

What's a semantic layer and why would I want one?

An increasing number of reporting tools allow users to make use of a semantic layer.

IPL's Chris Daniels examines the purpose of semantic layers and the benefits they deliver.



Introduction - the need for reporting tools

Perhaps, like me, you've noticed that software vendors can successfully sell computer systems which record data, without also providing facilities to make that same data available for reporting and analysis. Inadequate reporting capabilities seem to be a common characteristic of both off-the-shelf and bespoke systems. It's a real shame because it can prevent firms from exploiting the full potential of the information they hold.

Some of the reasons for this situation are as follows:

- Reporting is often seen as a secondary function by those purchasing or specifying systems. Indeed the reporting capabilities of a product are sometimes overlooked altogether when procuring a system.
- Reporting is often seen as a secondary function as far as vendors are concerned. Vendors may lack the expertise and the interest to produce reporting capabilities.
- The reporting needs of different customers can vary considerably and may be difficult to predict. Consequently vendors of off-the-shelf products will typically invest in "lowest common denominator" reporting only.
- Line-of-business systems are designed and optimised to support the transactions and events that occur every day. They're not built with flexible reporting and analysis in mind,

with the result that it may be inherently difficult to support such features.

Given these factors, it is inevitable that computer systems will continue to be delivered without adequate reporting facilities, and that organisations wishing to make the most of their information will need to find ways to plug the gaps.

Plugging those gaps might involve anything from creating of a set of standard reports to the creation of a data warehouse and a suite of analytic tools.

Whichever course of action is taken, the basic requirement is that business users need to be able to get data out of their systems to answer burning questions - questions such as "how much revenue did we receive last month from sales of baked beans?" To do that, they'll need some reports, and it's very likely that these will be created using a reporting tool.

Reports use queries to extract data from data stores, and then present the results to a user. At a low level, there are two main ways of creating queries. The first method is to use a query language appropriate to the data store, e.g. SQL or MDX. The second method is to use a semantic layer...

What is a semantic layer?

An increasing number of reporting tools allow users to make use of a "semantic layer". Some vendors give the semantic layer a name that is specific to their particular product (for example, Business Objects calls it a "universe") while others describe it as a business model or metadata layer.

A semantic layer is a business representation of corporate data that helps end users access data using common business terms. The aim is to insulate users from the technical details of the data store and allow them to create queries in terms that are familiar and meaningful.

How is the semantic layer created?

The semantic layer is configured by a person who has knowledge of both the data store and the reporting needs of the business. The person creating the semantic layer chooses to expose appropriate fields in the data store as "business fields", and to hide any fields that are not relevant. Each business field is given a friendly, meaningful name, and the business fields are organised in a way that will make sense to business users (they do not have to be organised according to the structure of the data store).

The person creating the semantic layer can also create pre-defined filters corresponding to common restrictions that users may wish to include within queries. For instance, a pre-defined condition named "Filter: Current Year" would restrict the results of a query to data for the current year.

How is the semantic layer used?

With a semantic layer in place, the user sees the data store not as a collection of tables but as a single list of familiar "business fields" which are organised into a meaningful folder structure. Users can specify a query simply by selecting a set of business fields and pre-defined conditions, and can see the results on the screen. Users can then use point-and-click techniques to customise the appearance of the report and the way in which the data is presented. For instance they can choose to have data presented in tables, crosstabs or graphs, and can add headings, text and images.

This technique is pretty powerful. It allows us to offer business users a "self-service" capability to create a wide variety of different queries and reports themselves, without the need to wait for developers to become available. Users can create reports without the need to learn a query language or understand the structure of the data store. From an IT perspective, this is sounding great – we can let business users create their own reports as and when they need them!

An ad hoc query environment that uses a semantic layer allows developers and analysts to work interactively with business users to prototype the reports that are needed.

Who would use the semantic layer?

Before we get too carried away, it's important to recognise one important point. **Most business users don't want to create their own reports!** They're much too busy dealing with other things. Reporting tools that utilise a semantic layer may be sold on the basis that they empower business users to create their own ad hoc reports, but the reality is that business users generally expect reports to be created for them.

So while the semantic layer is designed to be sufficiently intuitive to be used by business users, it will mostly be used by developers.

The case against semantic layers

Developers are generally comfortable with query languages, and may at first be suspicious of semantic layers and sceptical about the benefits they deliver. There is a learning curve associated with the creation of a semantic layer, and developers may feel that it would be quicker not to bother creating one in the first place.

Any initial scepticism can be reinforced when a developer tries to create a report that is just too complex to be created via a semantic layer. In such cases, it's not possible to delegate the task of query creation to the semantic layer, and the developer must instead create queries by hand.

As a result, it can be tempting for developers to revert to type and write queries by hand for each and every report.

Does this mean that the semantic layer technique is dead in the water? Certainly not!

Reasons to use a semantic layer

Semantic layers deliver a number of benefits, including the following:

- Firstly, an ad hoc query environment that uses a semantic layer allows developers and analysts to work interactively with business users to prototype the reports that are needed. This can be immensely valuable.
- Secondly, the presence of a semantic layer simplifies the creation of queries and increases productivity, accuracy and consistency between reports. This is clearly of benefit to all users who wish to create reports, not just business users.

- Thirdly, knowledge of the mappings between business fields and database fields resides in the semantic layer, not in the heads of developers. This makes it much easier for developers to understand and maintain reports that have been created by others.

- Finally, and perhaps most importantly, a semantic layer removes the dependency between queries and the data store. This is because queries that are built using a semantic layer are constructed (and stored) as sets of business fields and pre-

defined conditions rather than as statements in the query language. Consequently, a change to the data store can often be handled globally, without the need to modify individual reports. If you've ever had to modify the SQL within hundreds of reports to accommodate a database design change, or had to change reports to take data from a different source you'll know what a benefit this is!

Conclusion

There are real benefits to be gained through the use of semantic layers. Semantic layers are increasingly becoming available within all types of reporting environment and their use should be encouraged.



Chris Daniels is an Information Management Consultant at IPL, a leading UK IT services company specialising in the delivery of intelligent business solutions. He has over 10 years of experience in helping a wide range of high profile clients exploit the full potential of their information.

Chris has significant information management expertise gained across a variety of sectors, including manufacturing, finance, government, and telecoms. Chris has ongoing engagements in business intelligence, undertaking business analysis and the specification of information systems. He can be reached at chris.daniels@ipl.com.