

## Chapter 3

# The Business Intelligence Front-End

If the business intelligence *architecture* is like the engine of the car, then the BI *front-end tools* are like the body: sporty, sleek, fast, and where the razzle-dazzle of color, handling, and chrome finish all matter. You can have a perfectly architected data warehouse, and yet if you don't have the right BI front-end tools, you won't achieve business intelligence success. Technical capabilities matter here but so do subtle differences such as look-and-feel and ease of use. Conversely, while you can have a powerful, intuitive BI front-end, if you have not paid attention to the underlying technical components discussed in the last chapter, your initiative will fail and users will blame the tool for any underlying problems. You need to get both aspects right, even if it's only the tools that are visible.

This chapter describes the various BI front-end tools that are highly visible to business users. Chapter 12 discusses the importance of matching the tools with the right user segment and the role such tools have played in successful companies. As discussed in Chapter 1 (the section "Technology Changes Enabling BI"), vendors offer an increasing breadth of capabilities within one BI suite. Throughout this chapter, I will mention specific vendor modules to provide concrete examples. This list is not exhaustive and as vendors acquire each other and/or introduce new modules, specific names may change. For updated product names and modules, consult the BIScorecard web site.

### Business Query and Reporting

Business query and reporting tools are often referred to as "ad hoc query tools." This terminology is a little misleading, as in fact the queries are not always ad hoc (as in spontaneously crafted) but rather are often

## 36 Chapter 3

fixed reports. The difference is that a business user, usually a power user, may have built the report, rather than an information technology (IT) person. The business environment changes at a rapid pace, and unable to wait weeks or months for IT to develop a new report, business users often demand the ability to create queries and reports themselves. Business query and reporting tools allow for this and are most often used for decision-making and management purposes. The business query and reporting tool is a key module to provide users with self-service information access.

In some cases, a report is truly ad hoc; it's a one-off business question that will never be posed again. Ad hoc queries may be exploratory in nature as users try to find the root cause of a problem, test a theory, or consider changing a business model. Table 3-1 lists some sample fixed reports that may lead to an ad hoc query. As users explore the data, what started as an ad hoc query or one-time question may later become a fixed report. It's important to recognize the iterative nature of business intelligence and to ensure you have flexible business intelligence tools.

Getting to the data is just one capability of business query tools; the other aspect is presenting and formatting the data in a meaningful way, loosely referred to as *reporting*. The terms “query” and “reporting” are sometimes used interchangeably because a business query and reporting tool will have both capabilities—getting to the data and formatting it to create a report.

Fixed Report	Purpose	Related Ad Hoc Query
Inventory by Product	To determine if an order can be fulfilled today by the primary warehouse	If I'm short at my main warehouse, can I supply the product from elsewhere?
Top 10 Customers By Quarter and Product	To understand which customers generate the most revenue	Who fell off this quarter's list? Are there certain products we can cross-sell?
Raw Material Receipts and Delivery Times	To determine how long it takes to acquire raw materials and which supplier can fulfill purchase orders fastest	Are there other suppliers who can respond faster?
Patients Per Hour	To understand busy periods and wait times	Do staffing levels correspond to busy times?

**Table 3-1** Sample Fixed and Ad Hoc Reports

Business query and reporting tools vary widely in their formatting capabilities. The most basic of formatting capabilities allow for changing the font of column headings and making them bold and centered. Conditional formatting will, for example, display numeric values red when negative or below target and green when positive or above target. Simple report styles include displaying information in a cross-tab report, a chart, or a master-detail report with groupings and subtotals. Tools may provide a set of templates to create nicely formatted reports that use a consistent corporate look and feel. More complex formatting capabilities include the ability to present multiple charts on a page, perhaps coming from different data sources.

Examples of business query tools include BusinessObjects Web Intelligence, Cognos 8 Query Studio, and SAS Web Report Studio.

## A Business View of the Data

Business query tools allow business users to access a data source via business terms without having to write any SQL. The data source could be a data warehouse as described in Chapter 2, or it might be direct access to an operational system. A key feature of a business query tool is that it has a business view or metadata layer that hides the complexity of the physical database structure from the business user by:

- Using business terminology rather than physical field names. For example, a user may select a dimension such as Customer Name rather than a cryptic field such as CUST.L33\_NAME (the physical table and field name in the Relational Database Management System [RDBMS]).
- Automatically connecting related tables via joins.
- Providing metrics that may calculate and aggregate facts such as revenue, number of customers, number of orders, number of incidents, and average selling price.

Figure 3-1 shows an example of building a query with the BusinessObjects universe, one of the first products to introduce the concept of a business view.

This business view is the most important piece of your BI front-end tools and one in which the business and IT must work together to model. For integrated BI platforms, the business view is common to all the BI tool modules: business query, reporting, analysis, and dashboards. When the business view looks too much like the data warehouse

## 38 Chapter 3

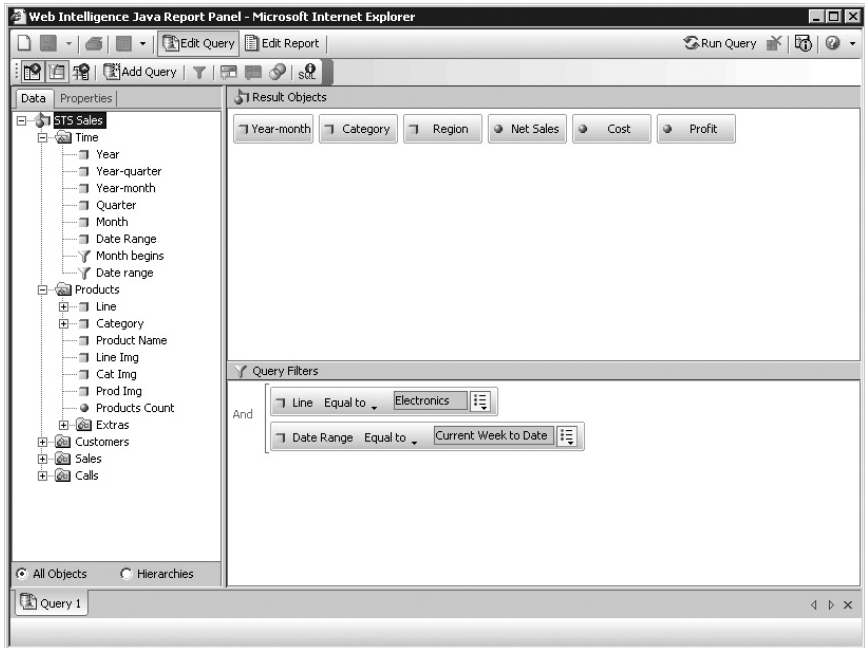


Figure 3-1 The BusinessObjects universe presents users with a business view of the data. (Reprinted with permission.)

### What Is Structured Query Language (SQL)

SQL, pronounced “sequel,” is a computer language used to communicate with a relational database.<sup>1</sup> SQL is a common language regardless if you use a database from Oracle, IBM, Microsoft, or Teradata. Querying a database with SQL can be fairly complicated. Business query tools will generate the SQL behind the scenes so business users don’t need to learn how to write SQL code. While there is a common set of SQL commands, such as SELECT and SUM, each database vendor may have its own SQL extensions or dialect. RANK, for example, is a popular SQL expression among business users but it is an expression that not all relational databases support. Sometimes when trying to develop a complex business query, you may run into limitations inherent in the SQL language. For example, a query about sales for this quarter would generate simple SQL. Asking a query about which products were cross-sold to the same customers this year versus last year would require very complex SQL and may be better answered in an OLAP database.

or source system with confusing table and field names, business users are overwhelmed and can too easily build incorrect queries. Poor business view design also forces users to put too much logic and too many calculations inside individual reports and dashboards. For these reasons, in some organizations, the power users within a business unit, function, or department, are responsible for building the business view or metadata layer; in others, it is the central BI group or data warehouse team that will build and maintain the business view.

## Production Reporting

Whereas business query and reporting tools allow for basic report formatting, production reporting tools have much more sophisticated formatting and design capabilities. Some people may refer to this category of tools as pixel perfect, operational, or enterprise reporting. Again, the terminology can be misleading as some business query and reporting tools can create pixel perfect reports, be embedded in operational systems, and are used across an enterprise. For lack of a better term, I will refer to this module as “production” reporting. Examples of production reporting tools include Actuate e.Report, BusinessObjects Crystal Reports, Microsoft Reporting Services, Oracle Publisher (which supersedes Oracle Reports), and Information Builders Web Focus.

A production reporting tool may access a transaction system directly to create a document such as an invoice, a bank statement, a check, or a list of open orders. When the reporting is not against the transaction system, it may be against an operational data store or detailed data within a data warehouse. IT usually develops these reports for the following reasons:

- The data source is an operational system in which you can't take the risk that “untrained” users may launch resource intensive and runaway queries with a business query tool.
- Reports are often accessed through and embedded within the transaction system.
- The information requirements are common to all users and departments and are static, such as for regulatory reports.

Because professional IT developers are often the users of production reporting tools, IT may also use these tools to develop management style reports, particularly when a company does not have a business query tool.

Characteristic	Production Reporting	Business Query and Reporting
Primary author	IT developer	Power user or business user
Primary purpose	Document preparation	Decision making, management
Report delivery	Paper or e-bill, embedded in application	Portal, spreadsheet, e-mail
Print quality	Pixel perfect	Presentation quality
User base	Tens of thousands	Hundreds or thousands
Data source	Operational transaction system	Data warehouse or mart, occasionally transaction system
Level of data detail	Granular	Aggregated
Scope	Operational	Tactical, strategic
Usage	Often embedded within an application	Most often BI as separate application

**Table 3-2** Differences Between Production Reporting Tools and Business Query Tools (Source: BIScorecard.com)

Table 3-2 highlights some key differences between business query tools and production reporting tools. None of these differences is an absolute, except that they serve the needs to distinct user groups and in many cases, distinct applications.

## Online Analytical Processing (OLAP)

Online Analytical Processing (OLAP) is a capability that focuses on *analyzing* and *exploring* data, whereas query and reporting tools put greater emphasis on accessing data for *monitoring* purposes. OLAP moves the focus from “what” is happening, to exploring “why” something is happening. To uncover the “why,” users may not know precisely what information they are looking for and instead will navigate and drill within a data set to uncover particular details and patterns.

OLAP provides interactive analysis by different dimensions (i.e., geography, product, time) and different levels of detail (year, quarter, month). For many users, OLAP has become synonymous with “drill-down” and “pivot” capabilities. Many BI products, though, will now provide drill-down and pivot capabilities without a full-blown OLAP engine or OLAP database on the back-end.

As the technology and users have evolved and matured, the distinctions between OLAP and reporting have increasingly blurred.

OLAP users want highly formatted reports that are based on multidimensional data; report users immediately want to drill when they see a problem with a particular metric in a report. They don't want to be forced to launch a separate tool as they move from reporting into analysis and exploration.

The following characteristics continue to distinguish OLAP tools from business query and reporting tools:

- **Multidimensional** Users analyze numerical values from different dimensions such as product, time, and geography. A report, on the other hand, may be one-dimensional, such as list of product prices at one point in time.
- **Consistently fast** As users navigate different dimensions and levels within a dimension, OLAP means fast—the speed of thought. If a user double-clicks to drill-down from Year to Quarter, waiting 24 hours, 24 minutes, or even 24 seconds for an answer is not acceptable. Report users, of course, do not want slow reports either, but some reports take this long to run and must be scheduled.
- **Highly interactive** Drilling is one way users interact with OLAP data. *Pivoting* gives users the ability to view information from different perspectives such as by geography or by product. *Slicing* allows users to filter the data within these dimensions such as sales for New York only and then for New Jersey only, or crime statistics for Leeds only and then Manchester only. This kind of interactivity within a non-OLAP report ranges from nonexistent to only recently possible.
- **Varying levels of aggregation** To ensure predictable query times, OLAP products pre-aggregate data in different ways. Reporting, to the contrary, can be at the lowest level of detail: rather than sales by product, you might have individual line items for a particular order number.
- **Cross-dimensional calculations** With multiple dimensions come more complex calculations. In OLAP, you might want to analyze percentage contribution or market share. These analyses require sub-totaling sales for a particular state and then calculating percentage contribution for the total region, country, or world. Users may analyze this percentage market share by a number of other dimensions, such as actual versus budget, this year versus last year, or for a particular group of products. These calculations often must be performed in a particular order and involve input numbers that users might never see. Detailed reports, however, often rely on simple subtotals or calculations of values that are displayed on the report itself.

In understanding OLAP requirements, it's important to distinguish between OLAP platform issues and OLAP user interface issues.

## OLAP Platforms

The OLAP platform is about how the data is stored to allow for multidimensional analysis. The cube shown in Figure 2-3, in Chapter 2, represents the OLAP database. On the one hand, business users should not have to care at all about how the data is stored, replicated, and cached, and yet the OLAP architecture greatly affects what you can analyze and how. The OLAP architecture also influences what OLAP front-end you can use.

There are four primary OLAP architectures as described in Table 3-3. Relational OLAP (ROLAP) platforms store data in a relational database so data is not necessarily replicated into a separate storage for analysis. Multidimensional OLAP (MOLAP) platforms replicate data into a purpose-built storage that ensures fast analysis. Hybrid OLAP (HOLAP) uses a combination of storage techniques. Dynamic OLAP (DOLAP) will automatically generate a small multidimensional cache when users run a query.

With each OLAP architecture, there are trade-offs in performance, types of multidimensional calculations, amount of data that can be analyzed, timeliness of data updates, and interfaces through which the data can be accessed.

Architecture	Primary Difference	Vendor
ROLAP	Calculations done in a relational database, large data volumes, less predictable drill times.	Oracle's BI EE, SAP Netweaver BI, MicroStrategy, Cognos 8, BusinessObjects Web Intelligence
MOLAP	Calculations performed in a server-based multidimensional database. Cubes provide write access for inputting budget data or performing what-if analysis.	Oracle's Hyperion Essbase, Microsoft Analysis Services, TM1, SAS OLAP, Cognos PowerCubes
HOLAP	Aggregations in a cache but with seamless drill-through to relational.	Microsoft Analysis Services, SAS OLAP, Oracle's Hyperion Essbase
DOLAP	Mini cache is built at query run time.	BusinessObjects Web Intelligence, Oracle's Hyperion Interactive Reporting (formerly Brio)

**Table 3-3** OLAP Architectures



Historically, many OLAP products used a MOLAP storage, which led to inflexible cube databases, management of more replicated data, and limitations on the data volumes and level of detail that can be analyzed. All of this has sometimes scared IT away from OLAP.

I would argue that every BI deployment needs an OLAP component; not only is it necessary to facilitate analysis, but also it can significantly reduce the number of reports either IT developers or business users have to create.

With OLAP, a report is just a starting view, say, Sales for 2007 by Country—a summarized starting point. As users click, drill, and pivot, the end result might be Sales, Unit Price, Volume for one quarter, for two products, in a particular city—a detailed, focused end-point. In a strictly relational reporting world, the starting view and end result would be two entirely separate reports, with dozens of iterations in between.

## OLAP Viewers

Microsoft Excel is one of the most popular interfaces to OLAP data. In fact, for three of the leading OLAP products (Oracle's Hyperion Essbase, Microsoft Analysis Services, SAP Business Explorer), the spreadsheet was initially the *only* interface. Users would open a spreadsheet and could immediately begin drilling within cells and Excel Pivot Tables to retrieve and explorer their data.

Today, Excel continues to be an important OLAP interface, but in addition, users can explore data via OLAP viewers. These OLAP viewers may be web-based (whereas Excel is desktop-based) and will have advanced charting and navigation capabilities. In addition, business query tools and production reporting tools may also be able to

### What Are Multidimensional Expressions (MDX)?

MDX is a query language similar to SQL but used to manipulate data within an OLAP database. Microsoft created MDX as a language to work with its original OLAP server, now referred to as SQL Server Analysis Services. As MDX gained industry acceptance, a number of other OLAP databases added support for MDX such that today OLAP viewers will generate MDX to access and analyze data in a number of different OLAP databases.



















