

ODBMS.ORG User Report No. 11/08

Editor Roberto V. Zicari- ODBMS.ORG www.odbms.org

July 2008.

Category: Industry

Domain: Finance

User Name: Richard Ahrens

Title: Director

Organization: Merrill Lynch, US

Richard Ahrens, Director at Merrill Lynch. I run the Global Trade Lifecycle development team within Equity Linked Technology. Prior to ML, I ran Automated Market Making and flow derivatives trading technology at Lehman Brothers. I have a total of 12 years' experience in the financial services industry. I'm a graduate of the Wharton School of the University of Pennsylvania, Philadelphia, PA.

Q1. Please explain briefly what are your application domains and your role in the enterprise.

Richard Ahrens: My team provides order management, pricing, and trade capture technology for Merrill Lynch's Global Equity Linked Products division. We have implemented a product-neutral order and RFQ (request for quote management) platform that allows us to manage customer orders and inquiries for equity derivatives from initial interest to trade booking.

Q2. When the data models used to persistently store data (whether file systems or database management systems) and the data models used to write programs against the data (C++, Smalltalk, Visual Basic, Java, C#) are different, this is referred to as the "impedance mismatch" problem. Do you have an "impedance mismatch" problem?

Richard Ahrens: We absolutely have this issue. In the equity derivatives business, agility and time to market are incredibly important. New products are introduced regularly, and flexible technology is required to adapt and keep up with the rapid pace of change in the industry. In our domain of order and quote management, having to maintain object-to-relational mappings across a wide variety of product and asset times is a drag on developer productivity and limits our ability to scale.

Q3. What solution(s) do you use for storing and managing persistence objects? What experience do you have in using the various options

available for persistence for new projects? What are the lessons learned in using such solution(s)?

Richard Ahrens: Our order and quote management system combines an embedded object-based continuous event processor with an embedded object database. This allows us to rapidly add new derivative products to our environment and keeps developers focused on writing code that adds direct business value. With our design, we have strived to eliminate "non-productive" development: keeping objects in sync with a relational data model adds no value to our business, so we rely on object database technology to make that problem go away.

We have found this approach not only enables us to deliver incremental functionality faster, but also reduces our testing burden since there are fewer moving parts for us to maintain ourselves.

Q4. Do you believe that Object Database systems are a suitable solution to the "object persistence" problem? If yes why? If not, why?

Richard Ahrens: In many cases, yes. I would still not recommend object databases for our most throughput-sensitive applications, but in situations where there's an acceptable tradeoff between time-to-market and transaction volumes, object databases can simplify life greatly. For us, the choice to use an object database has directly increased our ability to deliver business solutions in the shortest time possible.

Q5. What would you wish as new research/development in the area of Object Persistence in the next 12-24 months?

Richard Ahrens: We're very interested in the relationship between complex event processing of object streams and object persistence. It would be great to see thought leaders in these two fields collaborate to produce integrated products that combine the strengths of each projects.