BEATING RISK IN RETAIL POWER.

An In-Depth interview with Molecule Software CEO Sameer Soleja

MOLECULE SOFTWARE

Sameer, what are the unique, sector-specific challenges a vertically integrated power company faces with respect to trade capture and risk management?

The most obvious issue is the complexity of electricity in general. Customers in this space have to deal with different block sizes, MW/MWh conversions, the peculiarities of different ISOs, times of day, and the volume of data associated with all of that. This can be seen on both the retail and the generation side of the business.

The generation side faces issues that are quite unlike any other paper trading that we've seen - from dispatch modelling of the plants themselves to modelling of complex financial hedges and tolling agreements that help either feed the plant or sell its power.

Combining all that data, hedges, plant models, optimization plans, retail purchases and sales of power, and load forecasts can be a world of hurt that few other systems manage well.

What feature sets does Molecule provide to help retailers and generators deal with those specific issues?

First of all, the basic mark-to-market analysis for financial and physical trades is very easy to use and maintained to the point where it's almost self-driving.

Also, we've integrated a business intelligence solution in Molecule that helps us display and/or aggregate data in ways that would be hard for an external vendor to do or would require the hiring of a separate business analyst to arrange. We calculate metrics such as gross energy margin very easily, and our customers can define what they want to see in that report. We can then assemble that report quickly and without significant extra charge.

Finally, we can plug complicated external valuation models into Molecule for things that don't fit neatly into a mark-to-market or a Black-Scholes valuation, but that require things like a Kirk's Approximation, a dispatch

model, or perhaps something even more complicated. We can host your valuation models, ping them through our APIs, and get valuations and option Greeks out of them. We then combine them in our custom business intelligence reports with your plants, your generation capacity reports, your hedges, and your retail trades. They are displayed in Molecule in a custom view defined by you.

Who has traditionally championed Molecule amongst electricity clients?

The people who are always the most obvious beneficiaries of Molecule are the operations people. People in the middle office who have to munge all sorts of bits of data together to get analysis that work love Molecule because it literally takes hours out of their day.

We also serve retailers. For them, we combine their wholesale purchases with their financial hedges and even an element of their retail load. In the summer we will begin to automatically schedule their wholesale purchases on ISOs.

How does the Molecule VaR help trading managers deal with extreme volatility associated with rapid switching from hydro base load to wind variable load in countries such as New Zealand where renewables dominate?

When our clients started asking for a Value at Risk (VaR) module, we researched formulas to find the best predictive model. We found a range of shortcut analyses for VaR from delta-normal calculations to variance/

"You can trust that Molecule will alert you when something's not right."

Every day. We also see front office analysts who can better analyze their portfolio and managers getting great, low-touch, high-level metrics out of Molecule.

Does Molecule have customers who are both retailers and generators—companies who have to separate two sets of books? What features have been built for them?

We have customers who are retailers with varying levels of wholesale exposure. We also have customers who are generators with typical generation requirements. One of our largest customers is a global top five financial services firm that has a significant amount of generation. We have built a set of tools around gross energy margin for them and around their market data, which enables their portfolio managers to quickly get the analysis they want every day without having to munge a bunch of spreadsheets.

covariance calculations. We decided we wanted to give our customers the broadest, most reliably predictive toolset available—and we weren't bounded by horsepower—so we built a real-time Monte Carlo VaR. Our VaR runs throughout the day and re-runs as every position changes. We run ten thousand simulations at a time, and the VaR typically updates within minutes.

In the past, if you worked for a trading firm, you might see a sheaf of papers on your desk every morning with last night's VaR numbers. The moment you started the day, however, you'd be in the dark. With our real-time VaR, that all changes. If your VaR is updating every two minutes, you can see your exposure at any point in the day and limit your risk more effectively.

In addition, we're building automated monitoring metrics for our customers'

portfolios that'll do things like alert you when a number goes out of whack. Sometimes that's just a "fat finger" of a trade, but sometimes it's a real problem, and either way, we figure you should know. Molecule can really help you sleep better at night because it's doing a lot of the heavy lifting for you. If you're a risk manager, you can trust that Molecule will alert you when something's not right.

Talk us through the implementation process for Molecule. How long does it take? What are the steps? Who's involved? What level of input is required from the client. How does Molecule's process differ from the major vendor implementations that you've been involved in?

We can do implementations for customers in-house at no additional charge to our quoted package price - in 90 days or less. Our team in Houston typically runs two or three new customer onboardings at any given time. Within a few weeks your trades and market data will be coming in, your P&L will be producing accurate numbers, and custom reports will be built. By comparison, elsewhere in the market, implementations and upgrades take years.

Typically our implementation team will sit down with your key people to set and receive expectations. Together, we will outline key variables such as what reports need to be produced, where trades will be coming in from, and what market data will be used. We then start putting it together, and check in with you once a week.

We send you a login to Molecule once new trades start to arrive (typically within days). Soon afterwards, analytics start to become available. We then refine your books and reports over the next few weeks. Once all is well, we bring new production customers into our regular monitoring program where we look for exceptions and anomalies in their portfolio.

Finally, Molecule runs upgrades to the product about once every two weeks. They happen almost transparently to the customer - new features simply show up. The days of having multi-year projects for ETRM implementations are completely over.

So while Molecule fees are similar to those of other ETRM vendors, our implementation process dramatically drives down the total cost of ownership of our product and the amount of effort required from our customers to do their jobs. You don't need to bring together a custom implementation team.

Customers will typically be involved for up to eight hours in the first week of the Molecule 90-day rollout cycle. After that their involvement will be one to four hours per week (meeting with us for half an hour and working with the system for another couple of hours). At the very end of the implementation process, there is a 16 to 32 hour process of validating the reports we've built. So in all, the customer input required in a Molecule rollout amounts to between 40 and 80 hours.

Implementing Molecule doesn't have to be painful. We can do a quick wrap on a Saturday morning over breakfast or via Skype if you don't have time during the week. It's very simple.

By contrast, I worked on an OpenLink implementation that had 30 employees working on it for 18 months, and everybody seemed angry and upset all the time. Everybody went home stressed. Tens of millions of dollars were spent to get the solution working.

That experience was the impetus for starting Molecule. It just doesn't have to be that way and we've made a product that is far better.

Looking forward, what's on the roadmap for Molecule Software for 2017 in the electricity generation and retail space? What should people be getting excited about?

We're finishing up some more advanced valuation models for the generation side. We've been building a quasi-proprietary power plant valuation model and integrating the open-source QuantLib into Molecule so that all customers have dozens of valuation models available in this battle-tested library. So if you have a plant or a tolling agreement or want a funky dispatch model, we can run it through our API, integrating its results with all the rest of your analytics.

In the summer, we will roll out the view on our first power scheduling feature for U.S. power. Schedules will automatically be sent off to ISOs for confirmation so that users don't have to go and buy a million dollar copy of, for example, nMarket. This functionality will be transferable to European markets and those in the Pacific Rim.

ENERGY SOFTWARE THAT ROCKS.

MOLECULE SOFTWARE
1113 Vine Street
Suite 200
Houston, TX 77002
molecule.io
+1 832 464 4037
info@molecule.io

MOLECULE SOFTWARE