strategy+business

A CIO's View of the Balanced Scorecard

For further information:
Ilona Steffen, Zurich: ilona.steffen@booz.com
Booz & Company

booz&co.

Reprint

A CIO's View of the Balanced Scorecard

Defining the data and the timing of reporting is the hardest task, says a scorecard veteran.

by George Tillmann

When Robert Kaplan and David Norton published "The Balanced Scorecard: Measures that Drive Performance" in the *Harvard Business Review* in 1992, the idea of measuring business performance from financial and nonfinancial perspectives was novel. Their original balanced scorecard significantly advanced the notion that effective performance measurement must provide a view of both.

Today, the balanced scorecard is one of the most widely used and hotly debated management tools in the executive arsenal. Countless permutations of the original Kaplan and Norton framework, in an equally varied multitude of applications, have been tried by all kinds of organizations. Although the original balanced scorecard concept was meant to assess the health of an entire business, the basic concept has been adapted to fit business units and support organizations. When it works, the scorecard is a powerful resource to help executives understand past and current performance, and plan for the future.

Scorecards can be a great resource for managing the IT function. First, considerable numeric data is available to measure systems performance. Second, IT scorecards can be designed to measure end-user benefits and satisfaction. Third, scorecards can be a powerful vehicle to bridge the communication gap between IT professionals and the business customers they serve. So, when I took over as CIO of Booz Allen Hamilton in 2000, I wanted to set up a scorecard for my own management team.

Whether the cause is information systems' mystique, executives' technophobia, or poor marketing by information technology managers, for most senior business executives, delving into IT reports has limited appeal. IT performance reports that appear in a form comparable to reports of other business functions can offer a clearer window for businesspeople into the IT domain, particularly when it reveals corporate senior management and business unit leaders the value they receive from IT services.

Why Scorecards Fail

As is the case with any business tool,

George Tillmann
(tillmann_george@bah.com) is a vice
president with Booz Allen Hamilton in
McLean, Va. He spent his first 17 years at
the firm as a management consultant
specializing in information technology,
and the last four years as its chief infor-
mation officer.

however, the scorecard is not a magic wand; its value depends heavily on how it is implemented. I had the advantage of having worked for Booz Allen as a management consultant advising CEOs, CIOs, and other senior executives, so I knew what we were taking on.

Over the years, I had witnessed many attempts by large corporate IT departments — some successful, some not — to help implement scorecards or scorecardlike systems. Although I am a strong advocate of scorecards, I've seen enough unsuccessful ones in my time to have learned from those that didn't work.

Take the difficult and error-prone processes of defining the data (perhaps the hardest task of all) and the timing of reporting; many a scorecard has failed because the designers didn't get these two elements right.

The CEO of a U.S.-based health insurer I worked with wanted to create a business scorecard that provided more timely information than he was receiving from the traditional reporting system. The existing systems produced business unit performance data for review once a week. And although he received

would have a better chance of turning problems into opportunities. He concluded he wanted information that was no more than 24 hours old for all the business units, listed side by side on one report so it would be easy for him to compare data.

Even with the CEO requesting the scorecard, the project to create it failed. Initially, managers blamed the IT systems for technical problems with integrating data from the different business units to create the consolidated report the CEO had requested. I was on the team that was brought in to assess the situation and recommend a solution. As it turned out, the problem in producing the new scorecard had nothing to do with technology. Rather, it was related to defining the data. In this case, we found that no two business units defined data in the same way. In fact, none of the business units even shared the same definition of revenue, so comparisons were nearly meaningless.

At another company we worked with (a book and magazine publisher), the problem had to do with defining the time period for reporting.

When the company decided to

At one U.S. health insurer, business unit heads couldn't agree on the definition of revenue.

detailed reports from his four different business units, he didn't have consolidated information showing, at a high level, how the entire business was doing.

The CEO thought if he could follow performance trends daily, he

implement a scorecard, corporate leaders required business unit heads to submit their data to the scorecard team by the fifth day of each calendar month. To meet this deadline, the unit leaders, who wanted to review the data before forwarding it

to the scorecard team, asked their division heads to send the data to them for the previous 30 days by the 25th of each month. In turn, the division heads had their department leads supply them with data by the 20th of the month, and so on. This approach actually made some sense since many of the dates corresponded to the periods when data was sent to or received from external printers and distributors.

Although the scorecard accurately recorded events, it was still confusing for users. Since the end of the reporting time period varied by reporting level, there was no single frame of reference. Junior staff would see an event in one report, while the same event might appear a month later in the report for senior managers. This "when exactly did this occur?" problem strained communications and made comparisons difficult. The scorecard continued to be produced for a few years, but it was quickly marginalized by newer systems with more common reporting cycles.

Then there's the ever-present problem of "dueling data." Database administrators (DBAs) know that nothing can damage confidence in a database more than inconsistency. Having two pieces of the same data in a database is worse than having only one incorrect piece; any DBA worth his or her salt knows that if you have multiple pieces of the same data in a database, you should delete all but one. (One wrong employee number can be blamed on HR. Two employee numbers for one individual spells trouble for IT.) The same is true with scorecards.

And what if your scorecard data isn't consistent with the data in other, more established, reporting

systems? You're in trouble, as the backers of the more established systems (regardless of whose data is correct) will attack the newer scorecard application. Although both sets of data may be correct, if the two data sources appear to contradict each other, the veracity of at least one of the systems will be brought into question. Ultimately, the credibility of all the reporting systems, not just the scorecard, could be damaged.

Another recurring reason scorecards fail should be the most obvious, but it often isn't: Some people just don't want a lot of attention placed on their performance.

Consider the manufacturer of industrial products that was having a difficult time implementing an Executive Information System (EIS), an older, broader concept that is similar to a scorecard. The software, as well as the analysis performed by the company's IT department, was blamed. But a basic investigation turned up a different culprit: The head of one business unit never allowed data about his department to be sent to the chief executive officer. Instead, he took the reports personally to the CEO, reserving 90 minutes to walk the chief through the numbers. This way, the business unit leader could emphasize and spin the data in all the right places. He did not need, or want, a simple EIS, which would have allowed the CEO an unescort-

ed promenade through the data. He even told me that the EIS would be used "over my dead body."

When the marketing director for a building products manufacturer wanted to quash his company's scorecard project, he threw numerous roadblocks in its way and waited for the enthusiasm for it to die. I watched the marketing director go in for the kill, and then saw him formally dissolve the abandoned project.

Experienced project managers know that between the initial excitement and enthusiasm for a new project and the first deliverables, there are two critical periods. The

Inconsistency can damage confidence in a database. Having two pieces of the same data is worse than having one incorrect piece.

first is the honeymoon, usually at least 12 weeks, but rarely more than 26 weeks, when the project is running on the goodwill created during its kickoff. Project champions are still extolling the potential value, and critics, who lost the battle in opposing the project, find it safer to keep a low profile.

But when the sunny honeymoon ends, the dark clouds of the second period roll in fast and thick; enthusiasts become nervous and naysayers feel safe to come out and criticize. If the dark days last too long, say more than two or three months, a project that seemed to have much support may be marginalized or even canceled. Scorecard projects are especially vulnerable when there are major corporate distractions, such as an economic

downturn, a profit shortfall, or unexpected management changes.

Scorecard Wisdom

When I became Booz Allen's CIO, I was determined to introduce a scorecard for the IT department without making the mistakes of many of my former clients. First, we tackled data integrity issues, including dueling data and definitional problems. As luck would have it, considerable IT reporting was being done at the junior and middle manager level, but little was being done at the senior management level. So we had a relatively clean slate.

Where no data existed, we introduced it after studying and identifying appropriate key performance measures. There are two ways to look at the selection of IT reporting data: Key performance indicators (KPIs) give us the information we *want* to know; performance measures (PMs) are the actual information we *can,* or *are likely to,* get. KPIs are the performance knowledge we want to have about an asset, process, or group, such as sales growth, customer satisfaction, and e-mail availability. PMs are the reporting measures we can get our hands on — such as sales order totals, the number of survey respondents who liked the service, or the percentage of the time e-mail is available when users access it.

Corporate contentment happens when the KPIs and PMs are nearly identical. Major difficulties occur when their definitions diverge. To avoid this, we spent more time on data definition than on any other single facet of the scorecard. It is a process that never ends. As new systems, processes, or groups are brought under the scorecard, all definitions have to be reviewed and recertified.

Another critical step is defining the target audiences, which, for us, include corporate and business unit management, senior and junior IT managers, and journeymen IT staff. We wanted one report for all of them. But to create one report, you need to address two very different perspectives: IT shops think of themselves as selling technology, but users buy service. This distinction is at the heart of much of what goes wrong in managing IT, especially with respect to its relationship with business clients.

More IT departments are starting to appreciate how important it is to address these two perspectives when reporting IT's value to the business. Value is not a question of one perspective being more important than the other. Reporting on both is necessary. What needs to be different in the process is *how* each is reported. At Booz Allen, our IT scorecard is multilayered; the top layers have relatively few items and are focused on *service offerings* — the services a user wants and is willing to pay for. Examples are collaboration and communication, telephones, order processing, and financial accounting. Business users find the highest levels of the scorecard most useful.

Technology staff use the scorecard's lower layers to track the status and performance of each service offering's technology components, such as e-mail applications, directory services, WANs, LANs, servers, storage systems, and so on. Technology component performance is what IT shops traditionally are concerned with and what their internal systems traditionally report. Mixing the two can lead to confusion and mistrust. Tell a user who could not get into e-mail Tuesday that e-mail was 100 percent available last week,

IT shops sell technology, but users buy service — a distinction that explains many ills of IT management.

and you will know what I mean. He sees e-mail as being down on Tuesday; the IT manager sees e-mail as up and running. Why? Because the network, which is supported by a different IT group, was down. The result: irreconcilable differences.

High-level summaries with data that "drills down" into certain details are common in IT scorecards. What is unusual about our scorecard is the nature of the layered approach. Items at each level have been carefully selected for a particular audience. Indeed, creating a scorecard for multiple audiences with such diverse information needs meant we had to do detailed customer segmentation analysis to understand what key performance indicators users of the scorecard wanted to see monitored. Once the key performance indicators were understood, it was relatively easy to identify the performance measures. With a clear understanding of our audience, we were well positioned to identify the data various constituencies needed and develop the

right detailed data definitions for our scorecard.

At our company, we knew that if a multiyear IT scorecard project was to succeed, we needed to:

• Keep the communication program about the scorecard project in high gear, and keep pressure on the participants to produce.

• Ensure that all senior IT managers were publicly enthusiastic about the project (even if they were skeptical in private).

• Introduce (almost monthly) new functionality/features in the scorecard. Each change can be small; the advantage of gradual improvements is that mistakes will be smaller and easier to handle, so there's less risk of a grandstand collapse. Steady improvements can also reduce the project's dark days when enthusiasm and support wane.

• Maintain users' support though formal communication and demonstrations, informal dog-and-pony shows, and hands-on interaction. We try to emphasize the value of the scorecard and manage expectations. As is always the case, it is better to underpromise and over-deliver.

To date, we are pleased with the results. Our IT scorecard provides on a single page top-level information about our basic service offerings plus financial, customer satisfaction, and human resources information using a stoplight metaphor. (Green is good, yellow means caution, and red indicates there's a problem.) The colors are all backed up with detailed data (numeric where possible).

Level Two reports on the same 14 areas, but in greater detail — approximately 140 data points — using the same stoplight format. Together the top two levels com-

prise the organization's summary scorecard, which is sent to company senior management each month.

Starting with Level Three and down through the lowest levels of the scorecard, the focus shifts to technology components. The data is numerical, and it is linked to numerical targets. At the lowest levels, the scorecard is organized for use by specific IT managers, and can involve thousands of data points. All information is available to all IT staff. Critical to all levels and all data points is a set of definitions that include the information about what is being measured and what the service level targets are.

Continuous Improvements

For us, the next steps are quite clear. We are tightly integrating our monthly scorecard with numerous daily and near-real-time scorecards, monitors, and event management systems to prevent dueling data. We plan to enhance the scorecard's ability to predict rather than just report on status. More sophisticated trend analysis means our databases will have to better capture changes over time and draw the proper predictive conclusions.

We will also change our internal structure (processes, skills, budgets, even our organization) to better reflect our service offerings, rather than our technology components. This means adopting a matrix organizational structure, with our service-offering management structure orthogonal to the traditional IT management silos. Staff will need to develop the ability to report to two managers: the traditional technology manager and the new service-offering manager. And most important, the culture will need to change so that everyone takes to heart that

it is the service-offering dimension, and not the more comfortable technology dimension, that will drive our decisions and reflect our success or failure. The adoption of the service-offering orientation of our scorecard is an early phase of this culture change.

We recognize that if the IT scorecard is to be more than a simple report — that is, if it is to represent the core of who we are and what we do — then its refinement will be an ongoing process, and not just another short-term project. It is our goal to use the scorecard as the rallying point for improvements in IT service and beneficial change for our company as a whole. ✛

booz&co.