

1. Life after people

When asked previously to present, I have been hesitant as I felt that we were not done yet.

This was maybe because people kept on asking: "When will this project be done".

But then I realised that we will only be done, when business is done – actually some time after business is done as we will be providing the answers as to why the business does not exist anymore, I changed my perspective.

So, I think it is "safe" to say that we have a successful implementation.

2. Logo's

Before I go any further, I need to clarify who I refer to as WE.

The BI implementation in question is for Momentum Benefits At Work (BAW).

They are providers of employee benefits products and services, including retirement, risk and investment benefits; and administrative and consulting service.

I am Karen van der Walt, a BI architect from Alacrity (a software development company) and have been contracted to Momentum Benefits At Work (BAW) for over 7 years now.

I am responsible for the BI design, project management, business analysis, user support & basically all but the technical side.

The technical implementation is handled by BI in the Box. This BI consulting firm was founded by Clint Vosloo and employs BI specialists located in Cape Town and Johannesburg.

Specialist skills include Data (and specifically dimensional) modeling expertise, certified database administration (Sybase IQ and MS SQL server), ETL development, report development, expert SQL scripting and web development.

3. Schumacher on podium

When this project was initiated, I was fortunate to have a Business champion that really understood the reasons for and expected outcomes of a BI initiative.

She also had the vision to realise that embarking on such an initiative is a long-term commitment and should not focus on short-term fixes.

The project to setup a BI environment and implement a pilot was kicked off in September 2005.

4. Business Model

So there I was – without any resources or tools and desperate to get going. So I started at the point where no resources or tools were needed, and with something I was familiar with – Business Modeling.

Today I consider this starting point as one of the reasons for our ultimate success, but at the time I cannot remember having this end in mind when I started.

5. Business Executives

I was very fortunate to have the attention of the top management in designing the BM.

The theme of the workshops was WHAT information regarding WHICH entities is important as well as the RELATIONSHIPS between these entities.

The discussions were non-system specific and not subjected to any technology, systems, business, political (or even reality) constraints.

It was purely a documentation of WHAT defines the business and no reference was made to any specific reports that needed to be developed.

6. Report Mayhem

I believe that this placed the emphasis on the underlying data linkages instead of the generation of a multitude of seemingly loose-standing reports.

There was a general understanding that unless the data from the many disparate systems were transformed into meaningful information, no reports will be usable anyhow.

7. Detail pieces

Once the executive committee were satisfied that all aspects of the business were mapped, I met with each department in the business to fill in the detail.

Here the system and process specifics were added to the model.

I also documented the data mapping and additional requirements as it transpired.

With all the additional detail mapped, the then much enhanced business model were presented to the executive committee and signed off.

8. Ancient wheel

This was done in 3 months and was time well spent, as the business model still holds true today – 5 years later.

Even as systems and products were added and terminated.

Now this should not be a revelation, as a company's underlying business model does not change unless the business proposition changes radically.

It is important to note however that the BM was never envisaged to be implemented in one big never-ending project, but rather implemented in phases.

I use the original BM as 1st point of reference whenever functionality needs to be added to the DWH.

9. One size fits all / Hammer

This is probably a good time to mention that the modeling tool I selected is Sybase PowerDesigner.

The main reason I have selected PD is that it is an one-of-a-kind modeling tool as it is a business modeling and data modeling tool all in one.

From the Object Orientated BM, I generate the Physical Data Model and do the necessary changes. This Physical Data Model is then used to generate the DDL's used to implement the database tables.

So, with the WHAT sorted out, it was time to move onto the HOW.

10. Dinosaur

At that stage the technologies in use in the business were basically mainframe systems, some SQL2000 and an analytical system developed by the actuaries consisting of Essbase and DataStage.

As DataStage is the best ETL in the market, the ETL choice was easy.

So we only had to look for a database and Reporting Tool.

As the mainframe was not a consideration, the database technologies considered were:

- Hyperion (as Essbase was a current technology)
- SQL2000 (SQL2005 was not available then)
- Sybase IQ

11. Fast car

I got to know about IQ when looking at modeling tools and were really impressed with its performance and the fact that it is built specifically for analytics.

As you will see later on – we are not dealing with huge amounts of data.

We do however source data from multiple, disparate systems and the transformation and calculations on this data are quite complex.

Also from the onset we knew that a base requirement was providing analysis capabilities – and it was with this in mind that Sybase IQ was considered.

12. Last puzzle piece

The remaining BI component I had to choose, was the Reporting Tool.

As all the products thus far selected were supported by Sybase, I chose Business Objects (at the time also supported by Sybase).

Apart from BO having all the needed functions, having a single vendor really appealed to me.

We have subsequently (in 2009) migrated the BO reports to MS RS (after SAP purchased BO). We also use Dundas for graphs and dashboards.

13. Starting line

With all the “boxes ticked” we were ready to embark on the Proof of Concept. The aim of the POC was to test whether IQ could acceptably handle the analysis complexities.

The POC was done by BI Practice and consisted of the loading of 3 years of financial data and calculating profit measures on this.

The POC produced overwhelming results and the actuaries involved could barely contain their excitement!

Just to remain objective, Magmatec (who was at that stage responsible for systems development at Momentum), performed a similar exercise on SQL2000 – But needless to say – IQ was chosen.

We (BI Practice, a technical resource that has been assigned to the project and me) then did a pilot to test how all the technologies will work together. The pilot was done on the structural data (12 entities of the BM) and produced the desired results.

With the pilot completed by Jul 2006, it was time to plan the projects to implement.

14. Surfer Shark

But as the sole business resource responsible for profit reporting and analysis resigned, the decision was made for us – and it could not have been a more testing project!!

By October 2006 (I remember this date well as I went on maternity leave) the design and business specifications were done, signed off and development could begin.

Now from here I will speed up as I have only covered the first 8 months!

15. Stuck in the mud

Upon my return from maternity leave I found the implementation not delivering what was anticipated.

I found myself being assigned to another project altogether and being consulted from time to time.

By the end of that year (2007), the project was not going anywhere and I was made PM again.

Now I would like to think that the problem was me not being involved, but the problems unfortunately ran much deeper.

16. Jar of sweets

A reason for the project was not producing the desired results, was the data.

One of the actuaries correctly called it “the devil is in the data”.

An important aspect to remember when dealing with BI projects, is that the work is not done when the code does what the specifications required.

As we are in Data projects dealing with the “data waste” created by transactional systems, the project is only done when the last record of data has been handled successfully.

As it is impossible to analyse millions of records before coding can start, this phase (which I call Data Verification) is an important aspect to plan for.

17. Homer Simpson at work

Another reason for the project being stuck – and this is a critical success factor – was that the project was not resourced with the right people for the job.

I have found that someone that might be technically fine on other systems does not necessarily make a good data developer.

It is not just an attention to detail or understanding the repercussions of your actions that are important, but more an intuition on what is proper and what not.

For me it is a gut feel of which choice to make – probably a combination of theoretical knowledge and real-world experience.

18. Swiss army knife

For me the main characteristic of a performing team is having skill diversity and an agreed & accepted knowledge and respect for this diversity.

So often, people are employed for the specific skill/knowledge/attitude they can add, but once employed, they find themselves being Performance Managed to be clones of each other.

Maybe we escaped this due to the fact that we are all contractors and our performance is solely rated on output. Luckily for us, this works.

Getting back to the implementation –

I was fortunate to have had the Business champion's support to re-resource the project.

But before any new resources were added, we needed something extreme to get back on track...

19. Extreme makeover / Before & After / Chaos & Order

...and a project (I called it Project Extreme Makeover) was initiated.

The BI Practice performed various health checks on our implementation and presented us with quite far-reaching feedback.

From these findings we (Clint and I) constructed various interventions that included:

- Technical stabilization
- An implementation methodology
- Best practices for each technology used
- Resource plan.

By mid 2008 all our plans were in place and resources were on board.

The majority of that year was spent on fixing (and mostly redoing) the profit report and analysis system.

We he also added a Productivity system and Transformation system based on workflow and payroll data.

20. Data Dump

In 2009 the business environment changed when the 3 EB business units merged to form a single EB BU.

As we were suddenly being confronted with various “foreign” systems based on a similar BM, I decided to store this data in an one-to-one mapping from source on the DWH.

We called these Data Stores.

This turned out to be a good decision as it sped up the transforming of the source data into the BM tables.

We had the luxury of available technical resources and while the guys were adding this “raw” data to the DWH, I could concentrate on the critical data items that had to be loaded into the tables representing the BM.

This worked extremely well as we were able to anticipate what the business needed and produce almost instantly.

We ended up having Data Stores that are used to update the Business Model that is the source for specific Data Marts.

21. Compression

Now this would not have been possible without Sybase IQ's compression capabilities.

People with implementations on other technologies cannot believe that we have the same data items in the Data Store, then transformed into the BM tables and then eve have multiple data marts running.

Due to compression we are utilizing 47GB of data where it would have required 141GB in another technology.

This translates to a compression ratio of 33%

22. Bulk Loading

Sybase IQ's Bulk loader facility is another important aspect that allows us to store all these data.

For example, our two largest tables are currently updated in full-refresh mode:

- 155m rows loading at 40,000 rows per second with a 60% compression ratio (8.5GB)
- 98 m rows loading at 14,000 rows per second with a 11% compression rate (7.5GB).

23. Master chef

Another advantage of utilizing Sybase IQ, is that we are not forced to choose a particular design purely for performance.

As we are implementing a real-world solution to overcome real-world challenges, it is necessary to utilize all that is available depending on the situation.

We have therefore both relational and dimensional tables in a single implementation.

This is possible as a result of Sybase IQ's column-based indexing that facilitates fast access to all stored data – regardless of the underlying design.

(If you wonder about this picture – it was added since I compare people sticking to a theory to a cook that merely follows a set recipe.

This opposed to a master chef that knows what works when)

I have said before that our emphasis is on linking data items that might have no apparent connection.

And with sourcing from various systems where the update frequency and mode differs vastly,

24. Winston Churchill quote

we have run into referential integrity problems when adhering to (theoretical) correct design standards.

But luckily for us, with Sybase IQ converting keys to integers (even if it is stored as a character),

we are able to link data using meaningful foreign keys, instead of being forced to use auto-generated integer keys.

25. Chimpanzees

I have mentioned the right resources for the job before, but another important resource aspect is the utilization of the resources.

As Sybase IQ is not admin intensive at all, the certified IQ administrator on site is able to allocate 90% of his time to development.

We also do not need an expert data modeler – even I do data modeling!

26. For the technical minded

Where we are now is:

- Running Sybase IQ 12.7 on Linux
- Running DataStage 8.01 Server Edition on

27. Solution size

In order to support:

- 70+ Business users
- 725 tables
- Utilizing 47GB (without compressed GB)
- 272 reports
- Sourcing from 18 systems (7 different technologies)
- 791 DataStage jobs

28. In summary

Implementing the chosen technologies allowed us to focus on the business requirements without it becoming all about the technology.

Other important factors were:

- High level business buy-in from the start
- Having the BM as reminder of what the business is about
- Having the right resources and applying skills effectively and efficiently
- Having an anticipation of what is needed to generate business benefit

- Having the opportunity to fix what was not working – and through this implementing the needed BI governance.

29. Butterfly puzzle

But, ultimately I believe that the implementation is sustainable as it centres on the data connections rather than the generation of reports from unconnected data.

This “making something from seemingly nothing” made me think of a joke (I am sure everybody has heard before). But I will tell it anyway as I find it so relevant.

30. Kellogg’s rooster

And this is exactly what we are doing – Building colourful roosters from corn flakes!