Advance to GO

With a smart strategy of introducing its unique customized solution way before its official release date, the official UnRisk FACTORY 1.0 has skipped all players on the board and is headed straight for its pay off.

Rolling out the official UnRisk FACTORY 1.0 was really just a formality for the UnRisk Consortium. Already adopted by five pioneer banks in Europe since early 2007, the 1.0 version of the product is already an experienced soldier on the field and as its developers emphasized, it is 'bank proof'.

The UnRisk consortium is a partnership between MathConsult, headed up by Andreas Binder, and Uni Software Plus, both operating from Linz, Austria.

We speak to Herbert Exner at the UnRisk Consortium and find out how the FACTORY has matured over the last year and if there is something beyond the FACTORY.

Since we spoke last year, what have been the developments with UnRisk FACTORY?

We gained a lot of practical experience from our clients, making UnRisk FACTORY more comprehensive and faster. We entered into developing simulation capabilities as preparation for a very fast Value at Risk module for our pioneer customers, making it bank proof before it is even officially released as a product. What was released is UnRisk FACTORY 1.0, which is not your typical 1.0, since it is already bank proof.

The Austrian Central Bank, UnRisk PRICING ENGINE users, decided to enter into this pre-release program.

We have also started a Value at Risk module, which will be built on top of the FACTORY. With the VaR module, we profiled our pioneer users and built it based on their requirements, feedback, practical experiences, and based on previous product price. This is scheduled for release in autumn this year.

With the turmoil in credit and the US subprime issues bringing risk to the forefront, how will or has this affected the development of UnRisk FACTORY?

We emphasized on the foundation and the technical aspects, which we believe are very important. You can have very clever models, but if you do not represent them properly or solve the model correctly, you will have problems. We do not segmentize or say we need new strategies just because there is a crisis.

The crisis will probably motivate the regulators to force market participants into doing more comprehensive risk analysis and more quant work, which is good for us because that is exactly what we offer. Additional requirements could arise from this and some will be a challenge, but we will transform those into technical foundations.

If indeed, there are new requirements on the modeling, the algorithmic representation and algorithmic solving, we will look into it and do the best possible job because our strength is taking given requirements and apply the best mathematics we can apply.

So we are not trying to replace the financial strategies, instead we are trying to replace quantitative work to make everything more visible, understandable and valuable.

Besides the market conditions last year, are there other international financial events that may influence product development?

I think risk management will become much more important, although I cannot tell you what kind of unexpected events you need to take into account. However, with our rich mathematical foundations and background, we will be able to rise to the challenge. Right now, I think people fear everything, and they are likely to hedge everything against everything, which yields no profit. This is not really our area, but taking care of that fear is a new risk management level. The question is do we have the mathematical background, knowledge, and foundation to do this and we have.

Are there any significant partnerships that have been forged in the last year?

We entered into a partnership with Microsoft and Dell in early 2007, and we have designed and logically created this UnRisk FACTORY in a grid box, where we have specific configurations designed. This is an optimized utilization of our software on a specific hardware, we repackage it and Dell will offer it in a complete turnkey solution. This is one of the most important partnerships we entered into last year.

A boxed Dell cluster special in-a-rack, it will have Microsoft Windows Compute Cluster Server, Microsoft SQL server, special web services and of course the UnRisk FACTORY, we often call it a 'quant in-a-box'.

Can you give us a little history on how the partnership came about?

We deal with the Mathematica team regularly because our software has Mathematica as a platform. They, being in HPC computing, have contacts with all the HPC computing groups. We had the philosophy that finance and HPC are practically twins, and we believed that finance people should hone their skills in their expert finance work and not in IT. That is to concentrate on applying and exploiting different areas in finance, instead of setting up a complicated supercomputing environment.

This brought us into discussions with Microsoft who share our view and the idea of combining forces to do that. Microsoft asked if we had tried specific hardware pieces where these things fit very well together. In addition, most of our pioneer customers use Dell, which works well. Therefore, we came out with the idea to place a concrete UnRisk FACTORY in a grid box, based on Dell's power edge compute nodes, and then everything else is just marketing.

Dell and Microsoft then introduced us to the HPC center at Cambridge University, where we achieved comprehensive bench marking in finding an optimized box. That to me, means getting the linear speed up, and finding out where the linear speed up starts to flatten because of losses to communication areas. We tested in a 256-core sub-cluster and we found that the right size of the box comes down to a 24-core, 40-core or a 72-core configuration, depending on the size and requirements of customers.

Typically, we sell UnRisk FACTORY on any of such computer grids and we put together the pieces at the customers' site, but this Dell grid box will be shipped completely preloaded.

Herbert Exner

Basically, we preload the grid box with all the standard components, which is your Windows CCS server, Microsoft SQL server database, web services and the FACTORY, but it is optimized in such a way that we have pre-selected the correct sizes derived from the comprehensive analysis.

Is this solution targeted at a specific group of customers?
We are targeting the people who do not want to waste time with comprehensive IT stuff, or do not even have this comprehensive IT organization, such as private banks or small retail banks, capital management firms or funds.

With all this hype on HPC, how far are we from making HPC an everyday reality?
Everything is prepared; it is more an organization and offering issue because we have the clusters, which are built upon simple multi-core or single core machines. We have the software to put these things together. The multi-core revolution is a revolution, which simplifies HPC in a desktop way. It is more an issue of setting up and deployment than a technical one.

The components are there and they will be improved as always, but there is a psychological barrier. I remember during the stone age of IT, we talked about the mainframes, which meant more performance than on a dedicated workstation then. Now, people are talking supercomputing, about large clusters to be managed or optimized for deployment. If you take just the PC cluster, it already is an HPC unit, you just need to put the software together, parallelize it properly, which is not too complicated and viola! The psychological barrier is that people cannot believe that they can have something so powerful under their desk.

Viola! The psychological barrier is that people cannot believe that they can have something so powerful under their desk

What we need to do is to explain and promote that this is the real performance revolution–multi-core on one hand and clustering on the other.

You mentioned the psychological factor as a main barrier; could there be other issues in the way?
I think there are little barriers from the customers' side. However, I think the development community is still searching for the right methods and tools to parallelize their software. For instance, if you are a C++ programmer, you need to decide if you will do threading or use other distribution tools to parallelize your software. For us, it is very simple because our gridMathematica approach allows us to just input a few lines of codes to parallelize our software. This enabled us to put out a grid-enabled version of our PRICING ENGINE very quickly, and this is the basis on which the FACTORY was built.

Our approach not only explains and breaks the psychological barrier from the deployment and installation, it can also convince the development community to select tools to make these things simpler.

What are the new challenges that UnRisk FACTORY may face?
Being able to compute so much data, transaction processing becomes a real bottleneck, which we now really need to optimize.

The challenge for the future will be to bring in more expertise and knowledge of the requirements into the system. We have thought of developing something we might name UnRisk DIRECTOR, on top of the FACTORY, which will take the requirements of a term sheet and analyze it automatically. For example, it will ask a user for a risk profile they want to apply and then select the right models, the right calibration environment automatically, in order to get all the pricing, risk analysis and scenarios calculations from that.

Right now, every system in the world asks for specific inputs, such as which model for a specific instrument type, the type of calibration, market data etc. We would like to be able to automate that in the future, but this is very challenging. We have the basic technology for this, proven in technical systems, but we need to transform this into knowledge.

Again, our driving force is that we want to put quant work on a computer. I am a fan of the evolution of tools. First, we think of how to automate lower level work, then the next level. This will be the next level of decision-making.

Right now, we are concentrating on the rollout of UnRisk FACTORY 1.0. After that, maybe in a year or so, we might begin to have prototype results on UnRisk DIRECTOR, which really is still at a planning stage. We have all the ideas and decision support tools, it is a question of transform the 'what' into 'how'.

As you have mentioned, the ability to compute so much data brings about challenges in transaction processing; how do you deal with the ever-increasing data?
We do not provide data, we expect that there are systems around which do exactly that and are very good at what they do. What we have done is we have written a comprehensive adapter to our FACTORY, which is able to import information from any provider or any front to back office system. We concentrate on what we are good at and this is transforming information into good, quick and very accurate results.

We concentrate on what we are good at and this is transforming information into good, quick and very accurate results

How do you ensure accuracy is not compromised when speed is being increase?
From the beginning, we were forced by technical limitations to never waste a single core second. Our algorithms are optimized for single core. On the HPC level, we distribute this single application to different cores in parallel and at close to linear speed up. For example, if we have 32 cores, it is close to 32 times as fast. To take this further, we are able to calculate thousands of single valuation positions in about an hour.

With this in place, we do not need to make compromises in the models or algorithms replicating these models. Our approach is to keep the original accurate model and apply it to the core we have available. It is a very simple approach, but this is the way it should be done, from our point of view.

We do not deal with inner or fine grain algorithms parallelization, because you can lose a lot of speed through low-level communications. Instead, we do coarse grain parallelization, which is simple, scalable, and diminishes the risk of losing information in the algorithmic jungle.

Have there been any new adopters of UnRisk FACTORY in the past year?
We have an arrangement with our pioneer customers that we will not bring in any new customers before we roll out the system officially. Nevertheless, we have had several talks with interested parties.

One new development is with Solventis, one of our pioneer customers in Spain. They will adopt their business model and apply an application server with online capabilities to do risk analysis at their customers' site. With our web front-end, Solventis can run the FACTORY in their environment, put out specific functionality, which allows their customers to access the FACTORY on their own web browsers. Solventis uses the simple fact that UnRisk FACTORY follows its users.

What feedback have you gathered from your clients?
The most important feedback from our clients is on the usability of the FACTORY. Our clients are satisfied with the accuracy, speed of single calculations, and scalability with our grid computing approach.

The most feedback we have received is on how you manage all those things; how do you set up your users; how do you give roles to the users; administrate the right users, get the right information, and this is very selective. In addition, how we do that so that we can simplify and automate certain tasks like putting together your portfolio across scenarios; how do you call up information at your fingertips?

We also received feedback on the type of information we need to keep so that reviewers and auditors can look into and replicate what have been done in the banks, for example to do a risk analysis or a valuation.

These are the core feedbacks we have got and are indispensable. As a software maker, you might have an idea, but your clients can give you additional practical recommendations. This helps us turn the system into something comprehensive and easy to use.

How would you sum up UnRisk FACTORY in a single sentence?
Platform for risk analytics in real time.

What is the dream goal for the UnRisk Consortium?
Our dream goal is the UnRisk DIRECTOR, which will support the users in, not only valuation, insights on the right scenarios and results, but also supporting users in their decision-making. If you think of the name of our products: UnRisk PRICING ENGINE, which is a single machine tool; UnRisk FACTORY, which is like a real factory built of single machine tools; the UnRisk DIRECTOR will be something, which directs the FACTORY to do the right things with the right operation to make even better products. In our terminology, even better information on accessing risk spectral and better decisions. That is our vision for the near future.