CASE STUDY

SPHERA

Operating on modern market venues requires the ability to manage financial transactions in real time without any significant delay from the instant the order is received to the moment the order is sent to the market. Meanwhile, the trading system must ensure the suitability, appropriateness and best execution of the operations, in order to guarantee the compliance to the rules of each market and the national/international financial regulations. To be able to leverage its assets the end user also requires a plethora of different facilities, ranging from custom strategies (such as orders triggered by specific market situations) to risk and performance evaluations, even on massive historical data series, all to be performed without inflicting any significant delay on the market activity.

Sphera, the Kline Trading Platform, is an ecosystem of software products offering direct access to world wide financial markets and/or brokers, to trade assets and receiving real time updates from the markets. It encompasses clients for different desktop and mobile platforms, connectors built to leverage the different markets APIs, and a cluster of servers performing a wide range of real time checks on the customers' operations, all while evaluating and tracking the evolution of the markets. In addition, the system has to keep analyzing the incoming exchanges data to provide additional information to the end user in the form of plots and risk/performance indicators.

Our typical customer can manage up to many hundreds of requests per second. Each request requires a wide range of checks, ranging from database centric ones (the account exists and is enabled to operate on this specific market, the cash amount is enough to perform the operation, ...), requiring more than a single query, to algorithmic ones (e.g. to calculate the money necessary to cover a position in case the market collapses).

In 2015, Kline was quite close to reach the performance cap of its legacy framework and began planning a new iteration in the development of its trading platform. We decided to move away from the legacy architecture, in which the storage database was also a key component during the operations, choosing instead to separate the end storage assets from the operational components.

THE SOLUTION

It was clear that that solution should have to be based on in-memory Database, and after a round of tests involving several products, Volt Active Data proved to be the best candidate for the role of operational database to support our real time operations.

Volt Active Data was integrated in our ecosystem to leverage the characteristics of an in-memory database to cut down the time required to complete

the db-centric procedures, which could be a significant bottleneck in the previous disk based implementation. In addition, we have been able to maintain the advantages related to the ACID consistency and the chance to keep working with standard SQL, to which our developers are already accustomed.

Another advantage of the Volt Active Data implementation is the ability to scale with minimal effort our infrastructure to different machines, to adapt it to the requirements of each customer in terms of cost, performance and continuity of service. The latter point is by far the most sought-after in this market, and Volt Active Data implements natively all the facilities to enable it without any delay.

Finally, a multiple-hosts Volt Active Data instance can also be configured as a high availability cluster, in which the failure of at most a pre-determined number of nodes does not cause the loss of the cluster functionality. This is especially important for large customers that cannot afford even the slightest downtime; this also works well with our infrastructure, which has always been designed as a multi-node system with services replication for both load management and high availability.

THE RESULTS

After the core phase of the development of the new solution was completed, we devised a setup to allow us to gauge both the functionality and the performance of the new system. We decided to perform a campaign of parallel testing, by installing the new system along a legacy one in a real production environment on one of our customers' datacenter, while keeping it isolated (read only) from the production dataset. This way our new trading platform could be fed with the same input (both from markets and end users), while accessing the same initial data. The results of the analysis performed by the new system was fed into a separate storage and matched in real time with the ones coming from the legacy system. In addition, we started monitoring the timing data from both the new and old system, to perform a comparison. Here is a little sample of the results we obtained, separated by end market venue and operation.

We chose to split this way the results because the amount of work behind each operation is strictly bound to the kind of the operation itself and the market it is performed on. As shown in the table, the most demanding operation are creating or updating a market order, because that requires to perform all the tests related to the adequacy, appropriateness and best execution; conversely, the cancellation of an order only requires a smaller amount of checks and updating the portfolio status. At the same time, different markets have different requirements mostly because of their own way to calculate the money at risk during each operation. Each market provides its rules to determine the risk involved in an operation, and the calculus can be difficult and involve a fair amount of data.

As can be seen in the table, the result were astonishing. Thanks to the rebuilding of the system with more modern frameworks and the use of Volt Active Data as caching/operational database we have been able to improve the performance of the database driven operations by 30 to 50 times each, while continuing to assure the correctness and transactionality of the orders.

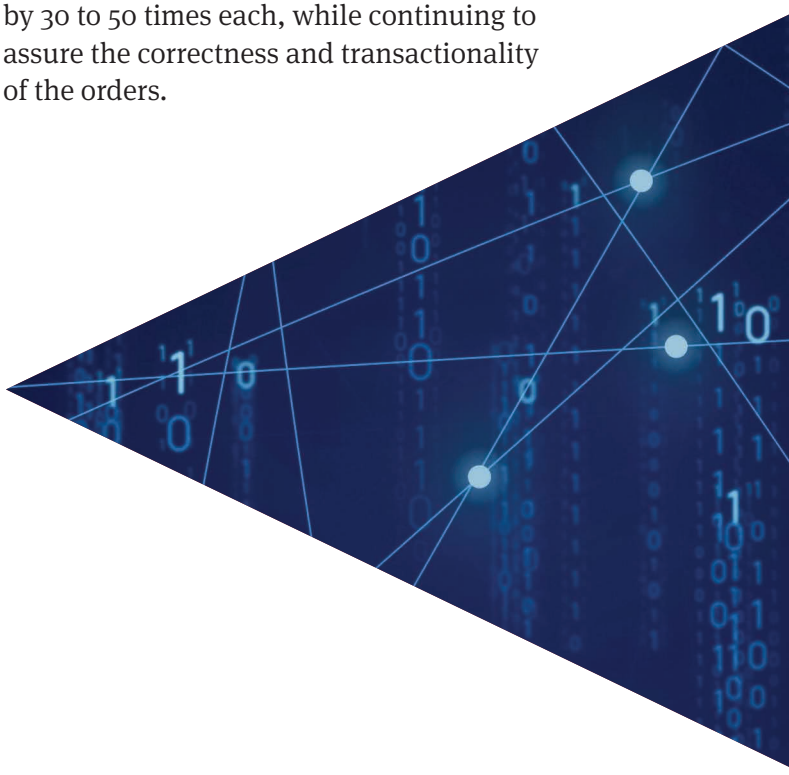

MARKET VENUE	OPERATION	OLD TIME (MS)	NEW TIME (MS)	IMPROVEMENT (TIMES)
MARKET 1	*Insert new order*	197	3	65.6
	Update order	201	3	67
	Cancel order	2	<1	3
MARKET 2	*Insert new order*	135	2	67.5
	Update order	150	2	75
	Cancel order	1	<1	4
MARKET 3	*Insert new order*	262	4	65.5
	Update order	262	4	65.5
	Cancel order	2	<1	5

NEXT STEPS

Now that we have been able to rebuild all of our trading infrastructure to leverage Volt Active Data, we are beginning to think about future features to implement by taking advantage of the new capabilities provided by this new component. Our research work is now aiming towards providing our users complex analysis of data in real time, something that the old infrastructure was unable to perform, to show indexes, market forecasts and historical plots and charts. Volt Active Data has proved very well suited to perform this kind of activity in real time, and our first tests are extremely encouraging.

ABOUT VOLT ACTIVE DATA

Volt Active Data enables enterprise-level companies to innovate faster, perform better, and create new revenue streams by unlocking the full value of their 5G data. The only data platform built for real-time, sub-10 millisecond decisioning, we empower companies to re-engineer their latency-dependent solutions to process more data than ever before at a faster pace than ever before, allowing them to not just survive but thrive in the world of 5G, IoT, and whatever comes next. By combining in-memory data storage with predictable low-latency and other key capabilities, we can power BSS/OSS, customer management, and revenue assurance applications that need to act in single-digit milliseconds to drive revenue or prevent revenue loss, without compromising on data accuracy. For more information, visit voltactivedata.com.