

note

What is Business Intelligence?

When IT professionals talk about *Business Intelligence* (BI) they usually mean the analysis and presentation of data stored in an OLAP cube.

In the recent past it was very expensive to implement a Business Intelligence solution.

BI Projects would often take many months to complete and involve large numbers of highly trained IT professionals to design (and extract data into) the OLAP cube.

The vision for the Excel data model is that ordinary Excel users can automatically create a ready-to-go OLAP cube almost instantly.

An OLAP pivot table (or OLAP pivot chart) can then be used to analyse and present the contents of the (automatically created) OLAP cube.

Microsoft sometimes call this concept: *Self Service BI*.

Lesson 6-3: Understand OLAP, MDX and Business Intelligence

What is an OLAP cube?

Dr E F Codd (an IBM scientist) invented the relational database in the early seventies. In 1979 Oracle released their first commercial relational database product.

By the mid to late 1990's relational databases had come of age. Almost all big corporations used a relational database to store their data.

With the adoption of relational databases came a problem. Companies found that relational databases were great at storing data, but made it difficult to generate management reports from transactional data.


If, for example, you needed to report sales for a specific month, the database server would potentially have to loop through millions of transactions in order to generate a total. On the hardware of the day this would often slow down the entire system to the point where it became unusable for critical business operations.

The solution to this problem was OLAP.

The idea behind OLAP was to pre-compute all of the totals and subtotals needed for reporting at night (or at the weekend) when the database server was normally idle.

The totals are stored in a special database called an OLAP cube containing *dimensions* and *measures* (see sidebar for an explanation of these terms).

An OLAP cube doesn't have to loop through any transactions because totals are all pre-calculated providing instant access.



An OLAP cube is a snapshot of data at a specific point in time (perhaps at the end of a specific day, week, month or year).

important

OLAP terminology

It is important that you understand OLAP terminology.

You'll find this terminology used in some of the OLAP functions that will be used later in this session.

You'll also find it used in Microsoft's help pages.

Dimensions (sometimes called Members)

As you know, OLAP cubes contain a large number of totals. Each of these totals correspond to two dimensions.

Dimensions are things such as: *country, branch, city, week, month, quarter or year.*

Measures

A measure is simply the thing that has been totalled.

Measures are values such as: *quantity in stock, quantity on order, invoice price or cost price.*

Hierarchies

A good example of a hierarchy is:

Months, quarters and years.

Items in a hierarchy have a parent-child relationship.

Hierarchies come into play when you "drill down" into data.

For example, you may look at sales for 2013 but then need to drill down into sales for May 2013. In this respect 2013 is a *parent* in the hierarchy and May is a *child*.


The standard version of Excel 2013 now provides a complete OLAP solution on your desktop!

In: *Lesson 6-2: Create a simple data model*, you actually created an OLAP cube without realizing it.

When you create an OLAP pivot table from a data model, an OLAP cube is automatically created in the computer's memory and is used to power the OLAP pivot table.

At any time you can refresh the cube using the current values in the source tables. With very large data sets it could take an appreciable amount of time for Excel to re-construct the cube.

How an OLAP cube works



To understand the diagram above you'll need to understand the key concepts of *Dimensions*, *Measures* and *Hierarchies* (see sidebar).

The diagram shows the constituent parts of an OLAP cube.

You can think of an OLAP cube as consisting of many boxes, each with a total inside. The values are said to be at cube *intersections* (of measures and dimensions).

For example, the cube in the diagram shows that in *January 2013* the *Sales* total for *Confections* was 2,233.

You can also see that *January 2013* has a *parent* called 2013 (where an annual total will also be stored).

MDX (multi-dimensional expressions)

MDX is a query language used to return the values stored in OLAP cubes. You'll see some MDX expressions later in: *Lesson 6-6: Convert CUBEVALUE functions to include MDX expressions.*

MDX is an open standard that is also used by many non-Microsoft products such as *Cognos Powerplay*.

An OLAP pivot table can (in theory) be used to report from any MDX compliant OLAP source.