



The traded risk architecture of the future

A strategic blueprint
for modernisation

Whitepaper





Content

Executive Summary	04
01. Risk Architecture in Traded Risk	06
02. Typical Challenges Observed in the Market	07
03. Building Blocks and Key Requirements for a Modern Traded Risk Architecture	10
04. Blueprint for a Future-Proof Architecture in Traded Risk	12
05. Takeaways	19
About KPMG	20
About ActiveViam	21

Executive Summary

Traded risk management is undergoing a profound transformation. As financial institutions face increasing regulatory demands, market volatility and technological disruption, the need for resilient, scalable and future-proof risk architectures has never been greater.

This white paper outlines a strategic blueprint for modernising traded risk infrastructure. It explores the motivations behind this shift, identifies common challenges in the market and presents actionable approaches for building robust, AI-enabled and governance-driven architectures.

Key takeaways:

- The imperative to treat technology as a strategic asset, not a cost centre
- The importance of unified data models, centralised computation and semantic consistency
- The role of modularity, scalability and real-time analytics in enabling agile decision-making, powered by AI-enabled insights
- A practical roadmap for transforming legacy systems into coherent, trusted platforms

The Digital Imperative: Transforming Traded Risk Architecture for Sustainable Advantage

The management of traded risk is entering a transformative era, one shaped by accelerating technological disruption, rising regulatory complexity and the growing need for agility in decision-making. For decades, banks have operated under the assumption that their core strength lies in managing capital and client relationships, not in building and maintaining sophisticated IT systems. But this paradigm is no longer sustainable.

Technology today is not just a support function; it is the engine of competitive advantage. The boundaries between banking and technology are dissolving, and the pace of change is relentless. With the rise of generative AI, agentic systems and real-time analytics, the rules of risk management are being rewritten. The uncomfortable truth is that no institution, whether bank, regulator or vendor, can predict with certainty how the landscape will look in five years, or even next quarter.

To stay competitive in this new reality, banks must evolve from being passive consumers of technology to active architects of their digital future. The old playbook, with incremental upgrades, tactical fixes and reliance on legacy systems, is no longer sufficient. Instead, institutions must embrace a radical rethink of their technology landscape and organisational mindset.

This transformation requires:

- Reframing risk IT as a strategic asset and not a cost centre
- Investing in resilient, scalable and future-proof architectures that support real-time analytics, seamless integration and rapid innovation
- Fostering cross-functional collaboration between IT and business teams to break down silos and accelerate value creation
- Promoting a culture of experimentation and continuous learning, where agility and insight are prioritised over tradition and inertia



1. Risk Architecture in Traded Risk

Traded risk encompasses the full spectrum of financial exposures arising from a bank's trading activities, in particular market risk, counterparty credit risk and valuation risk. These risks are inherently dynamic, complex and subject to intense regulatory scrutiny. Managing them effectively requires more than isolated technological solutions; it demands a coherent, multi-layered architecture that integrates governance, data, analytics and process design into a unified framework.

Unlike traditional credit risk, traded risk is defined by its immediacy and volatility. Exposures can change in milliseconds, across instruments ranging from plain vanilla swaps to exotic derivatives. The boundaries of traded risk are fluid, spanning front office innovation, risk management discipline and regulatory oversight. More than a mere technical domain, it is the pulse of modern banking, where strategic decisions and operational resilience converge.

Functional Scope of Traded Risk Architecture

A modern traded risk architecture must support a wide range of business functions, each with its distinct processes and regulatory implications:

Figure 1: Traded Risk: Business scope/major processes



Source: KPMG in Germany, 2025

¹ According to the KPMG risk function business catalogue; number depends on the product portfolio and the individual setup of the respective bank

Strategic Importance

In today's environment, the notion of risk architecture being a back office concern is dangerously outdated. A modern traded risk architecture is the difference between agility and paralysis, between regulatory compliance and costly sanctions, between actionable


insight and strategic blindness. It empowers banks to measure, control and optimise risk in a world where changes happen rapidly.

2. Typical Challenges Observed in the Market

Despite years of investment and regulatory scrutiny, the reality of traded risk architecture in most banks presents a sobering picture. The industry faces a paradox: while external forces like digitalisation, new business models and AI accelerate transformation, many institutions continue to wrestle with long-standing structural issues. Complexity, fragmentation

and inertia are not new problems; they are long-standing liabilities that have been tolerated, patched and postponed. What's different now is not the diagnosis, but the urgency. The industry is being forced to revisit old challenges with new eyes, and this time the stakes are existential.

Figure 2: Trade Risk Architecture: Observations in the market



Source: KPMG in Germany, 2025

² EUC = End-User-Computing or EUDA = End-User Developed Application

Data Management and Data Quality: The Achilles Heel

At the core of nearly every challenge in traded risk lies the issue of data. Banks operate with a sprawling landscape of market, reference, position and transaction data, sourced from a multitude of systems, vendors and business units. This diversity, while inevitable, is rarely harmonised. Intraday and end-of-day data are often misaligned; counterparty credit risk and minimum capital requirements rely on inconsistent and incompatible data sources; and data silos persist across departments and platforms.

The consequences are systemic: inconsistent models, duplicated data provisioning and endless reconciliation efforts that drain resources and undermine trust. Bitemporal data management, critical for handling corporate actions, restatements and historical corrections, is either neglected or implemented in fragments. The result is a brittle foundation where even the most advanced analytics engines are rendered ineffective.

Let's be clear: data quality and control are not optional. Without them, risk management frameworks lose their foundation.

System Complexity and Technical Debt: A Self-Inflicted Wound

Most banks' risk architectures are the result of years of incremental change, tactical fixes and project-driven development. The outcome is a patchwork of legacy platforms, parallel pricing libraries and bespoke solutions, each with its own quirks and dependencies.

There was one example of a global bank maintaining multiple separate pricing libraries for different asset classes, each developed by different teams. When a regulatory change arose that required consistent valuation logic across desks, the bank was faced with costly and time-consuming integration work.

This complexity is more than a technical nuisance; it's a strategic liability. Fragile integrations, high maintenance costs and reliance on a few experts create operational risks that are rarely acknowledged but widely felt. Shadow IT and end-user computing, especially Excel spreadsheets, further exacerbate the problem, undermining governance and increasing exposure.

The unfortunate truth is that many institutions spend millions on regulatory compliance but fail to address root causes, in particular fragmented data, technical debt and lack of architectural vision.

Reporting and Analytics: The Illusion of Control

Regulatory expectations for risk reporting have never been higher. Banks must deliver consistent, timely and auditable reports across jurisdictions and business units. However, they often still have divergent reporting logics, make manual adjustments and lack real-time analytics.

The inability to provide a single, consistent view of risk is not just a compliance issue; it's a business risk. Decision-makers often rely on outdated or incomplete information, limiting their ability to respond to market events or regulatory demands with agility and confidence.

Infrastructure and Scalability: The Limits of Legacy

Legacy systems are increasingly unable to meet the demands of modern risk management. Volatile markets, high data volumes and the need for real-time analytics expose the limitations of batch-based processing and rigid architectures.

In practice, this means that ad hoc analysis, so critical for decision-making under uncertainty, is still largely performed in Excel. Risk managers rely on spreadsheets and manual data extracts because



existing platforms lack the responsiveness, transparency and flexibility needed for exploratory analytics. This hinders insight generation and also introduces operational risk through versioning issues, lack of auditability and inconsistent logic.

While cloud technology offers a path forward, adoption is often hampered by concerns over compliance, security and integration. Many banks remain stuck in a hybrid limbo, neither fully modern nor reliably legacy.

Security, Governance, and Auditability: The Unseen Risks

As data volumes grow and architectures become more complex, the challenges of security and governance multiply. Many institutions struggle to implement robust data lineage (i.e. the ability to trace the origin, transformation and usage of data across systems), access controls and encryption across sprawling IT landscapes.

Auditability is often an afterthought, bolted on rather than built in. This increases the risk of regulatory sanctions and undermines the bank's ability to respond confidently to audits and inquiries.


Organisational Inertia: The Human Factor

Technology alone cannot drive transformation. Siloed thinking, resistance to change and insufficient collaboration between IT and business functions are persistent barriers. The traditional mindset, viewing IT as a cost centre, remains deeply entrenched.

The hard thing to accept is that the biggest obstacle to transformation is culture, not technology. Banks that fail to break down silos, empower cross-functional teams and embrace continuous learning will be outpaced by more agile competitors.

In Summary

The challenges facing traded risk architecture are well known and solvable. But they require more than incremental change. They demand a fundamental rethink of data, systems, processes and culture. The winners will be those who confront these challenges head-on, invest in resilient architectures and embrace technology as the crucial factor for future success.



3. Building Blocks and Key Requirements for a Modern Traded Risk Architecture

Modernising traded risk architecture requires more than just replacing legacy systems. It necessitates a fundamental rethinking of how risk is measured, managed and governed. This chapter outlines the essential components and guiding principles that define a future-ready risk architecture.

Data Consistency and Semantic Transparency

To eliminate fragmentation, institutions must establish a unified, version-controlled data foundation that ensures:

- A single source of truth for every point in time
- Historical accuracy with correction capabilities
- Consistent definitions for all metrics, dimensions and relationships to enable explainability
- Full traceability and auditability of data sources to meet regulatory and internal governance standards while supporting robust analysis

This foundation guarantees that all analytics and reporting are built on reliable, consistent data.

Scalable Infrastructure and Elastic Performance

Architectures must be designed to scale dynamically and adapt to fluctuating workloads:

- Elastic compute and storage capacity for stress testing, ad hoc analysis during volatile market conditions, and periodic regulatory reporting
- Cost-efficient resource management that releases unused capacity
- Predictable performance even under high load

Modular Expansion and Seamless Integration

To remain agile, architectures must support incremental growth and integration:

- Plug-and-play components for new business functions, models and analytics

- Standardised APIs and messaging frameworks to simplify integration with vendor solutions and reduce onboarding costs for developers
- High observability across distributed environments

This enables institutions to evolve without disrupting core operations.

Real-Time Analytics and Limit Monitoring

Effective traded risk management demands intraday visibility and scenario analysis across positions, exposures and P&L:

- On-the-fly recalculations using dependency graphs to optimise performance (e.g. for xVA pricing)
- Granular drill-down capabilities to trace data variations to their root causes
- Automated alerts for threshold breaches, securely routed to dashboards and mobile devices

AI and Machine Learning Capabilities

AI is no longer optional; it's a strategic imperative. A modern architecture must:

- Enable AI agents to interact with structured, governed data environments
- Leverage semantic layers for explainable logic and predictive modelling
- Support open MCP³-based platforms for secure integration of preferred LLMs
- Automate data quality validation and anomaly detection with real-time dashboards and workflows

³ See Agentic MCP (Model Context Protocol) Server protocol specification for open AI architecture

This empowers institutions to intelligently scale analytics and scenario generation.

Security, Governance and Compliance

Security must be embedded into the architecture, not added as an afterthought:

- Enterprise-grade access control and encryption across all layers
- Integration with identity and access management systems (LDAP⁴, AD⁵, SSO⁶) for consistent policy enforcement
- Compliance with regulatory standards like General Data Protection Regulation, BCBS239 and emerging AI governance standards

Auditability and traceability must be integral to every component.

Operational Stability and Simplified Maintenance

Reducing operational risk and technical overhead is critical:

- Eliminate shadow IT by replacing spreadsheets with governed, enterprise-grade tools
- Establish centralised and highly automated e.g. for daily risk and P&L sign-offs across global teams
- Establish continuous Integration (CI) and Continuous Delivery/Deployment (CD) pipelines and configuration-as-code for versioning, testing and deployment

This streamlines operations and enhances reliability.

Summary

Strategic transformation of traded risk architecture requires more than technical upgrades; it demands a shift in mindset, design philosophy and operational discipline. A modern traded risk architecture is not just a platform; it is the foundation for sustainable success.


⁴ LDAP: Lightweight Directory Access Protocol

⁵ AD: Active Directory

⁶ SSO: Single sign-on

4. Blueprint for a Future-Proof Architecture in Traded Risk

Figure 3: Traded Risk Architecture of the future



Source: KPMG in Germany, 2025

Architecture of the future

A traded risk IT architecture is designed to support a wide range of risk management functions (valuation, market risk and counterparty credit risk) by integrating diverse data sources, centralised valuation logic and advanced analytics capabilities into a coherent, governed platform.

At its foundation, the architecture ingests and harmonises data from multiple upstream systems. These include front office systems, which deliver trade and transaction data; market data feeds providing prices, curves and volatilities; and reference data sources such as counterparty and issuer information, contract terms, portfolio structures and other static data. This data is processed and consolidated in a dedicated data layer, where it is made consistent, traceable and ready for downstream consumption.

On top of this harmonised data pool, the architecture supports a suite of business functions that rely on a shared valuation and simulation engine. These functions include end-of-day valuation, independent

price verification (IPV), market conformity checks, sensitivity calculations, value at risk (VaR) computation and exposure simulations for counterparty credit risk and regulatory capital.

Each of these functions accesses the same data foundation and leverages a centralised calculation core, ensuring consistency in logic and results across the organisation.

Beyond the core risk computations, the architecture provides robust capabilities for reporting, analysis and monitoring. This includes regulatory and internal reporting, limit monitoring and breach detection, ad hoc analysis tools for risk managers and business users, and sandbox environments for AI/ML experimentation and scenario analysis. These analytical layers are tightly integrated with the semantic layer and governed by role-based access controls, ensuring that insights are consistent, secure and aligned with enterprise standards.

In the following subchapters, we outline the key architectural design decisions that guide the development of a streamlined traded risk IT platform and which are essential for fulfilling the requirements described above. This framework is built to deliver a single version of truth across data, computation and reporting, ensuring operational efficiency, regulatory compliance and long-term scalability. Each section provides a practical perspective on how to achieve consistency, transparency and performance across the entire risk value chain.

From Vision to Execution: Building the Foundation

At the heart of a robust traded risk architecture lies a well-governed, transparent and scalable data model. This foundation must support the technical demands of valuation and simulation as well as the regulatory requirements for traceability, auditability and long-term retention.

The first ingredient in this architecture is data lineage. From the outset, the system must be capable of tracking the complete journey of data, from its origin in source systems, through transformation, to its final use in risk calculations and reports. This lineage must be both technical, showing how data flows and is processed, and business-oriented, explaining why specific data was used in a particular context. For example, it should be possible to trace why a certain sensitivity was assigned to a specific FRTB bucket or which market data feed was used in a valuation.


This traceability is not just a nice-to-have; it is essential for regulatory audits, model validation and (especially)

daily business analysis. It also enables automated data quality checks and supports correction workflows with full audit trails, ensuring that any data issue can be quickly identified, understood and resolved.

A bitemporal data model, capturing both business time and system time, is essential for reconstructing historical states, supporting intraday snapshots and enabling retrospective or what-if analysis. This means capturing both the time when the data was valid (its business time) and the time when it was recorded (its system time). This dual-timestamp approach enables institutions to reconstruct past system states with precision, generate consistent intraday snapshots across multiple front office systems and provide reliable data slices for AI/ML sandboxing or any risk computation.

Data must be logically separated by type: input data (e.g. trades, market feeds), transformed data (e.g. sensitivities, intermediate calculations), and result data (e.g. VaR, limit utilisations). This tiered approach enables institutions to optimise storage, control costs and align retention policies with business value.

Furthermore, the architecture must be designed to scale horizontally. As data volumes grow and more users access the system concurrently, performance must remain stable. To meet long-term retention requirements - often up to 10 years - the platform should implement tiered storage, storing frequently accessed data on high-performance infrastructure and archiving historical data on cost-efficient platforms.





Smart deletion policies should be in place to ensure that data is only removed when it is no longer needed for business or regulatory purposes. Efficient storage formats combined with intelligent partitioning strategies help optimise performance and reduce infrastructure costs.

Governance is equally critical. Every data access and modification must be timestamped, versioned and auditable. Role-based access control ensures that users interact only with data relevant to their responsibilities, while integration with enterprise IAM platforms (LDAP, AD, SSO) enforces consistency across systems. These controls are not optional; they are the backbone of regulatory compliance and internal accountability.

Establishing a Single Version of Truth

Once the data foundation is in place, the next essential ingredient is consistency in logic and interpretation. This means ensuring that all business functions - whether regulatory, analytical or operational - rely on the same definitions, models and computational logic. Without this consistency, institutions face a proliferation of conflicting results, duplicated efforts and increased operational risk.

To avoid these pitfalls, the architecture must be designed around centralised components that are reused across all risk functions. This includes a shared calculation engine, a unified semantic layer and a governance framework that ensures alignment across the entire organisation.

At the heart of the architecture lies a centralised calculation core - a robust, scalable engine that encapsulates all valuation logic, simulation models and stochastic components. This core is a strategic necessity, not just a technical convenience.

For example, the same valuation models used for end-of-day pricing should also be used for counterparty credit risk exposure simulations and Pillar 2 stress testing under the minimum capital requirements framework. This ensures that risk figures are consistent across different regulatory and business contexts.

Similarly, the Monte Carlo simulation framework, used to model exposure profiles and risk factor evolution, should be centrally maintained and reused across modules. This avoids discrepancies in simulation results and simplifies model validation.

By consolidating these computational components, institutions reduce model risk, eliminate duplication, and streamline validation and audit processes.

The result is a coherent, trusted risk infrastructure that supports both regulatory compliance and internal decision-making.

While the calculation core does the heavy lifting, the semantic layer provides the business-facing abstraction. It acts as a central repository for all definitions, metrics and dimensional hierarchies, ensuring that everyone, from analysts to regulators, speaks the same language.

This layer defines key metrics such as:

- Value at Risk (VaR)
- Exposure at default (EAD)
- Aggregated sensitivities, such as the sum of credit spread sensitivities

It also maintains dimensional hierarchies, such as risk buckets, trading desks and time horizons, and encodes business logic used in reporting, analysis and regulatory submissions.

By exposing these definitions to all downstream systems, especially to reporting tools, dashboards, limit monitoring systems and AI/ML sandboxes, the semantic layer ensures that all outputs are based on consistent and validated logic. This eliminates the need for manual reconciliation and reduces the risk of conflicting interpretations.

Enabling Strategic Insight and Regulatory Transparency

In a modern traded risk IT architecture, analytics and reporting are far more than technical outputs; they serve as strategic enablers. They provide the visibility needed for risk transparency, support ad hoc decision-making and fulfil regulatory obligations. However, without a structured and governed approach, analytics can quickly become fragmented, inconsistent and inefficient.

To avoid this, analytics must be embedded in an architecture that is governance-driven, scalable and reusable. It must align with the broader principles of centralisation, traceability and performance, and be designed to support both standardised reporting and flexible exploration/ad hoc analysis, including AI/ML sandboxing and stress testing.

A robust governance framework is the backbone of trustworthy analytics. It begins with a clear separation of roles: defining who is authorised to publish analytical content and who may consume it. This distinction ensures that only validated, approved reports and dashboards are made available to business users and regulators.

Publishing new metrics or dashboards should follow a structured approval workflow, involving validation by subject matter experts and sign-off by data owners. This prevents the uncontrolled proliferation of reports and ensures that all published content meets quality and compliance standards.

Access rights must be inherited from the underlying data platform, ensuring consistency across the reporting tools and databases. This alignment simplifies access management and strengthens auditability, as every action is traceable to a defined role and user.

Defending the Single Version of Truth - Consistency Across All Outputs

Business intelligence of analytics tools must be closely linked with the central semantic layer, which defines all business metrics and dimensions. This ensures that:

- Definitions are consistent across all reports and tools
- Data is not duplicated or manually reprocessed
- Every analytical output can be traced back to its source logic and data

This architectural discipline prevents the “wild growth” of inconsistent analyses, conflicting results and divergent layouts that often emerge when teams build reports independently. It also reinforces trust in the data, as users know that every number is based on validated, governed logic.

Performance and Scalability – Supporting High-Volume, Multi-User Environments

Traded risk analytics must be designed to handle large data volumes, including intraday sensitivities, simulation outputs and time series. The platform must support concurrent access by multiple users and systems without compromising performance.

To achieve this, the architecture should leverage:

- Semantic caching to reduce query load
- Materialised views for frequently accessed metrics
- Partitioning strategies to optimise data access

Scalability ensures that performance remains stable even as data volumes grow and analytical demands increase - whether from regulatory teams, trading desks or AI/ML models.

Auditability and Traceability - Ensuring Transparency and Compliance

Every analytical output must be versioned, enabling institutions to reproduce past reports and roll back changes when needed. Reports and dashboards should be linked to their semantic definitions and data lineage, ensuring full transparency of logic and data sources.

All user interactions (such as report generation, metric changes and data access) must be logged and monitored. This supports internal controls and enables institutions to respond confidently to regulatory audits and inquiries.

Reuse and Modularity - Building Efficient and Maintainable Analytics

Analytics should be constructed using modular components that can be reused across departments and use cases. This includes:

- Reusable metrics such as VaR, expected shortfall and limit utilisation,
- Standardised visual templates to ensure consistent layout and interpretation
- Parameterised dashboards that adapt to different portfolios, desks or time horizons

This modularity reduces development effort, simplifies maintenance and ensures that analytics remain consistent and scalable as business needs evolve.

Self-Service Enablement - Empowering Users Within Governed Boundaries

Governed self-service analytics empower users to explore data and generate insights without compromising control. Users can access validated semantic definitions, build custom views and perform ad hoc analysis, all within a framework that enforces data integrity and usage policies.

Training and documentation are essential to support correct usage and encourage adoption. This balance between flexibility and governance enables faster decision-making while maintaining consistency and compliance.

AI and ML Capabilities

AI deserves special attention because it is new and evolving quickly. Financial institutions are in a pivotal moment: generative AI is maturing rapidly, offering exciting innovation and immense potential. But there are legitimate concerns about reliability, governance, IP protection and reputational risk.

The real breakthrough won't come from better models alone - it will come from enabling AI agents to interact with secure, governed and high-performance analytical environments with respected and trusted partners. AI agents need to interact directly to explore, drill down and analyse metrics dynamically - like a human analyst, but faster and without errors.

Semantic layers are crucial to enabling this because they provide structured context. AI thrives with clear structure and metadata. A semantic layer provides dimension names, measure definitions, hierarchical paths and data types and relationships. It also reduces ambiguity; for example, the AI won't have to guess what "revenue" or "active users" means because those terms are clearly defined in a semantic layer. Lastly, it facilitates predictive modelling, because AI can use hierarchical data for more context-aware predictions such as forecasting P&L by region and product hierarchy.

Open MCP-Based⁷ Architecture Enabling Seamless Integration of Preferred LLM and AI Agents

An open and extensible platform based on Agentix's MCP⁷ server protocol empowers AI agents that can reason, explain and act on financial data with confidence. As firms establish their AI policies, they will adopt a preferred (or even mandatory) large language model (LLM). Control over the LLM is essential for firms concerned about reliability, governance, IP and reputation risk. Firms will increasingly demand that AI-enabled architectures are open so that the chosen LLM can be integrated.

⁷ See Agentix MCP (Model Context Protocol) Server protocol specification for open AI architecture

Predictive Analytics, Scenario Generation and Model Explainability

AI-powered tools are improving the quality of predictive analytics and doing it at a lower cost. AI agents can synthesise historical data to generate plausible forward-looking scenarios. These intelligent systems can adapt to emerging risks and can also optimise scenario selection by focusing computational resources on high-impact, low-correlation cases while eliminating redundant or low-value simulations. This targeted approach significantly reduces compute time and infrastructure cost, enabling institutions to run more frequent, more granular and more relevant analytics.

Automatic Explainers, Data Quality Validation Built-In

AI-powered tools have an important role for data quality validation and automatic explainers. Data errors can slip into even the best infrastructure, and they are very difficult and time-consuming to find and fix manually. In this context, AI-powered tools can automatically check data for anomalies against a set of fixed and fuzzy rules.

Integrating this process with real-time data quality dashboards and an operational workflow where flagged data is reviewed and corrected is intrinsic to this capability. The anomaly detection piece alone is not a complete solution to the data quality problem.

Organisational Readiness: The Human Architecture

Technology alone cannot deliver transformation. Success depends on cultural alignment and cross-functional collaboration. Historically, IT and business teams have operated in silos, leading to misaligned priorities and duplicated efforts. A future-proof architecture demands a new way of working, one that fosters shared ownership, agile development and continuous learning. This requires:

- Empowering cross-functional teams
- Encouraging experimentation and learning
- Aligning budgets, incentives and KPIs across departments
- Adopting tools and platforms that foster the above

Reducing shadow IT is a key step. When business units build tools outside formal governance, they introduce risk and inefficiency. A centralised, well-governed architecture replaces these ad hoc solutions with enterprise-grade workflows, improving stability and reducing reputational risk.

Agile methodologies, CI/CD pipelines and configuration-as-code practices enable rapid deployment, testing and rollback. Sandbox environments enable experimentation with production-grade data, accelerating innovation without compromising compliance.


These practices reduce time-to-value, cut costs and enable rapid problem solving without compromising stability or compliance.

Conclusion: A Coherent, Trusted Traded Risk Platform

The blueprint for a future-proof traded risk architecture is built on three pillars: a governed and scalable data foundation, a unified logic and semantic layer, and an organisational model that supports agility and collaboration.

By aligning these elements, institutions can move beyond fragmented systems and reactive compliance. They can build a platform that is resilient under pressure, transparent across functions and ready for the demands of AI, real-time analytics and evolving regulation.

This is a strategic transformation, not just an IT upgrade. And it begins with a blueprint that turns architectural vision into operational excellence.



5. Takeaways

The journey toward a future-proof traded risk architecture is not defined by a single technology choice or a one-time transformation. It is a strategic evolution touching every layer of the organisation, from data and infrastructure to governance and culture.

This white paper has outlined the imperative for change. Traded risk is becoming more complex, more regulated and more data-intensive. Legacy systems, fragmented processes and siloed thinking are no longer sustainable. Institutions must move beyond tactical fixes and embrace a coherent, resilient and scalable architecture that can support both regulatory compliance and business innovation.

The blueprint we've presented is grounded in practical design principles: centralised data and computation, semantic consistency, modular analytics and elastic infrastructure. But these technical foundations must be matched by organisational readiness. Success depends on cross-functional collaboration, agile development practices and a mindset that views technology not as a cost, but as a strategic asset.

Institutions that align their architecture with these principles will unlock a range of benefits:

- Reduced fragmentation and duplication, leading to cleaner data and more reliable analytics
- Improved compliance and auditability, with full traceability across systems and reports
- Faster, more agile decision-making, powered by real-time analytics and AI-enabled insights
- Lower operational costs, through modular design, functional reuse and elastic infrastructure
- Greater adaptability, enabling institutions to respond confidently to regulatory changes and market shifts

Ultimately, the goal is not just to modernise IT systems but also to build a trusted, transparent and future-ready platform for traded risk. One that empowers risk managers, satisfies regulators and supports strategic growth.

The institutions that succeed will be those that treat architecture as a living capability, one that evolves with the business, scales with demand and adapts to the unknown. In a world where technology is the decisive factor, traded risk architecture is no longer a mere office concern. It is a boardroom priority.

Let's finish with arguably the most important insight: transformation doesn't require a mega-project. Change can and should begin with small, focused initiatives. Pilot projects and incremental improvements often deliver tangible benefits quickly, building momentum and trust across the organisation.

- Agile transformation enables institutions to adapt as they learn, avoiding the pitfalls of over-planning and under-delivering
- Early wins from small projects demonstrate value, reduce resistance and pave the way for broader adoption
- A clear roadmap is essential, but flexibility is key: organisations must be willing to adjust course as needs evolve

About KPMG

In navigating the intricate landscape of risk and compliance, organisations increasingly require robust risk management frameworks. KPMG banking risk professionals stand ready to assist, drawing upon extensive experience and technical capabilities to help you overcome multifaceted challenges. Whether grappling with regulatory complexities or evolving technology threats, KPMG specialists offer a suite of services to help fortify risk management practices and cultivate trust. KPMG banking risk professionals merge their experience in risk with transformative insights to offer strategies that can shield against uncertainties while also illuminating new pathways to resilience, growth and enhanced stakeholder trust. Leveraging powerful risk analytics, advanced modelling techniques and real-time risk reporting, KPMG professionals empower banks to integrate risk management into their daily operations. Through innovative services, KPMG professionals enable clients to proactively address emerging risks and seize opportunities for value creation.

Lukas Henatsch

Director, Financial Services
T +49 211 475-6386
lhenatsch@kpmg.com

Franz Lorenz

Director, Financial Services
T +49 89 9282-4542
florenz@kpmg.com

Dr. Barbara Götz

Senior Manager, Financial Services
T +49 89 9282-6798
barbaragoetz@kpmg.com

Dr. Marcel Kraus

Senior Manager, Financial Services
T +49 69 9587-1733
marcelkraus@kpmg.com

About ActiveViam

ActiveViam is a fast-growing financial data analytics solution provider. Built for and trusted by leading financial institutions, ActiveViam's flagship product Atoti delivers active intelligence for complex financial analytics. It combines unrivalled technology, continuous innovation and exceptional people to unlock the power of real-time and granular data at scale. Designed as a high-performance semantic layer, ActiveViam's Atoti enables clients to implement built-in front office and risk business solutions while accessing customisable technology.

ActiveViam is present in the world's leading financial marketplaces with offices in London, New York, Singapore, Sydney, Hong Kong, Paris and Frankfurt.

For more information, please visit activeviam.com or follow on [LinkedIn](#).

Florence Falck

Global Head of Partnerships
T +44 7769 216 031
ffa@activeviam.com

Chris Horril

Director of Product Marketing
cho@activeviam.com

Dominique Vigneaux

Key Account Manager
T +33 6 19 85 62 00
dvi@activeviam.com

Sarah Weeks

Global Head of Marketing
swe@activeviam.com



www.kpmg.de

www.kpmg.de/socialmedia



The information contained herein is of a general nature and is not intended to address the circumstances of any particular individual or entity. Although we endeavor to provide accurate and timely information, there can be no guarantee that such information is accurate as of the date it is received or that it will continue to be accurate in the future. No one should act on such information without appropriate professional advice after a thorough examination of the particular situation.

© 2025 KPMG AG Wirtschaftsprüfungsgesellschaft, a corporation under German law and a member firm of the KPMG global organization of independent member firms affiliated with KPMG International Limited, a private English company limited by guarantee. All rights reserved. The KPMG name and logo are trademarks used under license by the independent member firms of the KPMG global organization.