

Paper 334-2007

Evolution of Financial Risk Analytics at EnCana

Gord Anderson & Mark Perrin, EnCana Corporation, Calgary, Alberta

ABSTRACT

EnCana is one of North America's leading independent gas and oil production companies. We buy and sell product in the market place and need to offset and monitor risk from market movements, and assess risks associated with our customers, to ensure business is conducted within acceptable financial risk profiles. We have needed greater sophistication in our Risk Analytics to calculate "at-risk" metrics on revenue and mark-to-market including VaR, CVaR, PFE, CFaR along with other market, credit and liquidity metrics.

Our existing systems were not up to the task so we implemented SAS Risk Dimensions (RD). We started early in 2005 and had our initial implementation running by April. We then began enhancing the analytics within RD and infrastructure that support it.

This paper will discuss where we started with RD, the enhancements that have been made to date and where we are planning to go in the near future. This evolution will be presented from both the business and IT perspectives.

Gord Anderson, Market Risk Analyst, Financial Governance, will present the business perspective. Mark Perrin, IT Group Lead, Financial & Market Risk, will present the IT perspective.

WHY RISK ANALYTICS AT ENCANA?

We are active participants in the physical and financial energy commodity markets. This activity brings risk, both exposure and uncertainty, for us to mitigate. One of the ways we can mitigate risk is through hedging. It is important for us to have analytics to assess the impact of these hedge transactions prior to execution in order to achieve the desired result. Actions engaging customers, business segments, and individual transactions all contain risk. Depending on the appetite for risk in a given area, we need to assess the risks associated with transactions.

Our analysis provides better information on which management can base decisions. In the case of customers, we need to analyze the credit and liquidity risk to ensure we are not overexposed to a customer and can meet margin requirements if necessary. Potential implications to financial results due to market movements helps to identify the areas in which we are at risk, while assessing the impact of our corporate hedging strategy allows us to see how effectively we are mitigating those risks.

The why and how of analytics at EnCana is a constantly evolving process. As markets and fundamentals change, so do the objectives and requirements of analysis.

INITIAL PROJECT OBJECTIVES

At the start of the project in early 2005, there were several objectives. One of the primary objectives was to enhance the validity of our current at-risk metrics. These were being calculated using a variety of methods and technologies. Consequently, there was no consolidated corporate view. Individual business units, such as natural gas and crude, produced their own separate metrics. These types of metrics could not simply be added together, and needed to be calculated starting at a risk-consolidated basis. Most importantly, we wanted to produce a true corporate view in order to evaluate risk as it pertains to the financial outlook of the company.

To achieve this, we needed to combine the physical and financial transactions of our main commodities (natural gas, crude, and liquids) into one risk portfolio. Our goal was to apply a standard Monte Carlo simulation methodology to this portfolio to generate CFaR (Cash Flow at Risk), PFE (Potential Future Exposure; credit and liquidity exposure), VaR (Value at Risk), and ad hoc scenario and PL (Profit/Loss) Curve analysis. An additional objective was a faster desired turnaround time for our CFaR process – less than one day – to allow more responsive analysis. Most of the existing metrics were being run by business units through existing multifunctional systems. In achieving these objectives, it was also important that the environment for testing ad hoc scenario and stress analysis would not negatively impact the business resources, both people and systems.

In addition to the initial objectives of the project outlined above, reactions to changes in the business and a greater understanding of the tool led to additional requirements. Subsequent to the implementation of SAS®/RD at EnCana, new challenges and changing objectives have continued to necessitate a constantly evolving analysis. This is discussed in the final section of this paper.

HEDGE ANALYSIS: EVOLUTION AT ENCANA

One of the drivers of the implementation of SAS®/RD at EnCana was our hedging program. We will now examine one of the models and analysis revolving around this program. The main focus of this model is to project cash flows and allow for layering on hedging strategies to see the impact they have on our overall company position, while also seeing the impact of existing strategies in place.

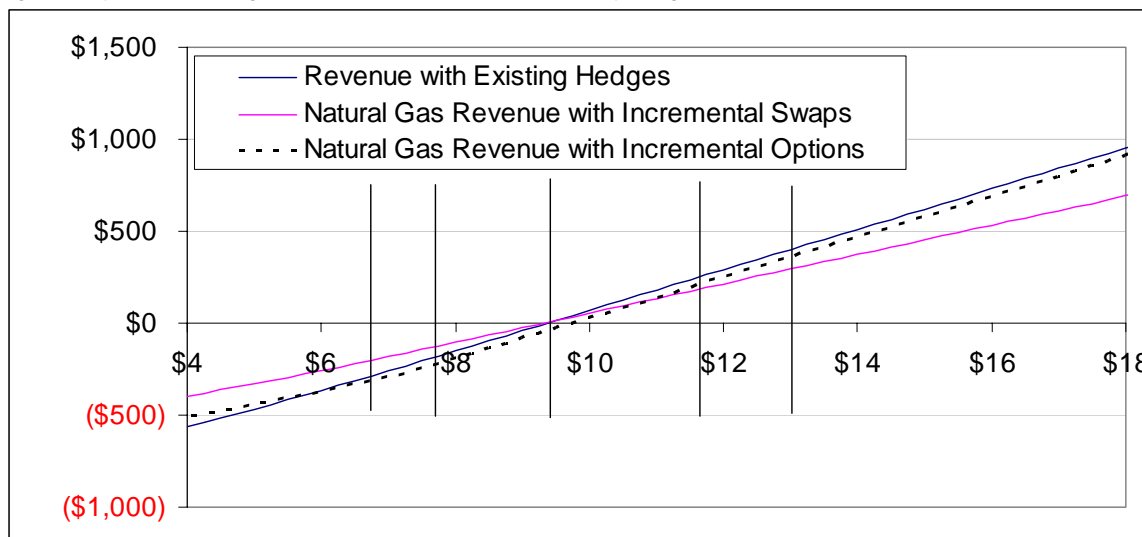
Our cash flow model originated as a parametric cash flow at risk calculation. This was achieved through the use of existing risk management applications and Excel. It was a very cumbersome process and did not incorporate all commodities to give a consolidated corporate view. This led to the introduction of SAS®, and more specifically Risk Dimensions, into EnCana.

Our first goal with the new implementation of SAS® was to take the results of our analysis and ensure they aligned to our financial statements. This also entailed validating our assumptions in the model. We used our budget group and their process as a benchmark to match against the base case. The simulation results are unique to the SAS® results and therefore have nothing to reconcile against. In the Challenges and Initiatives section this will be discussed more as to how we were able to true up the model that produces the simulated results.

Once we were able to align our base case with our financial statements, we moved it a step deeper and incorporated more detail into the cash flow model. We paid specific attention to tax and transportation and selling expenses. Initially these were taken as flat amounts, but as our simulated revenue changes, we modeled in fluctuating values for these pieces and a few others. This was dependent upon further alignment with both our tax and budget group. We needed to pay close attention to the complexity of the model in reference to the accuracy of output. One other important consideration was the ability to explain the model at many levels of management. Again this will be discussed further in the Challenges and Initiatives section.

NATURAL GAS REVENUE

In the first stages of developing the model, we needed to generate a new valuation. As will be discussed further in the SAS®/RD setup section of this paper, we initially looked at Mark to Market (MTM) only. It was identified that a revenue calculation in the pricing methods was a start to producing the cash flow model. As a result, we generated revenue based charts reflecting our existing portfolio and potential hedging scenarios. Observe in the chart below we have a swap and option scenario that are included with our base case. These illustrative revenue values are plotted against Nymex natural gas prices as our portfolio is heavily weighted in that area.



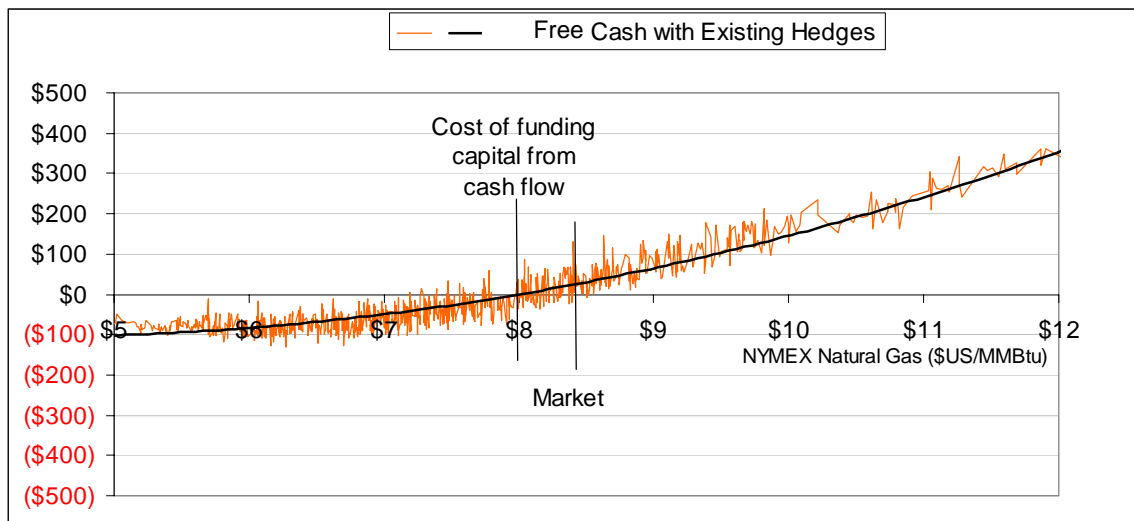
This was a preliminary graphical representation of our analysis. Given a growing adoption of the methods and analysis by management, our model evolved very rapidly. This was fueled by the ability to explain the model at every level of management.

CASH FLOW MODEL – FREE CASH

It did not take long for our initial revenue model to become what we now refer to our Free Cash model, based off of our CFaR analysis. Below is a high level of what is used to come up with our Free Cash numbers. The basis of this is, of course, our revenue simulation from SAS® /RD.


- Cash In
- Sales of energy commodities
- Cash Out
- Capital expenditures
 - Commodity transport and processing
 - Taxes
 - Administration
 - Share buybacks and dividends
- Difference: Free Cash

Free Cash is how much we would expect to have leftover at varying levels of pricing (Nymex natural gas). The goal is to have some degree of certainty that this number will not be less than zero. If this is not the case, the next step is to design strategies to ensure positive free cash or at least increase the certainty, or level of confidence. In the chart below is an illustration of how we present this. Recognize that if the current market prices stayed the same, we would be looking at a free cash value of approximately \$30 million. The cost of funding capital from cash flow is where this free cash value is 0. This also shows that no matter what happens, there is a great deal of certainty that we will not lose more than \$100 million. The example illustrates that there is a significant strategy already in place, and most likely the \$100 million loss is as a result of the cost of that strategy (i.e. buying puts). As with the revenue chart previously, this can also be shown readily with additional strategies.



Another consideration when looking at this type of analysis is what needs to be done if Free Cash is above or below 0? If it is below, do we simply borrow more? If it is above, what do we do with all that extra cash? We have stayed away from diving too deep into the effects of these questions. In the case of having extra free cash, and choosing to spend on share buyback, this could impact other pieces within the model. If share buyback is not done, then the extra cash would be sitting in the bank or put directly to debt, either saving interest or generating interest. If the choice was made to buyback shares, this interest generation or saving would not exist and result in a shift in the model. This is an area where we chose to rely on simple indications of costs of programs versus fully integrating them into our model. A choice of simplicity with ease of explanation over accuracy was chosen in this case. As illustrated (in the

next example), we show this by a graphical representation of the cost of what could be done providing free cash was available.



Keep in mind that the simple approach is not always the best one to use. Assumptions made need to be agreed upon by not only those involved in the model and generation of results, but also those using those results for decision making. It is also appropriate to vet and align with other areas that produce or use similar information, such as the tax department or budget groups.

CASH FLOW AT RISK CALCULATIONS

We also use our cash flow calculations to represent the cost and impacts of our hedging program. The chart below contains the illustrative values we generated. These represent the worst case cash flow, given a 95% confidence interval, and our total expected (base case) cash flow. Since the upside opportunity cost is a positive \$30 million in this example, we could have either swaps or options that are well in the money. This number is arrived at simply by subtracting the expected cash flow excluding hedges from the expected cash flow including hedges. In the case of the cash flow volatility mitigation, we see \$140 million. This is the CFaR excluding hedges less the CFaR including hedges. This tells us that not only have we reduced our worst case by \$140 million, but that the hedge program is not costing us anything, and is in fact generating cash, as a result of the current market prices.

	Including Hedges	Excluding Hedges
Total Expected Cash Flow	\$1,090	\$1,060
Total Worst Case Cash Flow	\$680	\$510
Cash Flow at Risk	\$410	\$550
Upside Opportunity Cost	\$30	
Cash Flow Volatility Mitigation	\$140	

These are a couple of simplified numbers that can come from the much more complicated analysis that illustrates the effectiveness of the hedging program. Even if in the above case the opportunity cost was negative, due to premium costs on options or out of the money swaps and options for example, this does not necessarily mean it is not an effective program if the mitigation number remains high. Ensuring that the cost does not exceed the mitigation is important, although these numbers cannot be looked at in isolation. The corporate objectives of what we are trying to achieve and the potential mechanistic approach to hedging also impacts the effectiveness of our program.

SAS@/RISK DIMENSIONS (RD)

There are many areas of setup required in SAS@/RD before analysis can be run. These range from price data to methods for valuing instruments (e.g. MTM). In order to perform scenario and what-if analysis within SAS@/RD, we needed to create a process for loading additional data. The IT infrastructure and reporting will also be discussed in the following sections.

SAS@/RISK DIMENSIONS – RISK FACTOR DATA

It is important to identify which risk factors to set up and the desired granularity of those factors. For our implementation, we identified energy commodity prices, exchange (fx) rates, interest rates, and volatilities as the risk factors. Volumes were also considered, but the data put aside as it was not the primary focus of the analysis. Time buckets or risk factor groupings are also important. We chose to set up our risk factors based on energy commodity price indices and the monthly time bucket (e.g. Nymex January 2008). Other possibilities included the set up of curves on an annual basis, or grouping indices together. For example, in the natural gas markets, we could have opted to group all Rockies indices into one blended index. Multiple considerations led to our decision to set up each index individually (CIG, NWP etc...). Primary among these considerations were the business transactions themselves and the portfolio of instruments involved. Balancing accuracy, simplicity, and transparency was also a leading factor in the decision to take the groupings to the lowest level possible given the existing risk factor data available. Although the primary objective was to provide a consolidated corporate view of risk, the decision to keep groupings at a granular level allowed us to build in greater flexibility for future analysis. The price tag, of course, was increased complexity.

One important note is that all the risk factor data were brought over from the source system as all-in values, as opposed to basis values to another reference point. For example, instead of Nymex Jan 2008 being \$8.00 and Aeco US being -\$0.80, the Aeco US price was brought in as \$7.20 = \$8.00 - \$0.80.

Much of the data is loaded through a source system in which we already capture our data. This load is done daily and is part of the Extract Transform Load (ETL) process into SAS@/RD. One set of data that is not directly loaded is historical volatility. Instead, we load the historical set of forward prices for a defined number of days back (e.g. 60). This gives us enough data to allow SAS@/RD to calculate the volatility of our risk factors. The implied volatilities are loaded via the same ETL process and source system. Since implied volatilities are based on the option data available, there are the only two sets of forward implied volatilities that we get out in the market: Nymex for natural gas and WTI for crude oil. In order to apply this to other risk factors, which remember tie back to an index and monthly bucket (e.g. Nymex January 2008), we adopted a method of translation to other indices. The formula used is as follows:

$$\text{Other Gas Index Vol} = (\text{Nymex Implied Vol} * \text{Nymex Forward Price}) / \text{Other Gas Forward Price}$$

This method of translation essentially implies no basis volatility, or in other terms, a static basis value. There are many ways to model in a translation effect. These implied volatilities will be used in our PFE and CFaR analysis, while the historical volatilities will be used in VaR.

For the purpose of calculating CFaR, we also adopted a method of scaling based on time to expiry. This was meant to adjust the volatilities from yearly to align with the time period defined by the monthly bucket. This was done as follows:

$$\text{Adjusted VOL} = \text{Annual Implied VOL} * \text{Square Root} [(forward term - effective date)/(days in year)]$$

The annualized volatilities are adjusted to daily for PFE.

The naming methodology of risk factors is also important, as we found out after implementation. We adopted a standard of <price type>_<commodity>_<index>_<term>. Example: P_NG_NYMEX_JAN2008. At implementation time, there were 3 price types; "P" is all-in price, "I" is interest rate, and "F" is fx rate. Remember, in the naming of risk factors, there may be a large number to consider and the SAS@/RD system defaults to an alphabetical norm. Since we included a text month (JAN for January, for example) and a number year (2008 in the example), the sort order is alphabetical, not chronological. For example, JAN2008 should come before APR2008, but it does not due to the format of the name. In our system, we have over 10,000 risk factors. This means having to scroll through columns and lists in queries in order to find the right risk factor. We also experienced difficulty when parsing out the date, as we then needed to convert it from its text form. One alternative would be to do a numeric date of yyyymmdd

(20080101, for example). Another method would be to use the Julian date format representing the number of days from Jan 1 1900 (MS Excel default), or the SAS® date format which represents a number of days from Jan 1 1960. It is important to remember this when approaching the risk factor naming conventions as it can have a significant impact on coding, reporting, and ease of use of the system.

SAS®/RISK DIMENSIONS – CORRELATION DATA

Correlation data is one of the key requirements to a simulation-based analysis. The way we chose to calculate this matrix is through the COVAR macro in SAS®/RD in conjunction with historical forward pricing. We also used the 60 days of historical forward pricing as a base set of data for the correlations. This generates a large 10,000 by 10,000 approximate matrix of bivariate correlations. This large matrix results from the relationship of one risk factor to another, calculated based on returns (which come from the day over day changes in risk factor data from the source system). An example of this relationship is the correlation between P_NG_NYMEX_JAN2008 and P_CL_WTI_JAN2008.

The correlation data is used in several analyses. In the case of CFaR, a methodology of scaling was used, similar to that for volatility, to come up with a custom set of correlations meant to model them as we go forward in time. The formula used is as follows:

$$\text{Adjusted CORR} = \text{Square Root } [((\text{CORR})^2 * t1 * t2) / (t1 + t2)]$$

– t1 is the term of risk factor 1, and t2 is the term of risk factor 2

Many assumptions need to be understood and accepted for correlation data, such as the scaling methodology for CFaR outline above or the time period selected for historical pricing. For example, we could choose 60 days of historical forward pricing for our VaR correlations, or choose a year or more of historical forward data for valuations such as PFE and CFaR. This is discussed further in the Challenges and Initiatives section of this paper.

SAS®/RISK DIMENSIONS – PORTFOLIO DATA

Portfolio data, or transactions, is another component of the equation. Not only is it required for calculating analytics, it also drives the need for risk factor data. For every instrument, the associated risk factor data is required to complete any valuations or analytics. We chose to split the data by instrument type to align with our pricing methods discussed in the next section. The data is loaded through the ETL process from the source system. We created custom code which, through the ETL process, appended specific fields and translated data that was needed for analytics later. One example where the custom code was required was the grouping of transactions that were not available in the source system.

Before we could group these transactions, we needed the ability to load data that was not coming from the source system. This gives us the ability to do what-if scenarios for potential hedging strategies. This entailed the setup of csv files that contained the data fields and formats required for each instrument. In the case of options, we modelled in the Black 76 valuation model to determine premiums based on a day's volatilities and market pricing. This is an area that we are looking to revamp as will be discussed later in this paper.

There are several pieces to the portfolio data section: data sets, filters, lists, and portfolio files. First are the data sources themselves, which are split by instrument type as previously mentioned (fixed swaps and physical index for example). The segregation of instrument data in the source system, and the different data fields required for different instruments, were two major drivers in setting the data up this way. Further we needed to set up an input list, which essentially pulled all of these data sources together into one for RD to use. We then set up filters to slim down the data for different types of analyses. The final piece are the portfolio files, which create independent data sets based on the list (data sources), and filters.

Within the actual Project setup within RD, a filter can be added, so what is the benefit of creating separate portfolio files? There is only benefit for those Projects that are running simulations. If the risk factors are simulated based on the required risk factors from the portfolio file, then the portfolio is filtered. For example, if we passed in a portfolio file with a fixed swap transaction from 2008 to 2010, all risk factors from 2008 to 2010 would be simulated. If we added a filter to the Project that looked for just 2008, this would not reduce the simulation time for the risk factors. However if we created a new portfolio file using that filter, the data passed in is only for 2008 and therefore only 2008 risk factors are simulated. The benefit is less needless processing time.

SAS®/RISK DIMENSIONS – PRICING METHODS

The pricing methods are used to value the instruments employing the risk factor data. The initial valuation that we set up was MTM. In the case of options, we incorporated the FEA module and passed instrument information into FEA to return the appropriate value. In the case of all other transactions, we valued the pay and receive sides of an instrument using the risk factors identified on the instruments, and any conversions required (fx, GJ to MMBTU etc.). The best way for us to set these up was to follow the data (instrument) breakdown. Given that each instrument type contained different data and valuation methods to get a MTM value, it was best to silo the code in the same manner. These two pieces are then tied together using the instrument type system variable and referencing the pricing method.

There is a system variable called `_Value_` and it is used by RD as the base valuation for PL Curve and scenario analysis. The only way to run a scenario or PL Curve is to use this variable. This means that if there is more than one valuation being done, there needs to be conditional logic created to pass whatever valuation needs stressing into the `_Value_` variable. After exploring a few ways to accomplish this, we chose transformation sets. This is discussed further in subsequent sections. By creating computed output variables, we can run simulations on more than just the `_Value_` variable. Simply add these variables as additional base valuations for simulation.

A lot of time was spent creating and tuning these pricing methods. We will be discussing how these have evolved and our current state later on in the paper.

SAS®/RISK DIMENSIONS – ANALYSIS AND PROJECTS

The Analysis and Projects are used to generate the results based on the base data (Risk Factor and Pricing Method). The Projects pull together several pieces including analysis, data, transformation sets, cross classifications etc. Not all pieces will be discussed here, but cross classifications allow groupings to be set up and this generates the summary results at these aggregate levels. The Analysis is where to define what type of analysis to run, whether it be a Monte Carlo simulation or a PL Curve. This is also where additional output variables can be pulled in to simulate and select the option to generate credit exposure (PFE). The Analysis area also provides the ability to define other parameters such as time horizons (1 day for VaR, several for PFE).

The Project section is where we specified our market data (risk factors), data portfolio, analysis, and correlation data. As discussed earlier, we have different correlations for PFE as well as for CFaR and similarly different portfolio of data (instruments). Output options can also be specified in this area. Be sure to understand what options are selected and the output that will be generated. In the case of PFE, with our current portfolio, if all output options are selected and because 60 time horizons are being examined, almost a full day would be required to run the analysis. This is not feasible if results are expected on a daily basis.

Take the time to isolate the analysis settings and the output data that is required. For example, any extra data, aggregate levels (classifications), or time horizons (simulations) will only increase processing time and disk space usage needlessly and reduce the amount of beneficial analysis that can be generated in a timely fashion. This will also be discussed further in the Challenges section.

SAS®/RISK DIMENSIONS – IT INFRASTRUCTURE

SOFTWARE

SAS® Risk Dimensions is a suite of SAS® software running on UNIX and Windows platforms. It includes most of the SAS® BI platform including the BI mid-tier components. Initially, our risk analytics relied on daily batch runs and did not involve the mid-tier. These batch runs consist of approximately 10000 lines of SAS® code and generate about 24 GB of results per run.

We have evolved to using more mid tier components: Web Report Studio, Web Report Viewer, the Portal and OLAP, and are running them separately from the core SAS® programs.

Change management practices at EnCana require separate environments for development, QA and production. CVS is used for source code control.

HARDWARE / OS / NETWORK

Base SAS® components hosted on a Sun Microsystems Sun Fire V440

4 X 1.6 GHz

16 GB RAM

64 bit Solaris 10 – virtual servers. All three environments are hosted on the same physical server as three virtual servers.

Network file storage

Prod 600 GB. 14 days results kept

Dev 250 GB

TQA 350 GB

Mid Tier components hosted on a Sun Microsystems Sun Fire V240

2 X 1.5 GHz

8 GB RAM

BI clients scripted and deployed to Windows XP desktops.

SAS®/RISK DIMENSIONS – REPORTING

Obtaining the data results and formatting them so that they could be used for management reporting and decision making was a requirement identified. We initially used somewhat cumbersome methods to obtain the results. We simply made csv export files, data sets combined with code, and pulled data directly into other reporting tools from SAS®. As a result, we needed to make a lot of presentation adjustments and relied heavily on typing in information and basic copy and paste. The efforts were not initially focussed in this area as our timeline for implementing SAS®/RD was quite short (a few months).

Once SAS®/RD was in place, and the analytics were in demand by management, the need to find a quicker, easier and better way to handle this was required. We adopted the use of Enterprise Guide (EGuide) and the Microsoft Excel addin. This allowed us to pull data directly into Excel in either a data set or pivot table form, and then use that to create several simplified charts for management. In the case of EGuide, we found it very useful to compile many SAS®/RD results and massage them for our use. The ability to link into other data sources for external data within EGuide was also used. We still use SAS® code to do some of the more static combining of data, but for the most part we use EGuide. The advantage to using EGuide is that while creating a Project which emulates a one step process that can be done in code, we also have the ability to break this process up into stages. For those with less coding experience, this tool provides greater transparency in the results and the reporting of them. Using this tool, we can isolate issues at any level in the data. Even if there are no issues to resolve, we still benefit from a one-click run of the entire process.

CHALLENGES AND INITIATIVES

Throughout the discussion thus far on the implementation of SAS®/RD at EnCana, we have referred to the challenges we have encountered. The challenges that we faced during the implementation, and continue to face as the use of SAS®/RD grows, affect both IT and the business users. Although discussed separately below, they go hand in hand.

In light of these challenges, we will examine some initiatives that have been completed as well as some that are planned. Many of these initiatives resulted from attempting to address the challenges that we will discuss. The goals of the initiatives are to grow the analytics in alignment with objectives, improve accuracy in the model, and maintain simplicity with ease of explanation.

IT CHALLENGES

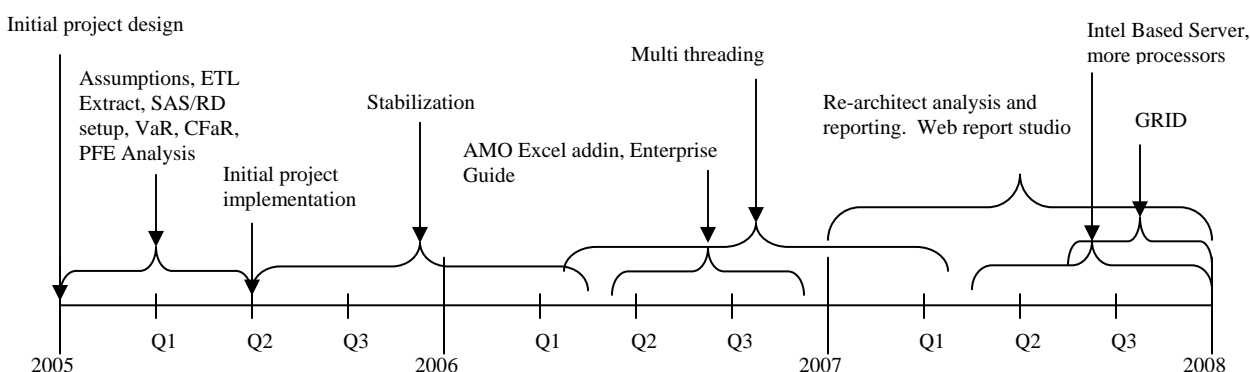
From an IT perspective, one of the growing concerns is how to balance the business need for increasingly sophisticated analytics with processing performance. We have to be ever more wary of what analysis we can run and how much output we can generate from both a processing time and disk space usage perspective. Our nightly run currently takes 9 to 11 hours and consumes 13Gb. Upcoming changes will increase this to 24Gb. In an analytics environment, the more input data we have and the more analysis results that can be generated, the better we can mitigate the financial risk. However, this takes more processing time and more disk space. The challenge we face is

how to address this performance and resource issue in order to grow our analytics?

Another IT challenge is reporting and maintaining data. The goal is to ensure simplicity in both the loading of additional data into the system, and the quick retrieval of information in a format suitable for presentation. As mentioned earlier in the paper, we discussed the inclusion of what we call “what-if” data. This is currently in a csv file import process that is somewhat difficult to use and maintain. We also use an Excel SAS® code-based import process for risk factor data. Our reporting process has definitely improved with the inclusion of the Microsoft add-in for SAS® and Enterprise Guide, but we still rely heavily on Excel for pulling the final pieces together. Part of this is to the result of not having all of the data needed in our model available in a single source. We are also much more comfortable with using familiar tools such as Excel for charting, and the audience is also familiar with seeing data presented in this way. How do we incorporate the additional data required into SAS® without an expensive development effort of a front end? Can we expand our knowledge of other reporting tools to present data in a more timely fashion?

IT INITIATIVES

The following timeline contains the initial project and IT initiatives a reference:



- **Multi threading:**
This splits the processing of intensive calculations such as PFE and Correlation matrix. It takes advantage of the multiple processors and cores available on our servers and reduces the run times.
- **Analysis & Reporting with BI clients**
We have been re-architecting from a purely batch environment to support BI clients like Web Report Studio, Web Portal, OLAP, Enterprise Guide and Access for Microsoft Office. These tools provide wider and easier access to the Risk Dimensions environment.

We are using SAS® Management console and the metadata server to manage security, library assignments and the autoexec processing,

- **Intel – based server, more processors**
We are testing a Linux-based Risk Dimensions environment on Intel. We are currently testing this configuration which runs 64 Bit Redhat Linux on a virtual server:
Hardware: HP ProLiant DL585 G1 or G2
CPU: AMD Opteron Model 885 (2.6 GHz), 4-CPU's (dual core)
OS Version (host OS): VMware ESX Server 3.01, 64-bit
OS Version (guest OS): RedHat Linux Enterprise Server 4, 64-bit

BUSINESS CHALLENGES

The first challenge we faced during the implementation was to define the objectives for our analysis. Before we could define objectives, we needed all levels of management to develop an understanding of the analytics we were generating. As mentioned throughout this paper, the challenge was to produce analytics that balanced accuracy with simplicity and ease of explanation. It is of course not realistic that we sit down with upper management and go

through all of the information in this paper with them. We had to develop a model that could be explained succinctly, on the “back of a napkin”, providing reconcilable numbers to move it from a conceptual example to a real example in our company. This was the first challenge in getting objectives outlined. Once we had successfully communicated the model and analytics, we needed to define objectives. To narrow the focus, we proposed a few areas of analysis, including Project Economics, Acquisitions economics, Cash flow maintenance and Manage equity markets. Based on the analysis discussed earlier in the paper, cash flow maintenance was an area that was selected, and the model enhanced accordingly. This is an area that continues to be reviewed and accepted.

We wanted to maintain flexibility for alternative valuation methods, given that this area, especially from a corporate perspective, was quite new in the company. We started by looking at MTM, and moved quickly into a revenue valuation as discussed. We have since explored and implemented other valuations such as discounted MTM and Credit value (MTM and AR/AP). What about other credit requirements? What about groups in the company that could benefit, like our budget group for example? The challenge was to keep the model open to incorporate new valuations for analysis, while enhancing the model based on the objectives already identified.

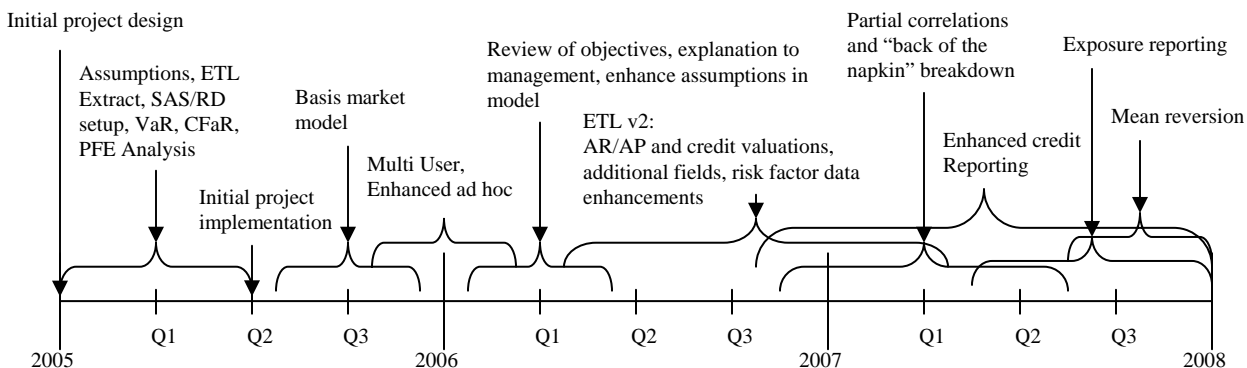
Most of what we have done has been strictly focussed on a simulation, scenario, or PL Curve based analysis. This only covers part of what is available to us with the use of SAS®/RD. How can the existing analysis be used more effectively in decision making and what other analytics, either alone or in tandem, can we use?

The question was raised whether we should rely on a purely market based analysis and incorporating some fundamental view of the market. One way we moved in this direction was to use the implied volatilities rather than the calculated volatilities using historical forward pricing; these are still being used in VaR. This prompted us to explore the base of our model, which led us to examine the correlation matrix. We understood that the correlations that were generated were bivariate, meaning a one to one relationship. But what about other risk factors that impact that relationship? What are some other ways that we can put a more fundamental view in our model? Our fundamentals group is responsible for making hedging recommendations and transacting based on that. It is important that our approach to the analysis used to make these decisions is accepted by that group as well.

Many assumptions were made with market data (risk factor data), that we reviewed and revisited. Was the 60 days of history enough for historical volatility and correlation calculations? If not, how much historical forward data do we have? In the case of the Crude market, how do we produce forward data for stream and field indices where there is not a published forward curve? How do prices correlate? Should fx have a correlation to gas prices? Gas prices to crude prices? What about the implied volatility translation assumptions? We scale correlations for our CFaR model; is this an appropriate methodology for what we are trying to show? What about the differences in forward data versus settled data? In the case of CFaR, we are looking at cash flow but using forward data to come up with correlations. The challenge here is ensuring the best return for effort. It would be great to tackle all of these market data challenges, but we need to balance accuracy, simplicity and the ease of explanation.

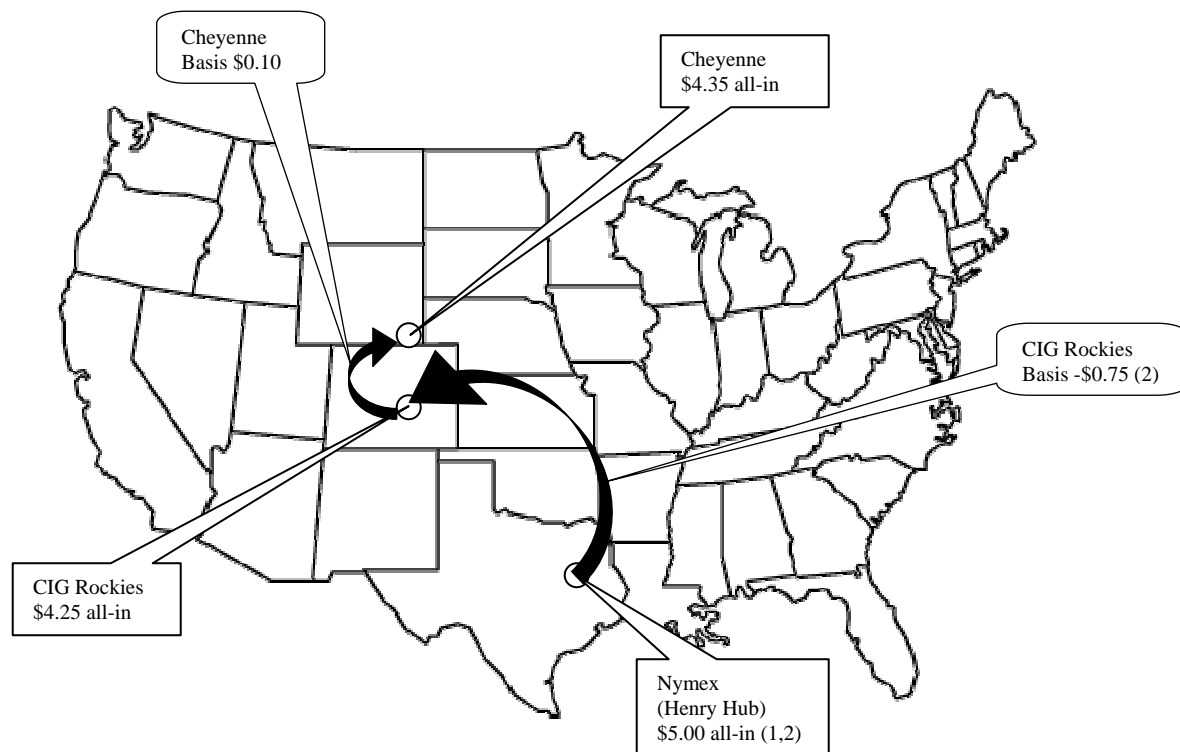
BUSINESS INITIATIVES

The following timeline contains the initial project and business initiatives as a reference:



One of the first initiatives that we tackled from the business perspective was the incorporation of the basis market model into the system. When working with PL curve and scenario analyses, we found that it was difficult to stress several risk factors. For example, consider the following natural gas price index chain:

NYMEX – CIG Rockies – Cheyenne Location Differential



These indices' associated risk factor (monthly bucket) data is stored as all-in prices. This results in two ways of examining the data:

- 1) ALL-IN (SAS®/RD setup)
 - Nymex = \$5.00
 - CIG Rockies = \$4.25
 - Cheyenne diff = \$4.35
- 2) BASIS
 - Nymex = \$5.00
 - CIG Rockies = $-\$0.75$ ($\$4.25 < \text{CIG all-in} > = \$5.00 < \text{Nymex from above} > - \$0.75 < \text{CIG basis} >$)
 - Cheyenne Diff = $\$0.10$ ($\$4.35 < \text{Cheyenne all-in} > = \$5.00 < \text{Nymex from above} > - \$0.75 < \text{CIG basis} > + \$0.10 < \text{Cheyenne diff} >$)

In a scenario analysis, if we wanted to stress Nymex by \$0.50, we would also have to stress the CIG Rockies and Cheyenne diff all-in prices. If we wanted to also stress CIG by $-\$0.10$, we would need to stress Nymex by \$0.50, CIG and Cheyenne by \$0.50 and by $-\$0.10$. For this example, we are assuming one risk factor per index, and only two price stresses. It would be much easier if we were able to stress Nymex by \$0.50 and have that flow through to the other prices. The complexity of stressing these all-in risk factors increases when moving to more than one risk factor per curve and also when looking at a PL curve analysis where several scenarios are done at once.

This led to the basis market model design and implementation. As mentioned earlier in the paper, using transformations was the simplest approach to solving this problem. A series of formulas were created, stating for example:

Nymex all-in = Nymex all-in
 CIG all-in = Nymex all-in + CIG basis
 Cheyenne all-in = Nymex all-in + CIG basis + Cheyenne basis

This formulaic method of calculating the all-in price when valuing an instrument through its risk factor, uncovered an important question: where do the basis values come from? In the extract process, logic was included to back calculate the basis values from the all-in values. We then had to create a risk factor for each of these. The risk factor naming convention was the same as outlined earlier in the Risk Factor section, with one exception. The preceding "P" was replaced with a "B" indicating basis price (B_<commodity>_<index>_<term>). We were now able to select a "B", or basis curve, from within a scenario or PL curve analysis and stress that specific curve. The transformations in place would then handle the propagation of that risk factor stress onto the rest of the market model in the chain. This is also controlled via the Project, where the inclusion of the transformation method can be specified. We discovered a technical limitation in RD as a result of creating these many transformations (one for every risk factor, which can be in the range of 5000). SAS® technical support worked diligently to discover the root cause, and have put in for a hot fix as of February 2007. This fix may be released before the SAS® Global Forum 2007.

One requirement that was not yet mentioned in the project was the ability to run ad hoc analysis. Essentially this entailed running the nightly process to extract the data from the source system, adding what-if data, and generating the volatility and correlation data within the parameters for term of data and history date. Since the initial focus was the regular nightly run, the whole system was set up to run in a single user mode. Once we considered rolling out the analytics in a customizable way to more than one user and other groups, we needed an enhancement for a multi user mode. This would enable us to process more than one run at a time, provided resources of processing and memory were available. From the ETL process from our source system, to logging into SAS®, we needed many enhancements. We took this opportunity to create a consistent file structure for storing the data and results for the nightly runs and user specific ad hoc runs. It was prudent to do this given the volume of data we were storing and the widening user community.

Our most recent initiative, which supports the growth of our analysis and reporting, was the ETL v2. The goal was to enhance our data source extract from an accuracy and performance perspective. We also required additional data to be brought into SAS®/RD for enhanced credit and exposure reporting. Many of these fields related to contracts, counterparty information, and credit ratings and limits. Initially, there was a strong focus on ensuring the market risk related analysis and reporting was in place. With good progress in that area, we added the focus of our credit risk analysis and reporting.

Given this added focus, we needed to identify what valuations in the pricing methods would be required to better enhance existing analysis such as PFE. We also wanted to replace existing manual reporting processes. We needed to isolate the components of PFE: AR/AP, MTM, and CVaR (Credit Value at Risk). In the case of MTM and CVaR, we already had these pieces calculated, both for credit and liquidity exposure. The challenge was obtaining the data in a timely fashion. Changing our options on our PFE Project to produce the simulated results at P5 and P95 would have increased the processing time. Being able to show the credit and liquidity in a breakdown of AR/AP, MTM, and CVaR allows for better credit risk reporting that enables us to see what makes up our credit and liquidity with counterparties at specific points in time. Ultimately, we will be able to put limits and guidelines around this data, similar to VaR, in relation to our counterparties.

In the case of AR/AP, we had not initially set up valuations in SAS®/RD. We adopted a 60 day rolling AR/AP methodology that introduced some assumptions of when AR/AP rolled in and out. A natural gas physical contract settles on the 25th business day of the following month, taking into consideration weekends and holidays observed in both US and Canada. To keep 60 days rolling, the current month instrument value needs to roll on for AR/AP when the 2 months ago position rolls off on the settle date. Traditionally, AR/AP would roll on when MTM rolls off (index settles) and gas starts flowing (cash month), but given the 60 day rule we needed to incorporate those assumptions.

The ETL v2 initiative also makes way for exposure reporting out of the system. We currently have a series of Excel macro spreadsheets which facilitate our exposure reporting. The required data, production and hedge volumes, exists in SAS®/RD, but is being generated and maintained outside of RD in Excel. The missing pieces for the reporting are certain instrument groupings, which are now coming across into SAS®/RD as part of ETL v2. The importance of reporting exposure out of SAS®/RD, aside from timeliness, is the ability to produce both the reports and the analysis from the same system, using the same reporting tools. This addresses the challenge of presenting the data in a consistent and easy to follow manner in order for management to make decisions.

We started to investigate using Principal Component Analysis (PCA) and partial correlation methodology. For partial correlations, we examined rolling the impact of risk factors on the correlation between two other (bivariate) risk factors. This entailed validating how much historical forward data we were pulling. We came to the conclusion that 60 days of history was not enough. The investigation also allowed us to look for ways to encompass more of a fundamental view into the correlation process.

The subject of historical forward data deserves further attention here. We realized that for our CFaR and PFE models, a 365 historical set of forward prices was more appropriate to use given our reporting terms and time horizons. This would also improve our ability to incorporate partial correlations which will be described shortly. How was increasing to 365 days of history going to affect our analysis? Re-running past analysis using the new settings, we were able to see less consistent results when changing market data only, holding cash flow deductions and instruments constant. Prior to the end of 2006, we began an initiative to include crude forward field price indices (risk factors), into the source system from which SAS@/RD obtained data. Up until this point, we had simply updated WTI and stream pricing. There are no published forward field prices for our use, so these were derived from an internal price model and forecast that incorporated items such as transportation and diluent costs. When increasing to 365 days history, we pulled in approximately 275 days worth of data that was not accurate to the market based on the internal forecast and model, introducing an improper increase to our correlations. Going through the process of making the 365 days historical forward data accurate improved our consistency from $\pm 6.7\%$ to $\pm 1.7\%$. This historical data is part of the foundation, so when using it to generate correlation and volatility data, it is extremely important to ensure its accuracy.

Once we had resolved the historical data accuracy issue, we could continue with the partial correlations. We generated the giant 10,000 by 10,000 matrix of correlations encompassing all of our risk factors. Identifying the key risk factors that represented most of our instrument data, we were able to slim this down to a few hundred – 10 to 30 indices. For natural gas, we considered grouping gas years together (November to October), or seasons (Nov-Mar for winter and Apr-Oct for summer). Since our analysis typically rolled up to the year, we decided that putting the years together would be appropriate going out 5 years. This created 5 groups of risk factors (10 to 30 indices multiplied by 12 monthly buckets in the year). The next phase was to run a series of regressions on the smaller subsets of historical forward data. The number of regressions depended on how much historical data was available to use. To calculate the number of risk factors that can be run through the regression at one time, the theory was to take the square root of the number of days of history. The square root of 365 is 19. The square root of 60 is 7, so imagine how many more regressions we would have needed going 7 risk factors at a time. This process is being fine tuned and at the time of this paper is more theoretical.

As an example, having 5 groups of 120 risk factors each and being able to run 19 at a time in a regression, works out to 6 or 7 regressions. Right? Unfortunately, this simply runs a regression on risk factors 1-19, 20-38, 39-57, 58-76, 77-95, 96-114 and 115-120. What about running a regression on risk factors 1-5 plus 23-26 plus 40-42 plus 100-106 (19 in total)? Enough regressions are needed in order to capture all the possible combinations, or at least reasonable combinations. Some risk factors may get thrown out due to little or no impact based on a high correlation to another risk factor. A convergence program was developed to come up with the ideal number of iterations; this took 15 days to run and would be best to run once the risk factors have been chosen for the groupings. Hundreds would probably be in order. A random shake, as we named them, was our solution to this problem. We ran the regressions to get a set of covariance values, randomized the order the risk factors were in and ran it again. The process repeated this to the number of ideal iterations, and averaged all the covariance values to get one final new covariance value for a pair of risk factors. All 5 groups were run through this process. Finally, the bivariate values in the larger original matrix were replaced with the new values. As we have sub-selected risk factors that we deem cover most of our portfolio, many of the values are not replaced with new values. The time it would take to do this process and generate a 10,000 by 10,000 matrix would be in the order of several days.

The result of this partial correlations code is the ability to incorporate impacts of selected risk factors on other risk factors (multivariate correlations). If we fundamentally thought that the price of pork bellies was a driver behind gas and crude pricing, we could pull pork belly prices into a grouping with Nymex and WTI to generate correlations that would reflect this fundamental relationship. At the time of this paper, we were in the initial testing phase with a potential production release before SAS@ Global Forum 2007.

The “back of the napkin” breakdown of our analysis, which was discussed earlier, is another initiative on which we continue to work. Since we start at the granular level and roll our analysis up, we are looking for a way to produce a “slim” model. When producing the results and providing transparency to the data behind it (e.g. simulated risk factor values), we also need to provide a way to make small tweaks and see the impacts without running through the entire process again. For example, if, for the year 2008, our analysis indicated that at a \$6.00 Nymex price, our WTI crude price was going to be \$85 a barrel, our fundamental group would have a problem with accepting that value. Keeping everything else constant while adjusting this WTI average price to \$50, what will be the impact to the results? This “slim” model is just a rough guideline, but will give us the ability to adjust our assumptions after the analysis is run. Using the exposure reporting, simplified breakdown of the analysis, and the cash flow assumptions, we can generate a cash flow impact on making these types of post analysis tweaks. Back-testing this process is of key importance to come within an acceptable variance of what the actual full model would generate for results.

Finally, another key initiative we are considering, but have not dug into at this point, is the concept of mean reversion.

Our simulation analysis assumes that prices (risk factors) continue to move up and down when moving forward in time. There is a fundamental argument that this is not necessarily true and the prices will come back towards an average value over time. Given our analysis can extend for many years, mean reversion is another way we can encompass a more fundamental view in our analysis.

TO WRAP UP

Any implementation of risk analytics at a corporation is an ever evolving process with continued explanation to management, review of objectives, acceptance of analytics, and growth in results generated and reported. One important theme we carry through from implementation to current and future initiatives is that accuracy needs to be balanced with simplicity and ease of explanation. If the model(s) cannot be explained and accepted by management, then the results that are produced will not be readily accepted. The goal is to continue to produce improved analytics and reporting, in line with objectives, to make informed business decisions.

Take the time to truly validate and understand all assumptions in the setup of the system and analysis/models that are created. Align with other groups within the company working with similar data to insure inputs and assumptions do not contradict other areas, making it easier to gain acceptance from management. When looking at an implementation of risk analytics, IT infrastructure is important. It is not always a matter of throwing more resources at the problem, and to grow in the area of risk analytics may require compromise and ingenuity. Analytics produces lots of data and consumes much processing power, generating the need for proper planning and foresight for possible exponential growth.

SAS® and Risk Dimensions have been invaluable tools in reaching our risk analytics goals within the company. What comes out of our financial model(s) is a reflection of what goes in. We hope this information is useful in pursuit of risk analytics initiatives. Our goal is to promote discussion and potential collaboration on what is still a relatively new concept to the energy and utilities industry. Any and all questions and suggestions in regards to this topic are greatly appreciated.

ACKNOWLEDGMENTS

Dominic Pazzula – SAS – SAS/RD setup and coding
Akbar Shahmoradi – SAS – Partial correlation methodology
Scott Dalton – EnCana Corporation – Advisor, Market Risk
Derrick Boyd – EnCana Corporation – IT Infrastructure

CONTACT INFORMATION

Your comments and questions are valued and encouraged. Contact the authors at:

Name: Gord Anderson, Market Risk Analyst, Financial Governance
Enterprise: EnCana Corporation
Address: 1800 855 2nd Street SW
City, State ZIP: Calgary, AB CANADA T2R 0K4
Work Phone: 1-403-645-4353
Fax:
E-mail: gord.anderson@encana.com
Web: www.encana.com

Name: Mark Perrin, P.Eng., IT Group Lead, Financial & Market Risk
Enterprise: EnCana Corporation
Address: 1800 855 2nd Street SW
City, State ZIP: Calgary, AB CANADA T2R 0K4
Work Phone: 1-403-645-3939
Fax:
E-mail: mark.perrin@encana.com
Web: www.encana.com

SAS and all other SAS Institute Inc. product or service names are registered trademarks or trademarks of SAS Institute Inc. in the USA and other countries. ® indicates USA registration.
Other brand and product names are trademarks of their respective companies.