

Don't Discard Those Spreadsheets: The Power of Excel-Friendly OLAP

The focus within the business performance management market on getting rid of spreadsheets is misguided. With the right database connectivity, spreadsheets really can be an enterprise technology.

SHOULD EXCEL BE A KEY COMPONENT OF YOUR company's business performance management (BPM) system? There's no doubt how most IT managers would answer this question. Name the top 10 requirements for a successful BPM system, and they'll quickly explain how Excel violates dozens of them.

The user community also has concerns about using spreadsheets for business performance management. As companies become larger and more complex, some suggest, they outgrow Excel. Managers need information right away, from diverse sources, and they need the information to be accurate. But spreadsheets don't scale well. They can't be used by many different people at once. Excel reports are notorious for having errors and for being ugly. Security is a joke. Consolidation efforts occupy a large corner of "spreadsheet hell." And Sarbanes-Oxley has changed everything. Or so we're told.

For these reasons and many more, a lot of companies have concluded that it's time to replace Excel for BPM. But before your organization takes that leap, perhaps you should take another look at Excel — or, more precisely, at the capabilities of spreadsheets that are enhanced by an Excel-friendly OLAP database. Excel-friendly OLAP helps eliminate many of the classic objections to using spreadsheets for business performance management.

What OLAP Can Do for Excel

OLAP is the name for a type of database technology that stores information in cubes, rather than in lists. A company might keep its general ledger accounting data in a simple OLAP cube that includes three dimensions: account, division, and month. At the intersection of any particular account, division, and month you would

by Charley Kyd



Charley Kyd is president of ExcelUser.com, which focuses on the business use of Microsoft Excel. He is a Microsoft Excel MVP and has used spreadsheets for management reporting and analysis since beta-testing VisiCalc in 1979.



find one number. Most cubes have more than three dimensions, and they typically contain a wide variety of business data, not merely G/L data. Users with access to an OLAP cube of corporate

information can define any consolidation hierarchy for the cube's dimensions. For example, in the "month" dimension, every month could roll up into quarters, which could roll up into years. Or months could roll up into year-to-date numbers. The individual data points (the cube's "leaf members") and the consolidated numbers would be equivalent sources of data. So users generating a report could choose data from a leaf member like Mar-2007 just as easily as they could choose from a consolidated member like Mar-2007-YTD.

There are probably 50 OLAP products on the market. While many can exchange data with spreadsheets, only a few offer spreadsheet functions that read data from OLAP cubes into Excel and write information from Excel to the OLAP database (see *Excel-Friendly OLAP Products*, at left). The difference between this type of "pull" technology and other OLAP products' "push" functionality is significant to users.

Database products with push functionality typically require users to choose data through their OLAP product's user interface, then write that information as a block of numbers to Excel. If a report relies on five different views of data, users must take this action five times. Worse, when the data is imported into Excel, it typically isn't placed where it's needed within the body of the report. It's parked in one area of the spreadsheet for use somewhere else. It really doesn't matter whether the data is imported as a text file, copied and pasted, generated by a PivotTable, or pushed to a spreadsheet by some other means; the data lands in Excel in the same inconvenient format.

In contrast, OLAP products that I call "Excel-friendly" pull data into Excel. Rather than writing numbers to a spreadsheet when a user issues a command, Excel-friendly OLAP provides current data to the spreadsheet whenever Excel users recalculate the workbook. Spreadsheets linked to Excel-friendly OLAP databases don't contain data; they contain only formulas linked to data on the server or in a local database. Suppose that an Excel dashboard presents performance information for a specific corporate division for a specific month. And suppose that users want to show the performance of each division in the company in this same format, and they want to be able to easily update the spreadsheet with new numbers at the end of every month. To do so, they might designate a "month" cell and a "division" cell, then reference these cells in every formula throughout the spreadsheet that links to the OLAP database. With this design, they could change the month cell from "Feb-2007" to "Mar-2007" and the division cell from "Northeast" to "Southwest." Then, by simply recalculating the workbook, users could update the report to reflect the new settings. They even could automate the process so that Excel would print a report for every division for a given month.

Using this type of pull technology, spreadsheet users can write formulas that reference any number of cells in any number of cubes in the database. A single cell in the spreadsheet can contain a formula that pulls data from several cubes. To illustrate, one formula could show labor costs (from the G/L cube) per full-time-equivalent employee (from the head-count cube). Another formula could show the ratio of total company sales (from the G/L cube) to the sales of its publicly traded competitor (from a competitor cube). Excel-friendly OLAP also typically allows users to write from their spreadsheets to the database.

When spreadsheets contain sensitive information or are able to write data back to corporate databases, security obviously becomes a concern. Within Excel-friendly OLAP systems, read-write security

Excel-Friendly OLAP Products

The best-known Excel-friendly OLAP product is Analysis Services, which is included with Microsoft SQL Server. Excel 2007 includes a variety of functions that read data from Analysis Services, but it won't write to Analysis Services. Users of Excel XP and earlier versions need one of several third-party Excel add-ins to use spreadsheet functions that read from Analysis Services. These products — which include XLCubed, IntelligentApps, and BIXL — also enable spreadsheets to write to Analysis Services data cubes.

Two other Excel-friendly OLAP products offer a sharp contrast to Analysis Services. My tests show that both products return data to Excel about 100 times faster than Analysis Services does. In addition, unlike Analysis Services, both products can be administered by knowledgeable users rather than the IT department. These two products are TM1, from Applix, and PowerOLAP, from PARIS Technologies. TM1, probably the first OLAP product, offers approximately 30 read-write spreadsheet functions. PowerOLAP offers more than 60 read-write spreadsheet functions. PARIS Technologies has also partnered with a subsidiary of Hitachi Ltd. to produce a Japanese version of the product.

These vendors' spreadsheet functions work slightly differently, but the most-used spreadsheet function for each product looks something like this:

```
=GETDATA(database, cube, member1, member2, ...)
```

This function could return a number from a cube named "GL" in the "finance" database, for account 1234, from the Southwest division, for July 2006 using this formula:

```
=GETDATA("Finance", "GL", "1234", "Southwest", "Jul-2006")
```

Like all Excel functions, this one typically would contain cell addresses or range names, not the literal values for each argument. This would enable the database, cube, account number, division, and date to be referenced in one location — and so an update to that one location would result in the updating of all calculations throughout the spreadsheet.

The functions that tie an Excel spreadsheet to a TM1 or PowerOLAP data cube won't be familiar to businesspeople who've never used Excel-friendly OLAP before. But the concepts behind them will be. These technologies present a manageable way to give business users throughout the company access to corporate data in real time.

typically can be defined down to the cell level in the OLAP cube. The OLAP administrator can ensure that only certain analysts can write to a forecast cube, or that a department manager can read only the salaries of people who report to him.

How Much Truth?

It's common these days for database vendors to talk about having "one version of the truth." What's less common is for anyone to ask these vendors exactly how relevant the truth in their systems is. This is a critical question for managers looking for BPM information, and for their staff — usually Excel users — who must provide performance information. It's true that corporate data warehouses typically contain massive numbers of transactions. But this exhaustive detail is largely irrelevant to BPM, which typically relies on detailed summaries of data. A well-designed OLAP database incorporates data that is not in most corporate data warehouses for a variety of reasons:

Data silos. Many information systems — both old and new — rely on databases that never will be added to the data warehouse. But often HR, CRM, supply chain, and many other siloed applications contain data that can help managers properly oversee business performance. Most provide some way to export data into the OLAP database that the organization's performance analysts use.

Mergers and acquisitions. When two companies merge, the new organization suddenly has two data warehouses, not one. Even if each organization has one version of its own truth, neither has one version of the whole truth. This is not an easy problem for IT to solve. I know of one company, for example, that has five ERP systems on four continents. For most of the past decade, one goal of this organization's IT department has been to create a single data warehouse within the following two years. Unfortunately, users and their managers need summary data to be fully available immediately, certainly by the end of the month in which a merger or acquisition closes. Decisions can't wait until the IT systems are fully merged years later.

OLAP databases speed up the data integration process considerably. One company that uses Excel-friendly OLAP purchased a billion-dollar subsidiary several years ago. By the board meeting two weeks after the transaction closed, the finance staff had printed more than 200 spreadsheets that reported both consolidated and consolidating reports for the new company, down to low-level summaries. All financial data was expressed in terms of the parent company's chart of accounts. The staff was able to integrate the disparate systems so quickly for at least two reasons.

First, the parent already was using an Excel-friendly OLAP system. Users mapped the subsidiary's metadata (e.g., G/L codes, department codes) to the parent's metadata. They imported the subsidiary's financials to a new "slice" in the parent's general ledger cube, translating the metadata on the fly. Then they printed their standard spreadsheet analyses, all 200 pages of them, while adding a few new Excel analyses specific to the new subsidiary. If the staff had relied on spreadsheets alone, the reporting process would have been a nightmare.

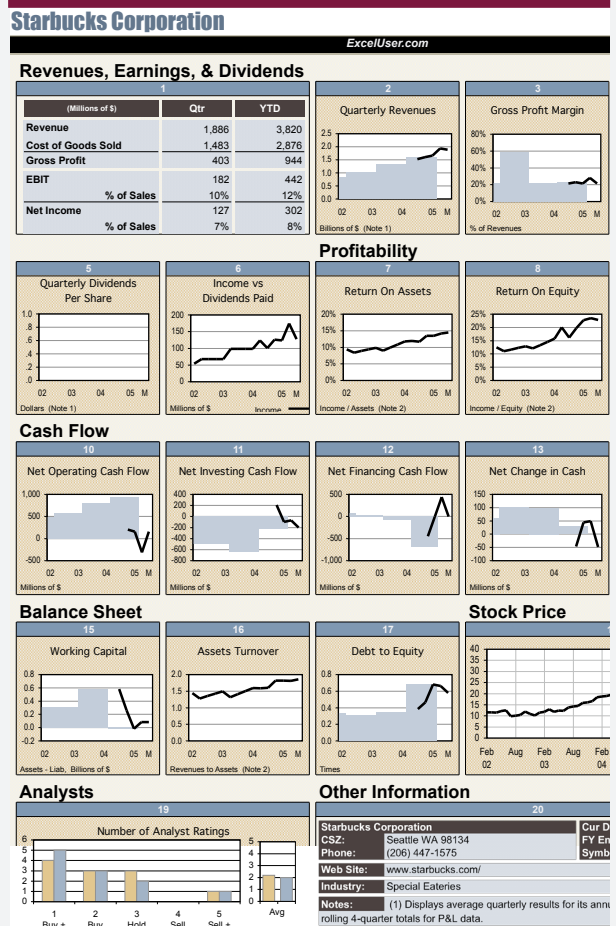
The second reason the new company was able to consolidate its reports so quickly is that its subject matter experts were doing the work. As the finance team experienced the many problems inherent in a project like this, they could quickly and effectively solve those problems. Where compromises were necessary, the OLAP users themselves could make the best choices, and when they made mistakes they could recognize them quickly. In contrast, if IT technicians had been doing the same work creating consolidated reports with the same Excel-friendly OLAP technology, they would not have recognized many key problems in the data and metadata — and problems they did notice would have required a series of meetings to discuss, followed by redesigns and more meetings.

System conversions. When a company purchases a new ERP system, it creates at least two problems for BPM reporting. First, the company

It's common for database vendors to talk about "one version of the truth." But how relevant is the truth in their systems?

Exhibit 1

Dashboard Formatting in Excel



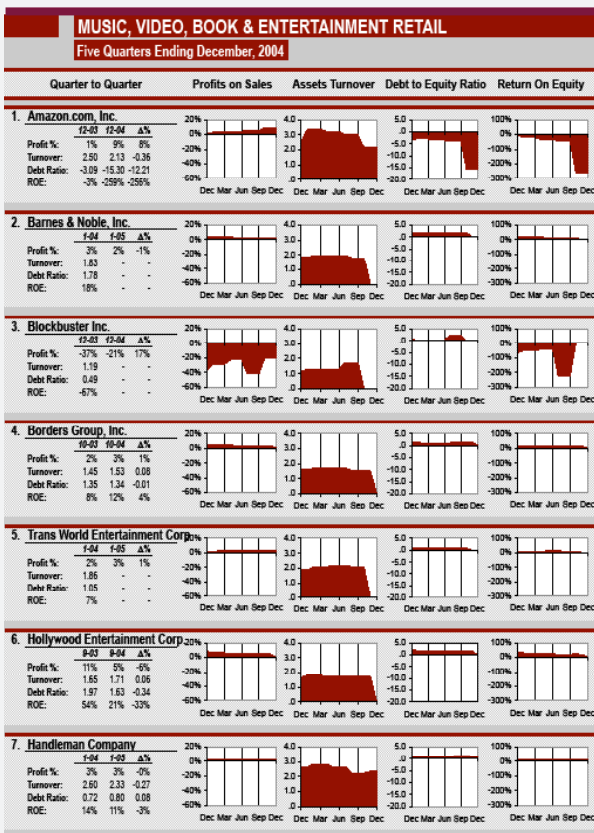
typically converts the fewest months of historical data it can get away with. For financial systems, companies often convert only one year of history prior to the current fiscal year. But for performance management purposes, data about past performance is very useful, frequently even crucial. For example, monthly time-series forecasting requires at least 30 months of historical data, preferably more. New products and new sales offices within an organization often follow a consistent pattern for both revenue growth and startup expenses, but those patterns can be discovered only by analyzing data for startups over the course of several years. And the analysis of trends in cost-volume-profit relationships during past downturns can serve as a guide to cost-reduction efforts during a current downturn. For companies that have an Excel-friendly OLAP database in place, conversion to a new ERP system has less impact because all of the historical performance data within the OLAP cube remains available. Better yet, managers continue to receive their standard Excel reports, which can contain data from both ERP systems.

The second performance management problem caused by a new ERP system is that transactions may be classified differently than they were in the old system. Excel-friendly OLAP systems can help users find the inevitable classification errors. To illustrate, I know of two large companies whose system conversions went significantly over budget. The accountants for both companies had specified that all account-department combinations that had not explicitly been allowed were to be rejected by the new systems. But to reduce costs, both systems were set up to allow all such combinations that weren't specifically prohibited. As a consequence, many transactions each month were automatically booked to incorrect account-department combinations. If they had not been using an Excel-friendly OLAP system, the accountants in each company would have had to manually inspect more than 100 million account balances to find G/L accounts whose transaction patterns had changed when the accounting systems changed. This would have been an impossible task, of course.

Fortunately, both companies were using Excel-friendly OLAP systems before their conversions began. Each created a simple spreadsheet that returned the monthly transactions for any specific account, department, and division for the 12 months prior to the conversion and for all months after. Then using standard Excel statistics functions and simple spreadsheet automation, the companies' analysts were able to loop through every combination of account, department, and division and to list all questionable combinations. The staff quickly corrected the obvious mistakes and researched the others.

Exhibit 2

Excel Report on Competitive Financial Data



External data. Managers often need to see their performance reported within the context of their business environment. That environment can be described by the financial data of the company's publicly held customers and competitors, by local and regional economic data, by population trends, and by other measures. IT doesn't control such data, and few IT managers understand it. That's not their job; it's the business user's job. In most companies, if business users don't create and maintain cubes of external data, no one ever will. It's not unusual for a knowledgeable user to create an OLAP cube on her local computer, populate it with public data, and then test its use with various spreadsheet reports. Once the cube is tested, she can work with the database administrator to move the cube to the OLAP server.

Forecasts. Most data warehouses provide empty buckets for budget data, but they typically don't capture the wide variety of forecasts that companies generate. Nor do they help to generate those forecasts. Excel-friendly OLAP can do both. For instance, Excel-friendly OLAP users easily can generate forecasts, compare the forecasts with one another to find conflicts, and then revise the forecasts after researching the differences.

To prepare manual forecasts, they would send a forecasting spreadsheet to the company's salespeople — but unlike most forecasting spreadsheets, this one would include formulas that write the salespeople's data to the appropriate area of an OLAP cube on the server. The finance department could simultaneously create a spreadsheet that used statistical methods to project forecasts based on past sales performance. This spreadsheet also could write the forecast to an area of the OLAP cube.

Excel is becoming less popular for reporting performance information, in part because spreadsheet reports are ugly. But they don't need to be.

Then, using automation, users could apply this spreadsheet forecast to each product and region. Finally, users could set up a spreadsheet to compare the manual and statistical forecasts. Again using automation, they could calculate the workbook for every combination of product and region and automatically note where the two versions varied by an unusual degree.

Data corrections. Forecasts, analyses, and management reporting all can be seriously flawed if analysts rely on historical data that includes errors and oversights. But correcting financial data can be tricky because, for a variety of reasons, managers, investors, and auditors all take a dim view of prior-period adjustments to the general ledger. One way to handle this problem, particularly for data that will be used in forecasting and analysis, is for users to maintain an error-correction entity that can be consolidated or ignored, depending on the circumstance. Of course, the corrections must be managed carefully. It would be very easy, after all, for indications of real problems to be “corrected” out of existence. But when statistical corrections are tightly controlled, they provide the only practical way to analyze past performance as it actually happened, not as it was mistakenly booked at the time.

Excel Dashboard Reporting

Excel is becoming less popular for reporting performance information, in part because spreadsheet reports are usually ugly and difficult to read. But they don't need to be. Exhibit 1, on page 17, is an Excel dashboard report. I created this report from public data on Starbucks Corp. completely in Excel, with no assistance from third-party tools. My spreadsheet could generate an equivalent report for any public company whose financial information is on the two Web sites from which I imported Starbucks' data. In a business environment, a report like this could show performance for a department, division, product line, or entire company, drawing data from an Excel-friendly OLAP database rather than the Internet.

One significant advantage of using Excel for this type of reporting is that Excel users can change the report quickly and easily without involving the IT department. In fact, assuming that the necessary data already resides in the OLAP database, an Excel user typically could replace one measure with another in less than 10 minutes. Another significant advantage is that the report — even a single graph in the report — can display data from many original sources. To illustrate, a graph could show the trend in operating expenses (from the G/L cube) per current customer (from the sales cube). Another figure could show the ratio of unit sales for a relevant product (from the sales cube) to the number of potential customers (from the market cube).

Excel dashboards can compare the same measures for many different products, divisions, departments, and other entities. To illustrate, exhibit 2 on page 18 compares the same financial ratios for each of seven public companies. The same workbook could generate dashboards for any number of companies, seven at a time. There is virtually no limit to the variety of appearances that an Excel dashboard can take. I often get ideas for dashboard designs when reading business magazines.

As a general rule, Excel output tends to be on paper rather than on screen. Paper reporting may seem primitive, but many managers prefer to view hard copies of reports. Excel-friendly OLAP vendors are making progress toward Web-based reporting. PARIS Technologies allows Excel reports to connect to an OLAP database over the Internet. This gives Excel users read-write access to their cubes from anywhere they have Internet access. Several vendors offer interactive Web implementations of Excel reports linked to Applix TM1 cubes. And Office 2007 offers ways for Excel users to interact with Microsoft Analysis Services over the Web.

On the other hand, the high-tech solution might not always be the best solution. One large company decided to take a low-tech approach to online management reporting. Each month the company automatically captures nearly 10,000 bitmaps of Excel reports of its OLAP data, then displays those images on the corporate intranet.

Whatever the format of the reports they generate, Excel-friendly OLAP databases have many benefits. They provide a familiar interface for business users across the company to interact with corporate data while requiring little training in most cases. They provide quick access to company data for both ad hoc analyses and standardized reporting, without requiring IT's help for either. They provide the security that stand-alone spreadsheets sorely lack. And they reduce the likelihood that companies will suffer the “multiple versions of the truth” problems for which Excel BPM efforts are notorious. At the very least, Excel-friendly OLAP should convince you to take another look at Excel for business performance management. **bm**