

Why Open-Source Can Win the Quant Infrastructure War — and What Comes Next

The quant infrastructure stack is being rebuilt from first principles. Transparent algorithms, cloud-native delivery, and AI-powered interfaces are converging — and the institutions that understand this shift can define the next decade of risk management.

By Dr Stuart Smith

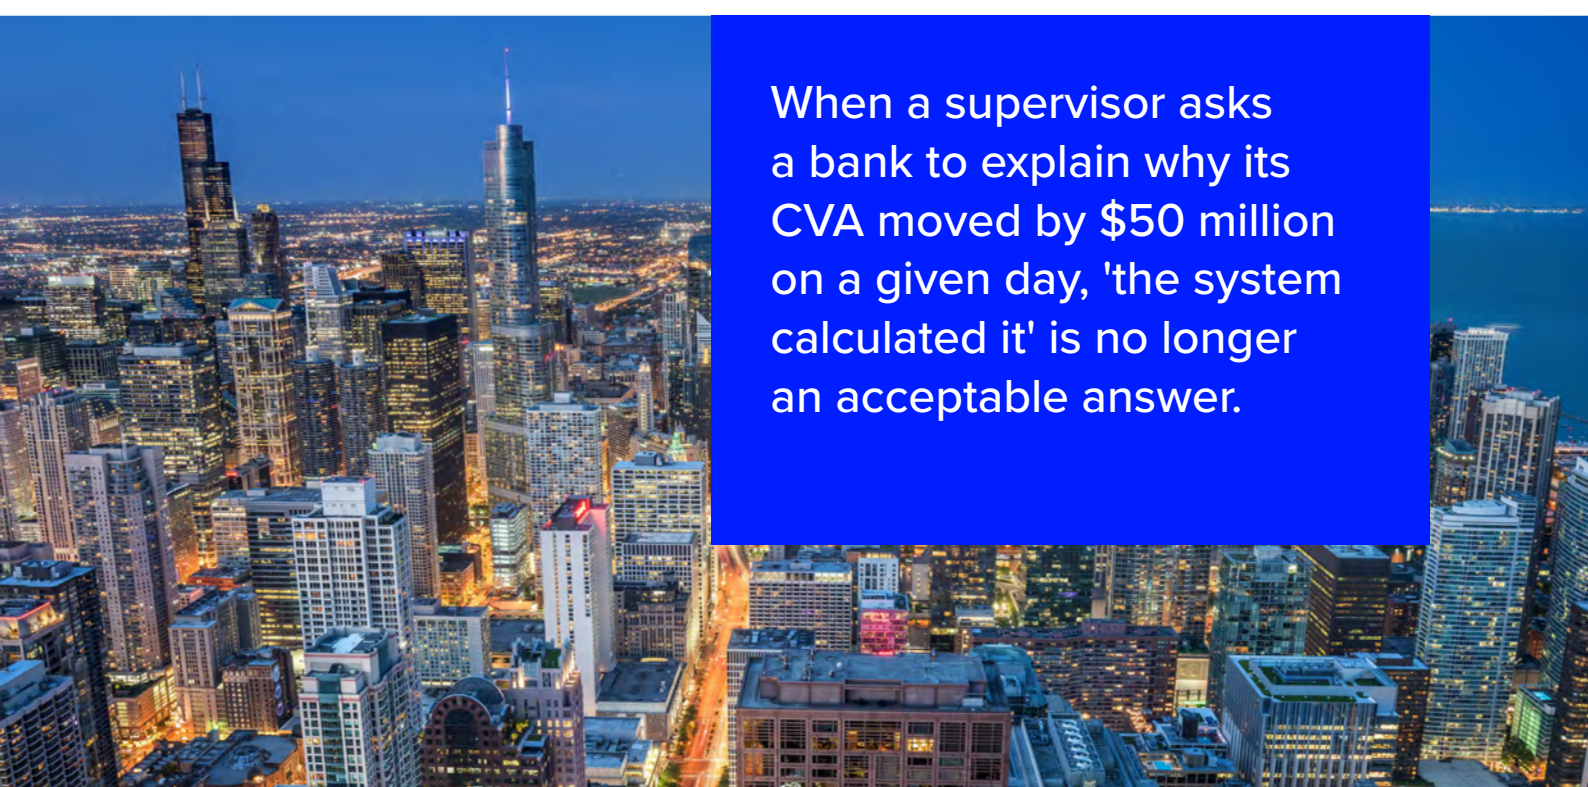
LSEG Post Trade Solutions



The proprietary era is ending

For three decades, the dominant model in quantitative finance infrastructure was simple: buy a black box. A vendor provided the models, the calibration, the pricing engines, and the risk aggregation. The buyer accepted the answer, trusted the number, and moved on. This arrangement made sense when quant talent was scarce, computing was expensive, and regulators asked few questions about methodology.

That world is ending — not because vendors have failed, but because the industry's requirements have fundamentally changed. Basel IV, FRTB, SA-CCR, and the expanding scope of XVA have made model transparency a regulatory imperative, not a luxury. When a supervisor asks a bank to explain why its CVA moved by \$50 million on a given day, 'the system calculated it' is no longer an acceptable answer. Institutions need to see inside the calculation.

An aerial night view of a city skyline, likely Chicago, with numerous skyscrapers illuminated. A large blue rectangular text box is overlaid on the right side of the image.

When a supervisor asks a bank to explain why its CVA moved by \$50 million on a given day, 'the system calculated it' is no longer an acceptable answer.

Open-source infrastructure addresses this directly. When the mathematics is published, peer-reviewed, and openly available, the model validation conversation becomes a genuine scientific dialogue rather than a negotiation with a vendor's professional services team. This is why the adoption of open-source quant libraries has accelerated so sharply among institutions over the past five years — not primarily because it is cheaper, but because its code is accessible and auditable.

ORE and the QuantLib lineage


The story of open-source quant infrastructure is in many ways the story of QuantLib. Launched in 2000 and now with over two decades of continuous development, QuantLib demonstrated that a peer-reviewed, community-maintained pricing library could match proprietary alternatives and, in some cases, offer greater transparency and flexibility — proprietary alternatives for the vast majority of financial instruments. It became the de facto reference implementation for interest rate derivatives pricing.

The Open-Source Risk Engine (ORE), originally developed by Quaternion Risk Management and now stewarded by LSEG Post Trade Solutions (formerly Acadia), extended that lineage into the counterparty risk and XVA domain. Built on QuantLib, ORE provides a complete framework for Monte Carlo exposure simulation, XVA calculation, sensitivity analysis, historical simulation VaR, and ISDA SIMM — all with openly published methodology, open-source code, and an active user community spanning sell-side banks, buy-side asset managers, and central counterparties.

The parallel to Linux is not accidental. Linux did not win the server infrastructure war because it was free. It won because open development is widely regarded as contributing to robust, secure and adaptable platforms through community collaboration a more robust, more secure, and more adaptable platform than any single vendor could maintain alone. The same dynamics are playing out in quantitative finance: the collective scrutiny of hundreds of quant developers can produce better algorithms than a proprietary team working behind closed doors.

ORE at a Glance

- Full-revaluation Monte Carlo exposure simulation using multi-factor models across IR, FX, Equity, Credit, Inflation and Commodity
- American Monte Carlo for accelerated exposure simulation
- XVA suite: CVA, DVA, FVA, COLVA, KVA, MVA with ISDA SIMM dynamic IM & AAD for fast sensitivity calculation
- Market risk: Sensitivity analysis, Historical Simulation VaR, stress testing
- FRTB and SA-CCR capital calculations
- Scripted trade framework for exotic product coverage without recompilation
- SWIG language bindings for Python and Java; REST API via ORE web service



How open risk analytics solves the infrastructure problem

Despite its analytical depth, open-source software has historically stopped at the boundary of deployment. A bank that adopts ORE still needs to solve the hard infrastructure problems: cloud provisioning, data pipelines from market data vendors, API management, high-availability scheduling, and incremental calculation architecture. These are not quant problems — they are engineering problems, and they are expensive.

This is the gap that Software-as-a-Service delivery closes. The calculations remain open, transparent, and auditable. The deployment, scaling, data integration, and availability become managed services. Quant teams can focus on model governance and business logic; technology teams are freed from maintaining bespoke infrastructure that is not a source of competitive advantage.

LSEG's position in this transition is structurally distinctive. The combination of ORE's open-source calculation depth with LSEG's Models as a Service (MaaS) marketplace means that the Open Risk Analytics (ORA) offering is not merely a hosted version of open-source software, rather it's part of a coherent ecosystem of complementary models governed and delivered by LSEG.



The AI layer:

From calculation to conversation

The next frontier is not faster computation or more complete data — both are rapidly becoming commodities. The next frontier is accessibility. Today, the full power of an XVA engine is accessible only to quants who can write code, configure XML, and interpret multi-dimensional output files. Risk managers, traders, and executives who most need to understand the numbers are typically one or two steps removed from them.

Large language models change this equation. An AI interface built on top of a structured calculation API can translate a question like 'what are my top five CVA concentrations and what macro scenarios drive them?' into a precise sequence of calculation calls, execute them against live market data, and can generate a structured explanation through a conversational interface answer — in seconds, through a conversational interface. This is not a futuristic vision. The underlying capability exists today.

The calculations remain open, transparent, and auditable. The deployment, scaling, and data integration become managed services. Quant teams focus on what only they can do.

LSEG's integration with industry Large Language Model (LLM) providers through the Models Context Protocol (MCP) brings exactly this capability to the ORA platform. LLM's reasoning ability — their capacity to understand financial context, decompose complex queries, and explain analytical results in plain language — combines with ORE's rigorous open-source algorithms to support new ways of accessing and interacting with risk analytics: risk analytics that are simultaneously institutional-grade in methodology and accessible in interface.



The key open-source difference here is that the LLM's can train on the open-source code and documentation. This means that those LLM's already understand ORE code and interfaces, and can build them quickly and easily. By comparison with a proprietary black-box engine the quality of any LLM interface is limited by the transparency any vendor is willing to give to the 3rd party model and may be more limited where transparency into underlying models and documentation is restricted as a true open-source implementation.

The implications for model governance are as significant as the operational ones. When an AI can explain not just the number but the logic — walking a risk manager through why a particular trade contributes disproportionately to CVA, or why a scenario produces an unexpected PFE spike — the gap between quant production and business understanding narrows materially. Regulators who ask institutions to demonstrate comprehension of their own risk models may find AI-augmented explanation a more satisfying answer than a methodology PDF.

What institutions should do now

The transition to open, API-accessible, AI-augmented risk infrastructure is not a distant disruption — it is happening in the current technology cycle. Institutions that are rebuilding for FRTB, re-platforming post-IBOR, or replacing end-of-life legacy systems have a one-time opportunity to adopt architecture that can serve them through the next regulatory and market cycle, rather than one that is likely to require replacement in five years.

The practical steps are more accessible than the magnitude of the shift might suggest. Starting with a specific, bounded use case — intraday CVA monitoring, SA-CCR capital optimisation, or FRTB sensitivity generation — allows teams to validate the ORA model, build internal familiarity with API-driven risk calculation, and demonstrate business value before committing to full portfolio migration.

The open-source foundation matters here. Because the algorithms of ORE are published and peer-reviewed, institutions can validate outputs against their own models without dependency on vendor disclosure. The transition from in-house or legacy-vendor calculation to LSEG's MaaS platform can be managed with the same rigour applied to any model change and the documentation exists to support it.

Conclusion:

The stack is being rebuilt

The quant infrastructure war is unlikely to be won by the institution with just the most sophisticated proprietary algorithms. It almost certainly will be won by the platform that combines open, auditable calculation with cloud-native delivery, institutional-grade data, and AI-powered accessibility. That combination makes risk analytics faster, more transparent, more explainable, and available to a far wider audience within a financial institution than current architectures permit.

Open-source did not win the server infrastructure war in a single release. It won through a sustained accumulation of trust, capability, and community — until the proprietary alternative could no longer justify its cost or its opacity. Quantitative finance infrastructure is at an equivalent inflection point. The institutions that recognise this transition early, and position themselves on the right side of it, could gain a structural advantage in the decade ahead.

About the Author

Stuart Smith is a Director at LSEG Post Trade Solutions, responsible for management of the risk, reconciliation and data businesses. LSEG Post Trade Solutions is the steward of the Open-Source Risk Engine (ORE) and provides risk analytics infrastructure to financial institutions globally through the LSEG Risk Analytics platform.

Contact us

To learn more, contact us at analyticssales@lseg.com or visit our website:
<https://www.lseg.com/en/data-analytics/lseg-analytics-solutions>

Americas +1 646 989 2200

Tokyo +81 3 6441 1015

EMEA +44 20 7334 8963

APAC +886 2 8729 5130

© LSEG 2026. All rights reserved.

London Stock Exchange Group plc and its Group companies ("LSEG") do not guarantee that any information contained in this document is and will remain accurate or that use of the information will ensure correct and faultless operation of the relevant service or equipment. LSEG, its agents and employees, accepts no liability for any loss or damage resulting from reliance on the information contained in this document. This document contains information proprietary to LSEG and may not be reproduced, disclosed, or used in whole or part without the express written permission of LSEG. Any software, including but not limited to, the code, screen, structure, sequence, and organization thereof, and documentation are protected by national copyright laws and international treaty provisions. This document is subject to U.S. and other national export regulations. Nothing in this document is intended, nor does it, alter the legal obligations, responsibilities or relationship between yourself and LSEG as set out in the contract existing between us.

LSEG Data & Analytics is part of London Stock Exchange Group plc Group companies.