



Data Virtualization Redefines the Stock Exchange

ABOUT NYSE EURONEXT

NYSE Euronext operates the world's leading and most liquid equities and derivatives exchanges. The company is comprised of six equity exchanges and eight derivatives exchanges located in the U.S. and Europe, including the New York Stock Exchange (NYSE), NYSE Arca, NYSE Liffe, Euronext and American Stock Exchange (AMEX).

Over 8,500 issues are listed on NYSE Euronext exchanges and cover an extensive and diverse set of products, such as stocks, bonds, exchange-traded funds (ETFs), exchange-traded notes (ETNs), options, open funds, warrants, commodity futures and other derivative products. The exchanges handle over four billion transactions per day with an average daily value of \$153 billion.

CASE STUDY BACKGROUND

The Global Data Services team at NYSE Euronext evaluates, designs, develops, and implements services, technologies, and architectures for the entire firm. Areas of responsibility include business intelligence; common data services; data architecture; data access and integration; data warehousing; ETL/ELT design, development and integration; and grid/cloud computing. The team's scope is managing post-trade data, that is, data immediately after a trade executes through to data warehousing, data delivery, reporting and analysis.

THE PROBLEM

A major challenge for NYSE Euronext is the sheer complexity of its business and operating environment. There are several contributing factors:

1. The organization went through many mergers and acquisitions in its evolution. This resulted in significant complexity in terms of the ability to effectively integrate data across multiple enterprises and exchanges.
2. NYSE Euronext trades 14 different types of products, from the equities market to complicated derivatives, plus commodities and futures. These all have different data structures, increasing the integration challenge.
3. NYSE Euronext deals with massive data volumes, producing an aggregate of 2TB per day across all of its exchanges and markets. Therefore, it is a challenge for the firm to ingest that much data on a daily basis, synthesize it, get it into a data warehouse, and make it consumable by downstream applications for regulatory, research, capacity planning, marketing, and many other purposes.
4. The need to meet rigid service-level agreements (SLAs) with business units for delivering and retaining data. Thus, it is difficult work just to "keep the lights on."

"External customers need fast access to this information. With data virtualization, we have the opportunity to engineer the application properly, use the appropriate technology stack, reduce cost, and improve performance. These enhancements will result in increased value and good visibility for the firm."

Emile Werr

Vice President of Global Data Services,
Chief Data Architect,
NYSE Euronext

AT-A-GLANCE

Industry

Financial Services

Business Problem Solved

NYSE Euronext needs to remain innovative in a competitive, highly complex environment while transforming the business. Automation, standardization and performance are critical to the success in meeting these objectives.

Data Integration Patterns

- Data Virtualization
- Data Federation

Composite Software Products

Composite Information Server

Data Sources

- Transactional systems – quotes, orders, trades, acknowledgments, receipts, etc.
- Reference data systems – listings and member data, customer and products data, corporate actions data

Data Consumers

Custom reporting tools and query tools accessed both by internal and external users

THE PROBLEM CONT.

From a strategic point of view, NYSE Euronext is also involved in redefining the purpose and scope of an exchange in an effort to transform the business into a broader services provider. The goal is to provide new products and services for existing customers and to find new customers for existing products and services. Clearly, customers have options for trading securities, NYSE Euronext must be innovative to attract the trading – along with the associated fees and opportunities for other forms of revenue – to its exchanges. NYSE Euronext needs to remain innovative in a competitive, highly complex environment while transforming the business. In addition, all of the information generated by these new business services has to funnel into the NYSE Euronext data delivery environment.

THE SOLUTION

The NYSE Euronext solution is an enterprise-wide data virtualization layer, built using Composite Software's Data Virtualization Platform, which functions as a virtual data warehouse to provide access to post-trade data for analysis and reporting. The solution has been named TORQCA (pronounced "torkah"), an acronym for the major data transactions that comprise the business: trades, orders, reports, quotes, cancels and administration (admin messages).

NYSE Euronext has multiple data centers located around the globe (each exchange has its own data center) and many different data warehouses and data marts. Given the volume of data and the strict SLAs for retaining data online (data needs to be quickly accessible for at least seven years to meet regulatory requirements), multiple physical data warehouses often exist to store different time slices or levels of granularity of the same data. Performance is very important, not only because of data volumes but also because response time for analytics must be fast. The company, therefore, is a strong adopter of massively parallel processing (MPP) technology.

In the NYSE Euronext environment, the data virtualization layer has many functions and roles. It provides all of the following:

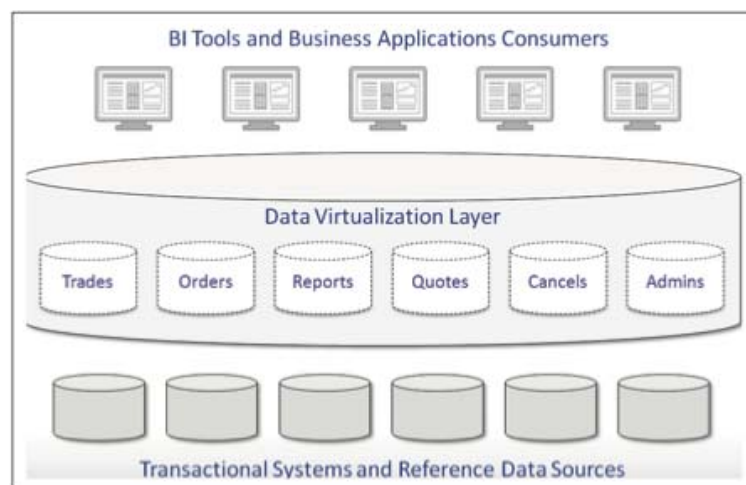
- Standardized data access for connectivity to all data sources.
- Virtual data warehouse for federating data through logical views.
- Centralized repository for common metadata, application/business logic and connectivity, and data services.

Changes to any of these centralized components are made in the data virtualization layer, eliminating the need to redeploy code to applications.

This makes it easier to embed business requirements directly into the development life cycle without going through multiple layers of translation at the application level. The data virtualization layer also makes the environment extensible and provides significant opportunities for reuse. These are all important drivers for faster systems development going forward.

Three primary factors will be critical to the company's ultimate success with data virtualization:

- Ability to migrate existing SQL without modification — Pass-through SQL will accelerate the ability to unwind a lot of the point-to-point systems and redirect them to the data virtualization layer without having to retrofit all of the applications that contain embedded SQL.
- Ability to capture how applications use data — Every request will go through the common data access layer, giving the company knowledge about all system activity with the opportunity to proactively evolve and improve the data architecture and performance architecture.
- Composite partnership for data virtualization — Using the Composite platform enables the company to centralize and provide universal access to federated data without reinventing a technology that works well.



A virtual layer based on the Composite Information Server.

THE RESULT

- Anticipated savings of over \$4.5 million annually