



A Platform for Effective Risk Management

Increasing market risk managers' effectiveness

This white paper outlines a systems architecture for market risk managers that provides for speed, scale and operational integrity, without prejudice to its capacity for flexibility and change.



XENOMORPH®

Table of contents

Introduction.....	3
IT Requirements of a Typical Market Risk Management.....	4
Key questions risk managers need to answer	4
Answering the questions.....	4
Using the computation results	5
Summary and review.....	7
Key Elements of the Market Risk IT Architecture.....	8
Market Reference Data Database.....	8
Historic Market Data Database	8
Market Data Update Service.....	9
OTC Reference Data Database	9
Position Database.....	10
Market Data Scenario Generation Service.....	10
Valuation Model Library.....	10
Pricing Service	10
Risk Calculation Service	11
Risk Management Database	11
Reporting Service	12
Task Orchestration Service	12
Implications for Providers of Risk Management IT	13

Introduction

Effective risk management is impossible without effective information technology. The IT industry has responded accordingly with risk management conferences crowded with salesmen demonstrating risk management systems.

But are these technologies truly effective? For many real-life risk managers, whose job it is to identify and to quantify the rapidly changing risks facing their companies, risk management systems are not facilitators but are at best a constraint, and are at worst a cause for concern. For many trading desks developing new business, the inflexibility of risk management systems is one of their greatest impediments. Given the billions of dollars of recent investment in risk management information technology, these results are disturbing.

The purpose of this paper is to suggest that risk management information technology has become divorced from risk managers and their daily needs, whose mundane reality is often a million miles away from the 'advanced risk analytics and scenario capabilities' which risk management IT vendors may seek to deliver. This rift is the most important reason for the IT profession's failure to deliver value.

We seek to heal the rift by re-stating the real technology requirements of risk managers and by outlining the key features that risk IT must provide if it is to be effective. We will then compare these features with well-known risk management offerings.

The approach proposed by this paper is based on three principles.

Breadth. Our scope spans all the functions of a market risk management group without exception. Breadth of coverage is crucial to the effectiveness of any risk management IT system. Otherwise, risk managers have to expend a disproportionate amount of valuable manpower on relatively unskilled activities – such as investigating and reporting instrument or reference data problems - activities that are not addressed by the sophisticated risk analytics and scenario functionality provided within vendor risk management systems.

Use of tight and/or loose coupling. One the biggest sources of tension in the risk management IT world is the conflict between the need traders and risk managers have for innovation and flexibility and their equally-important requirement for mature, well-tested risk systems with rich audit trails. Risk managers require an architecture in which inputs and outputs can be tightly-coupled where appropriate (e.g. when interfacing to market data sources or transaction systems which are stable and permanent) or loosely-coupled if required (e.g. when providing ad hoc computational results to risk managers via Excel functions).

Promiscuity. No real-life risk manager can solve all his or her problems with a single data vendor or a single analytics provider: they require the capability to be flexible, pragmatic and promiscuous. The "glue" with which they integrate different data sources and different analytical tools under a unified umbrella is therefore just as important to them as the individual components.

Having set out the three guiding principles, we will now explore the IT needs of a typical market risk management group, the key elements required to meet those needs and how mainstream risk management solutions match up.

IT Requirements of a Typical Market Risk Management

Key questions risk managers need to answer

Essentially the task of a risk management group is to answer these three questions:

- **What is my trading position?** What positions and instruments are in my trading position and what are the key position measures – be they notionals, PV01's, bucketed vegas, or others that describe them?
- **How is my day-on-day P&L moving and why?** What market factors are driving the market value changes in my trading book?
- **What could happen to my position and P&L if markets moved differently in the future?** What would be the effect of large or small moves in different markets on my trading book?

Risk managers need to be able to answer these questions using IT systems that are sufficiently automated and integrated so as to be robust, cost-effective, auditable and reliable. But, at the same time, risk managers need to retain sufficient adaptability and flexibility to continue to provide the answers as their organisations' trading activities, circumstances, or market environments change.

Answering the questions

In order to answer these questions, risk management systems typically generate:

- Trading position reports
- P&L attribution back to market movements
- Value at risk calculations
- Value at risk back-testing against realised P&L
- Stress testing results

The data sources and processing methods required to support these are illustrated in Table 1 overleaf.

Table 1. Information Processed by Risk Management Systems

Information Generated	Data and Processing Required
Trading position reports	<ul style="list-style-type: none"> • Positions and traded instruments • Instrument or security static reference data • Trading portfolios and their reporting hierarchy in the firm • Position measures such as notional amounts, trade counts, Greeks, etc • Integrity checks on transaction universe and accuracy; exception follow-up
P&L attribution back to market movements of different kinds	<ul style="list-style-type: none"> • Current and previous close of business market data datasets • A valuation model library • All the position, instrument and portfolio data identified above as required for position reporting • Pricing service workflow to revalue the portfolios systematically and record the results in line with market data changes day-on-day
Value at risk calculations	<ul style="list-style-type: none"> • Reliable, clean historic market data on any significant risk factors identified by the P&L attribution process • A scenario generation service to perturb the most recent close of business market dataset into a set of random or historic test scenarios • The valuation model library, as identified above • The position, instrument and portfolio data identified above • A pricing service which can combine a perturbed market data dataset with a set of transaction data to generate a hypothetical portfolio P&L figure for each perturbed market dataset and to store results in a database • A risk calculation service to extract and manipulate calculation results – such as 99th percentile losses – and store the results of its calculations
Value at risk back-testing against realised P&L	<ul style="list-style-type: none"> • Historic data on realised P&L & VaR, from the two previous steps
Stress testing results	<ul style="list-style-type: none"> • A tool to perturb market data datasets with absolute or relative shifts and to save the results • Historic market data, valuation library and pricing service as identified above, together with the position and portfolio data identified earlier

Using the computation results

The computational results above are only a beginning. In their search to obtain value from the results, risk managers undertake multiple data output, investigation and reporting tasks: ad hoc queries, formal report-writing and distribution, data production or download to other systems and exception processing.

Examples of these categories of activity are illustrated in Table 2 and are not, of course, exhaustive and the range of potential additional requirements is vast. No off-the-shelf risk management system can guarantee to directly provide the variety of reports that its users might use in the future.

Rather than depending on a single system to satisfy their needs, risk managers therefore require the capacity to satisfy their own reporting requirements using standard tools and programming interfaces such as Excel, Access, Crystal Reports, Visual Basic, Java, C#, etc. Given that many risk managers are not experts in technology, the most important by far of these tools is Excel.

The risk management data described in the previous section of the paper, together with the results of data enrichment activities such as user comments, must therefore be held in an open format with an API for all of these tools and interfaces.

Table 2. Output Requirements for Risk Management Systems

Category	Example Output or Activity	Delivery Methods Typically Required
Ad hoc queries	<ul style="list-style-type: none"> History of VaR or daily P&L or VaR limit for a portfolio over previous month Comparison of interest rate deltas on a portfolio between two dates Position listing for a portfolio with the credit spread delta for each position List of trading books and host systems contained in a trading portfolio 	<ul style="list-style-type: none"> Pre-prepared queries running directly into Excel or Access, or another informal tool, on-demand from risk managers Other ad hoc queries built as required, with data running into similar tools
Formal Reports	<ul style="list-style-type: none"> VaR and daily P&L for the most recent two business days for all trading portfolios at the top two levels of the hierarchy P&L attribution report for a portfolio Total position reports for each asset class, run at various levels of the portfolio hierarchy 	<ul style="list-style-type: none"> Formatted report, delivered onto Web server for users to query directly, once data are validated and checked Formatted, saved and mailed out to user list, after validation
Data production and download	<ul style="list-style-type: none"> Flat-file download of risk statistics for other corporate systems 	<ul style="list-style-type: none"> FTP of flat file or other reporting format, scheduled overnight
Exception processing	<ul style="list-style-type: none"> Review list of significant changes in VaR or Greeks, by portfolio. Verify that data are correct and release formal reports, or correct and re-run Review list of limit breaches, investigate, comment, and release formal report. 	<ul style="list-style-type: none"> Web-based or Excel-based exception delivery, with facility for users to save comments back to the reporting database if they have appropriate permissions Report generation itself as outlined above under "formal reports"

Summary and review

The requirements above may appear rather slight, so what problems lie under the surface?

Typical sources of difficulty fall into four areas:

- **Scope.** Solutions to parts of the risk management requirement, such as calculating VaR, ignore other parts of the problem - market data cleaning, static data update, exception reporting, commentary and ad hoc queries to name but a few - which occupy surprisingly large portions of risk managers' time.
- **Valuation model diversity.** Solutions that appear feasible for known products today tend to break down when valuation models for new products are introduced.
- **Scale.** Solutions that appear feasible for portfolios with limited numbers of instruments tend to break down when many thousands or hundreds of thousands are to be processed.
- **Market data modelling.** Although valuation models for new products may be available, the market data associated with them cannot be collected and processed to feed the valuation models.

The data storage, data maintenance and computational elements required to fulfil these needs are described in the next section.

Key Elements of the Market Risk IT Architecture

Key elements of the risk IT architecture that addresses the requirements of the proceeding sections are illustrated below:

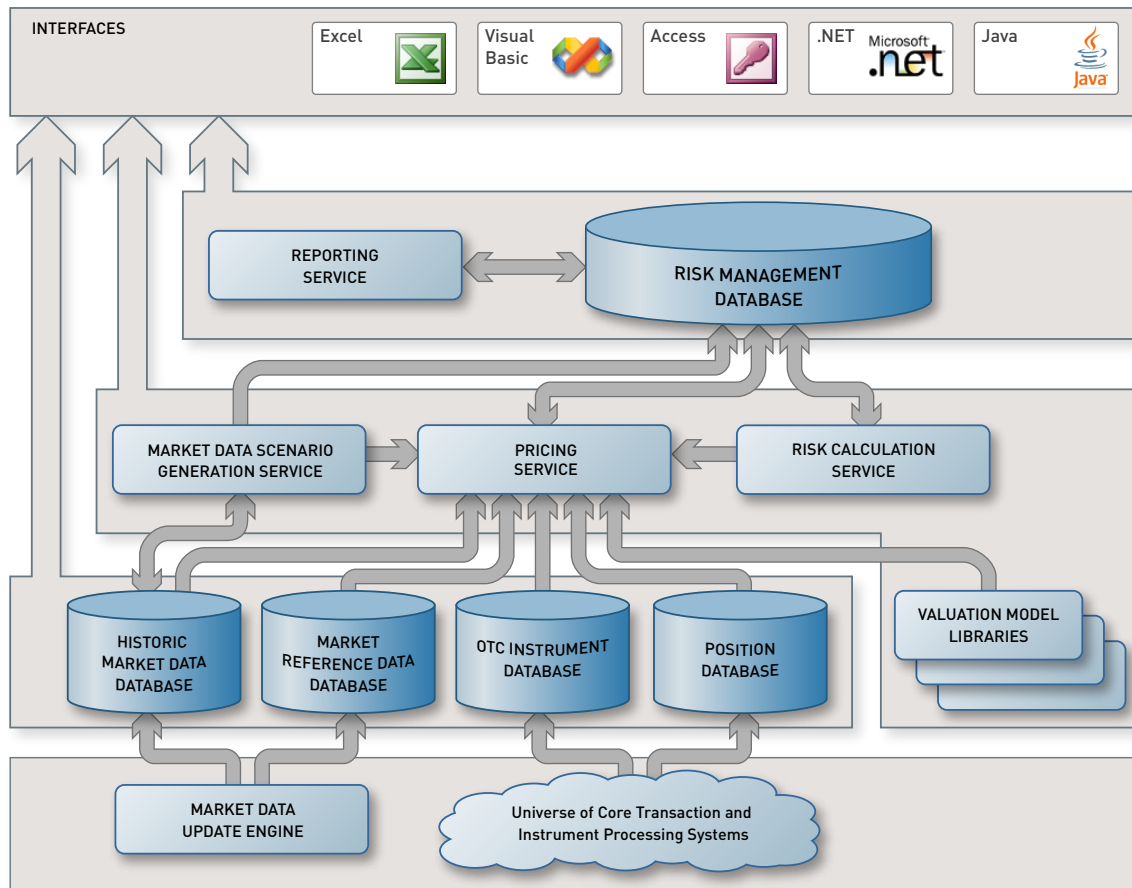


Figure 1. Market Risk IT Architecture Overview

Market Reference Data Database

This is a reference database for market-derived static data and needs to cover items such as: security IDs and their terms; exchange traded futures and options and their terms; equity or credit derivative indices and their components; issuers, ratings and rating sources; industry sectors; currencies with their quotation conventions and codes; country names and their codes; etc.

Historic Market Data Database

This is a database with two types of historic data: data captured directly from external systems and data derived internally via secondary calculations.

Typical examples of captured data are: historic prices and rates for traded products, vanilla or exotic, as sourced from the major data vendors; prices and rates extracted from the trading room's internal records on a daily basis; traded volumes and open interest figures; bid/offer

spreads; etc. The same data items would almost always be captured from multiple sources for verification, data cleaning and validation purposes.

Examples of derived market data are parameters such as implied volatilities and implied correlations or bond credit spreads, generally corresponding to other prices or rates captured directly, but enriched with additional computations. In order to carry out these computations, a pricing service, described later in this section, is required.

Market Data Update Service

This service has three functions. First, it updates the market reference data overnight and produces exception reports on data that has changed.

Second, it updates the historic data captured overnight from any relevant internal or external data sources and produces exception reports on data which warrant further investigation, perhaps because their daily movements are too large in relation to their previous volatility, perhaps because they are unchanged or perhaps because different vendors disagree. Significant exception reporting and processing functionality is required at this point in order to facilitate rapid review and, where necessary, correction.

Finally, the market data update service uses the pricing service to update derived market data such as implied volatility and store them back in the historic database, generating any relevant exception reports in the process.

OTC Reference Data Database

The most challenging task of the risk IT architect is to provide a database with the terms of traded OTC instruments. This database must contain sufficient information to feed a valuation model.

The difficulties in the creation of an appropriate instrument database are two-fold. For vanilla products, problems can centre around volume and scale: some institutions may run portfolios of hundreds of thousands of instrument transactions and this volume places strain on any risk management database.

For new and exotic products, issues revolve around how to represent unusual terms and conditions within a traditional database data model. Multiple and changing underlyings, complex put and call schedules, historic matrices and curves are not necessarily easy to represent in a pure relational database model. This problem is exacerbated further by the simple business fact that new and innovative financial products require a constantly changing and adaptable data model of instrument data.

This can lead to the database designer becoming the bottleneck through which all new product development must pass, hitting time to market and profitability. Alternatively, a trading desk may simply decide to avoid integrating new products within the firms risk management framework, increasing operational risk through both over use of spreadsheets and having risk systems that report on a very incomplete view of the firm-wide trading book.

Position Database

In comparison with the instrument and market data reference databases, the position database is relatively simple. It contains positions in every traded instrument, with the relevant trading book details. No other descriptive data are required. If counterparty exposures are of interest, its contents can be disaggregated to carry these positions per trading counterparty. For audit purposes the positions database must be tightly integrated with settlement and accounting systems, so that they can always be reconciled.

Market Data Scenario Generation Service

The scenario generation tool is a calculation tool that produces hypothetical, simulated market data: it can extract a recent market data dataset from the market data database, perturb it in some way, store it for later use in the historic market data database or send it to the pricing service.

The perturbations might be determined in one of three ways. First, they might be generated randomly via a Monte Carlo process that would itself be based on volatilities and correlations from the market data database. Second, they might be simulated using historical shifts, again from the market data database. Finally, they might be determined by users who wish to input their own relative or absolute shifts in yields or prices and save the resulting market data datasets for stress testing.

Valuation Model Library

The valuation model library is a calculation tool-kit for producing instrument valuations and sensitivity measures such as the "Greeks". Early designers of risk management IT placed the valuation library at the heart of their offerings in the belief that risk managers needed to be able to value any product traded in their institution and that provision of the necessary models was an essential task of the risk IT infrastructure.

In fact, the opposite turns out to be the case. Risk management vendors do not need to provide a complete set of these valuation libraries for the simple reason that the set provided will always be out of date relative to what is traded on the desk. The libraries of models needed to price these instruments already exist both in the front office, and in mainstream transaction processing systems, as well, perhaps, as in the finance function and off-the-shelf via vendors such as Financial CAD, ITO33, Numerix or Monis.

Pricing Service

Central to the computation of stress results, P&L attribution, value at risk, and many other risk results is a pricing service to produce instrument and portfolio valuations and sensitivities. This service brings together instrument reference data, market data and the model valuation library and orchestrates the calculation of the desired result sets.

Ideally the pricing service should be able to value the same instrument using different pricing methods, different valuation libraries, different data sources and on both a current and historic

basis. Additionally as the pricing service is the most computational intense part of the risk architecture, it should ideally make use of grid/clustering technology to ensure that calculation load can be distributed in order to scale as portfolio sizes increase. These requirements are summarised in Table 3 below.

Table 3. Pricing Service Functionality

Pricing Service Functionality
Fast integration of multiple pricing methods for the same instrument type
Easy and quick to integrate multiple valuation libraries from multiple vendors
User-configurable switching across multiple pricing methods and valuation libraries
Scaleable and fault-tolerant calculation distribution to cope with portfolio growth
Data source preference and data rules for dealing with missing data
Theoretical instrument valuation on both a current and historical basis
Instrument pricing available outside of the risk architecture directly into tools such as Excel

Ultimately, the results from the pricing service would be stored in the risk management database for later manipulation by the risk calculation service for calculating value at risk, or the reporting service for illustrating stress test results.

Risk Calculation Service

Pricing is not the same as risk computation and a calculation service is needed to extract and manipulate pricing results, such as computation of 99th percentile losses from value at risk scenario results, and to store the results of its calculations back to the risk management database. The risk calculation service is also the agent for most of the computation of limits against positions.

Risk Management Database

The risk management database is the core of the risk IT architecture described in this paper. As a general principle, any important risk calculation input and output should be held within it for later reporting and analysis. Specific examples of the data held in the risk management database are given in Tables 4 and 5 below.

Table 4. Risk Management Database – Static Data

Static Risk Management Data
Trading books with their purposes, their owners and the portfolios where they belong
The trading portfolio, trading desk and business unit hierarchy
Trading limits and limit types
User names and email addresses required for routine reporting
Stress test parameters
Country risk classifications
Industrial sector hierarchies for risk management aggregations
Rating bands for risk management aggregation

The risk database also holds risk measures and outputs as computed daily by other elements in the risk IT systems. Many of these data, illustrated below, are stored on a historic basis for back testing, for audit or for ad hoc enquiry.

Table 5. Risk Management Database – Dynamic Data

Dynamic Risk Management Data
Daily P&L attribution data per book
Daily VaR data, per book and per portfolio at each level of the hierarchy
P&L commentary made
Value at risk back-testing against realised P&L
Any position calculation results (notionals, deltas, vegas) which are the subject of trading limits, per trading portfolio
Standard Greeks for each position in the position database
Individual scenario P&L results from the value at risk calculation, for later extraction of key percentiles or for investigation of anomalous value at risk figures
Commentary made by risk managers on limit breaches and action to be taken
Daily stress test results, per book and per portfolio at each level of the hierarchy

Reporting Service

Once all of the data is in place within the risk management database, a powerful reporting service is required to extract meaning from all of the data present. Many of the risk reporting requirements will fall into the common slice and dice nature of reporting analysis (e.g. risk by currency, by book, by underlying etc). However, some reports will be of an adhoc nature involving complex data, which means that it is essential to provide external access to the reporting service within tools and interfaces such as Excel.

Task Orchestration Service

In practice, risk management groups run hundreds of data capture, cleansing, computation, database update and reporting tasks on an automated basis every day. They therefore depend heavily on a service to automatically orchestrate these various tasks, their hierarchies and dependencies and to automatically alert task owners should a process fail to complete successfully.

Implications for Providers of Risk Management IT

Some risk management IT vendors can provide some elements of our ideal risk management architecture to a high standard but, to the best of our knowledge, no single vendor covers the entirety of the requirement and some parts, such as exception processing and commentary, are almost universally ignored. Furthermore, none of the mainstream risk management IT specialists approach the problem with the open architecture that would allow risk management clients to adopt the eclectic, promiscuous approach advocated above.

The first vendor to fill this gap might not get rich because its clients, being promiscuous and eclectic, would only use parts of its offering in conjunction with parts sourced elsewhere. But it would provide an essential service to the risk management profession. Such a vendor would never be poor and would heal the rift between risk managers and their IT departments.

About Xenomorph

Xenomorph delivers data and analytics management solutions to the financial markets. Our TimeScape technology leverages our clients' proprietary expertise, enabling them to analyse and manage more data with greater control and transparency.

Our focus is to make our clients more successful by closing the productivity gaps between high performance database technology, data management and end-user data analysis. Through unified and transparent access to data and data analysis, our clients achieve even higher levels of financial innovation, business process efficiency and regulatory compliance.

Trading, research, risk, product control, IT and back-office staff use Xenomorph's TimeScape data and analytics management platform at investment banks, hedge funds and asset management institutions across the world's main financial centres.

Established in 1995, Xenomorph has offices in London, New York and Singapore.

For more information on our company, our clients, services and solutions, please see www.xenomorph.com and contact us at info@xenomorph.com.

Xenomorph®, TimeScape™ and SpreadSheet Inside™ are registered trademarks of Xenomorph Software Limited.

For more information about Xenomorph:

Web: www.xenomorph.com

Email: info@xenomorph.com

London: +44 (0)20 7614 8600

New York: +1-212-401-7894