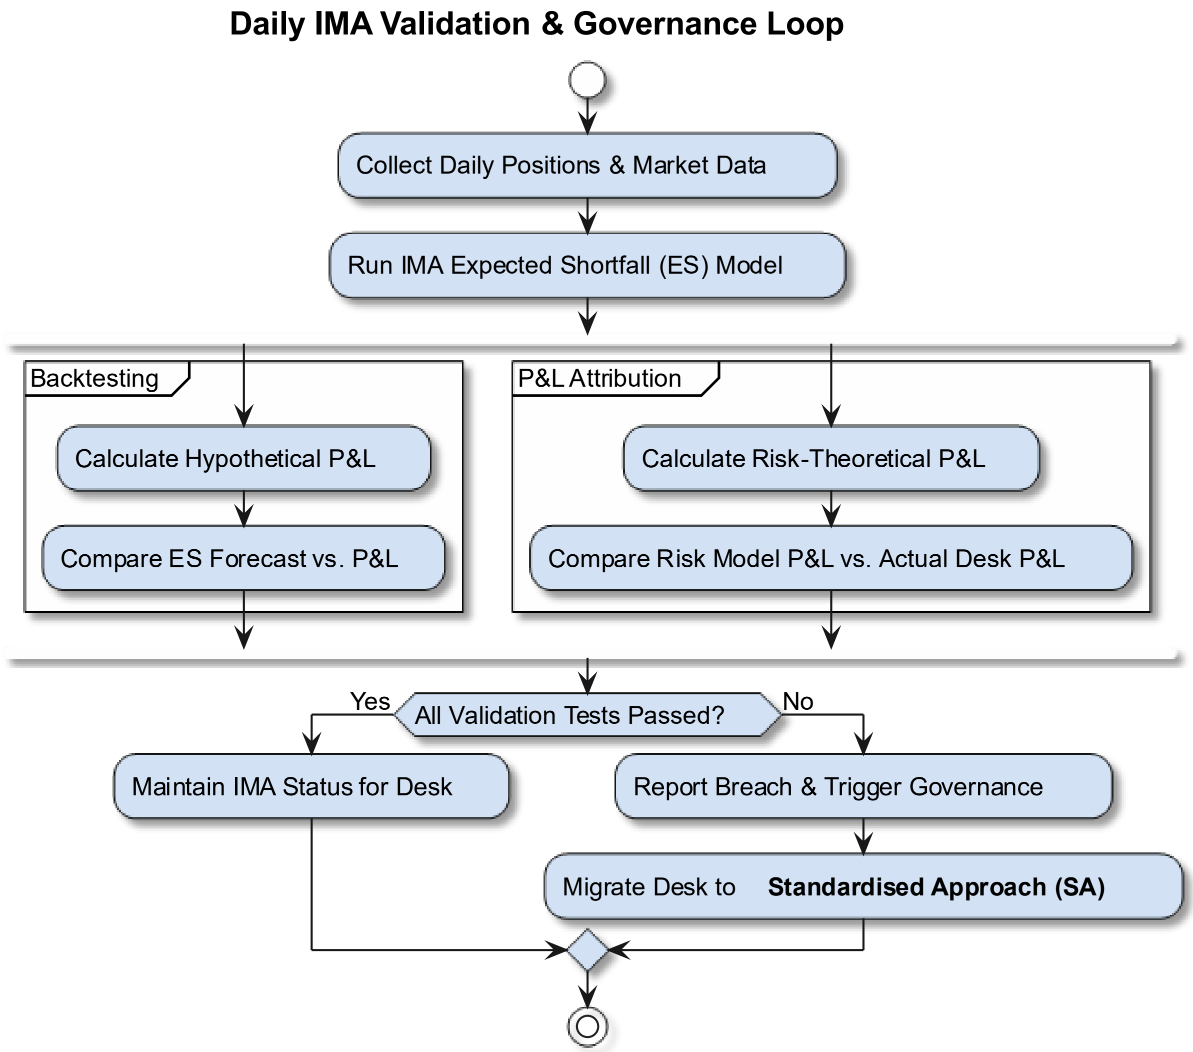

Figure 3: The Daily IMA Validation Loop. *Desks approved for the IMA must complete this rigorous validation workflow daily. Failure in either the backtesting or P&L attribution tests can force a desk back onto the more conservative Standardized Approach, highlighting the significant ongoing operational burden.*

Governance extends beyond models. FRTB makes the trading desk the unit of measurement: capital, model approval, and data quality must all be managed at this level. Desk heads are responsible for understanding how their trades drive capital. Senior management must oversee cross-desk consistency and ensure that risk governance frameworks are harmonized across jurisdictions. Banks should establish committees to manage model approvals, handle requests to move positions between trading and banking books, and monitor risks not captured in the internal model (often called RNIM, or Risk-Not-in-Model).

Why Visualization Matters for FRTB Validation

FRTB validation is a daily, desk-level practice—not an occasional audit. Clear visuals shorten the path from data to decision, letting you spot model/FO misalignment, exceptions, and modellability gaps in time to keep an IMA desk onside—or to revert cleanly to SA when required. These visuals plug directly into the daily loop you already run (ingest → calc → backtesting/PLA → report), and make that loop easier to govern and defend.

- **PLA scatter + UPL timeline**. A scatter of HPL (x) vs RTPL (y) with a 45° reference line shows alignment at a glance; a companion timeline of unexplained P&L (HPL−RTPL) with ±2σ bands flags outliers to investigate (mappings, risk factors, pricing models). Use this pair in the daily validation loop for each IMA-aspirant desk.

- **Backtesting dashboard (99% VaR).** Overlay APL and 1-day 99% VaR to highlight breach days, and track the cumulative exception count against green/amber/red thresholds. This keeps breaches actionable (what happened, where, and why) and supports governance when a desk drifts toward amber or red.

- **RFET heatmaps**. Visualize modellability across your factor dictionary—modellable vs. NMRF—by tenor, instrument, or liquidity horizon. You'll see where observation density is thin, prioritize data remediation, and quantify the capital impact of NMRFs before approval discussions.

- **Capital-attribution treemaps**. For SA, decompose by SBM / DRC / RRAO; for IMA, show ES and NMRF add-ons. Desk heads can trace capital to drivers (risk class, bucket, curve node) and test what-ifs (hedges, risk transfer) against both SA and IMA views.

How to make this operational? Build these views on a single, scripting-based ("scriptable") platform shared by SA and IMA so you reuse the same data ingestion, factor taxonomy, controls, and reporting. Emit immutable artifacts (images/data) in your nightly run, tag them by desk/date, and keep units and definitions consistent (APL/HPL/RTPL). This reinforces auditability and reduces effort when rules or jurisdictions shift (Basel/CRR3/PRA variants).

Example Backtesting dashboard

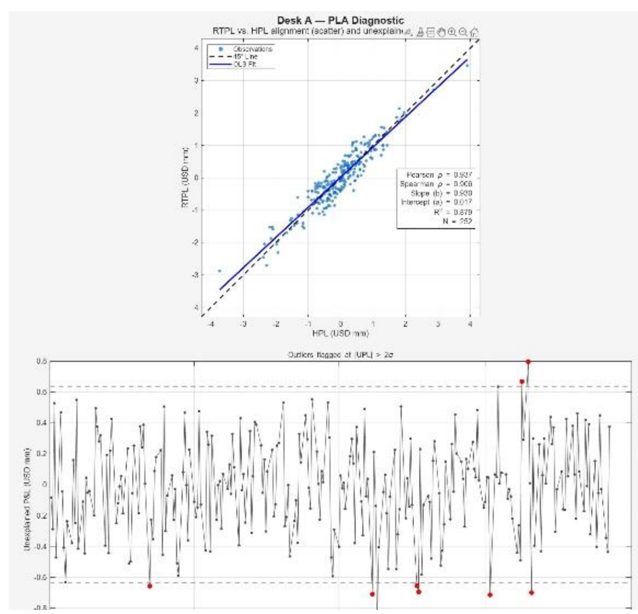

Example PLA Scatter and UPL Timeline

Desk-level Focus and SA versus IMA Comparison

FRTB requires risk and capital to be measured at the level where it is taken: the trading desk. This design addresses a key failing of the past, where losses were hidden in aggregated portfolios. Under the SA, each desk calculates capital using prescribed formulas, and no regulatory approval is needed. Under the IMA, the desk builds internal models but must also compute the SA as a floor and fallback. This means even desks that aspire to use the IMA must first implement the SA as a reference and contingency plan.

Comparing the SA and IMA at the desk level highlights several trade-offs. The SA imposes fixed risk weights and correlations, which can misrepresent risk compared to tailored models. By contrast, the IMA allows banks to model full distributions and to include non-linear sensitivities such as vega and curvature explicitly. However, this freedom comes with heavy validation burdens. Many firms are intimidated by the IMA and are choosing the SA for its predictability, while larger institutions may adopt a hybrid strategy.

Desk heads need tools that illuminate how their trading activities translate into capital. They must be able to see their desk's contribution to the bank's total market risk capital, identify the drivers of that charge, and assess the marginal impact of new trades. Such insights support better pricing, limit management, and hedging decisions.

How to Build a Scalable, Modular FRTB Technology Stack

A successful FRTB program requires technology that is transparent, scalable, and scriptable. Using the same platform for both SA and IMA approaches can provide operational efficiencies, consistent data handling, and simplified model governance across methodologies. Many banks are building integrated architectures that separate data ingestion, calculation engines, model prototyping, and validation layers.

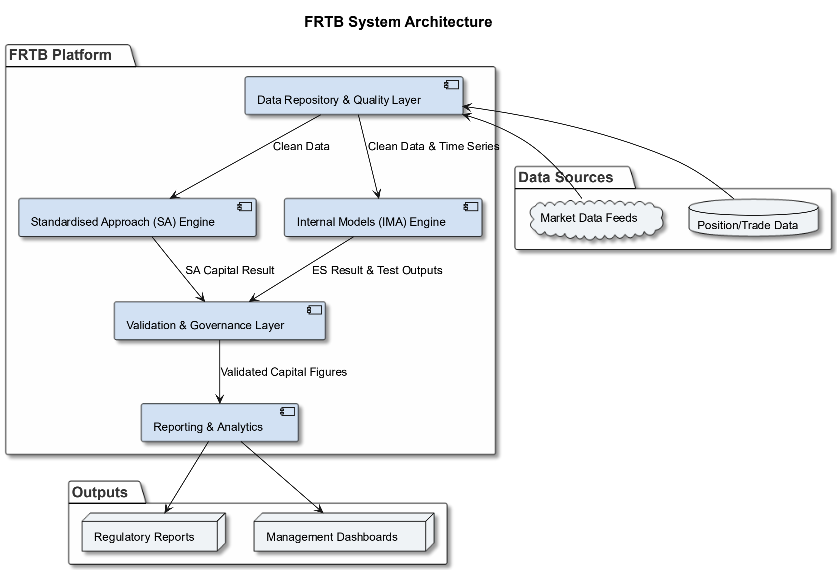

Figure 4: A Scalable FRTB Technology Architecture. *A robust implementation requires a modular architecture like the one shown. Data is ingested, cleansed, and fed into parallel engines for both SA and IMA calculations, with dedicated layers for validation, governance, and reporting to ensure accuracy and auditability.*

Banks should favor scriptable engines for the SA. For example, platforms like MATLAB, with specialized packages like the Financial Instruments Toolbox, can provide such an engine. They can accept standardized CRIF files, construct a regulatory object representing the portfolio, compute all required charges, and aggregate them into a total capital requirement. The same engine can be configured for Basel, CRR3, or other regional rules, offering a consistent calculation core. Because it is scriptable, quants and risk managers can embed the engine within automation pipelines, integrate it with data-quality routines, perform what-if analysis, and create insightful visualizations that help stakeholders understand risk drivers and model behavior—critical capabilities for the complex validation processes FRTB demands.

Beyond the SA, toolboxes such as Risk Management Toolbox (which provides VaR and Expected Shortfall backtesting capabilities that have been available for years, with new validation features in recent releases) can provide backtesting utilities and risk-analysis functions that serve as building blocks for IMA prototyping. By combining a standardized calculation engine with a flexible environment for model development, banks can prepare for both SA compliance and a potential transition to internal models.

Options and Strategic Considerations

Banks face choices not only between the SA and IMA but also in how they sequence their implementations and allocate resources. A pragmatic, step-by-step approach would be to:

- **Establish a Baseline:** Implement the SA as the baseline across all trading desks. This ensures a compliant capital charge can be calculated for the entire firm and that the necessary data infrastructure is in place.
- **Identify High-Impact Desks:** From this baseline, conduct gap analyses to identify desks where the potential capital benefit of using the IMA significantly outweighs the implementation and maintenance costs.
- **Prioritise Investment:** Use scenario and sensitivity analysis to quantify potential capital savings, inform business cases, and prioritise the development of internal models for the most promising desks.
- **Maintain SA as Fallback:** Remember that the SA remains the mandatory fallback for all desks. Maintaining a high-quality SA engine is therefore essential, even for desks that successfully move to the IMA.

Another consideration is jurisdictional alignment. A global bank may choose the IMA in one jurisdiction and the SA in another. When planning across jurisdictions, firms must map portfolios to local versions of the rules and consider differences such as probability-of-default floors or index look-through provisions. They should also be alert to future changes; regulators have already adjusted timelines and may do so again.

Technology choices should support agility. Scriptable engines — platforms that allow users to write and execute custom risk calculations through programming interfaces— allow firms to implement regulatory changes quickly. Modular architectures let teams swap components without rewriting the whole system. Investing in cloud capacity can reduce time-to-market for simulations and support scalability. The ultimate goal is not merely compliance but a risk management function that can adapt to new regulations, products, and stress scenarios.

Conclusion and Outlook

The FRTB is a transformative regulation that compels banks to adopt more risk-sensitive measures, improve data quality, and embed rigorous validation into daily processes. Its phased implementation adds complexity but also provides an opportunity to sequence investments. The Standardized Approach offers a clear, formulaic route to compliance and will be broadly adopted. The Internal Models Approach remains attractive for sophisticated desks but demands a disciplined model validation culture and robust infrastructure.

Whichever route banks choose, they will need to focus on granular, desk-level capital allocation, data governance, and scalable systems. By leveraging scriptable calculation engines, flexible modelling tools, and modern computing resources, firms can turn FRTB compliance into an opportunity to build more insightful and resilient risk management functions. Continuous monitoring of regulatory developments will remain essential as the Basel III reforms continue to evolve.

Accelerate Your FRTB Implementation:

https://www.mathworks.com/discovery/frtb.html

Contact your MathWorks Account Representative, or email compfin@mathworks.com