CHAPTER 11

How to Create Effective Performance Metrics

One of the most common questions people ask about performance dashboards is "How do we create effective metrics?"

Change Agents. Metrics are important because they govern how employees do their jobs. The adage "What gets measured gets done" is true. Metrics focus employees' attention on the tasks and processes that executives deem most critical to the success of the business. Metrics are like levers that executives can pull to move the organization in new and different directions. In fact, among all the tools available to executives to change the organization and move it in a new direction, performance measures are perhaps the most powerful.

For that reason, executives need to treat metrics with respect. As powerful agents of change, metrics can drive unparalleled improvements or plunge the organization into chaos and confusion. If the metrics do not accurately translate the company's strategy and goals into concrete actions on a daily basis, the organization will flounder. Employees will work at cross-purposes, impeding each other's progress and leaving everyone tired and frustrated with little to show for their efforts. In short, the company will be efficient but ineffective.

A trucking company, for example, that measures performance by the percentage of on-time shipments may drive hauling costs skyward because the metric does nothing to discourage dispatchers from sending out half-empty trucks to meet their schedules. To keep costs in line, the company needs to add a second metric that measures the percentage of unused cargo capacity in outgoing trucks, and it needs to revise the first metric so it emphasizes meeting customer expectations for fast, reliable shipments rather than just on-time deliveries. This combination of metrics gives dispatchers leeway to contact customers and renegotiate shipping schedules if they know the customer may be flexible.

More Art than Science. Crafting sound metrics is more an art than a science. Although a metrics team may spend months collecting requirements, standardizing definitions and rules, prioritizing metrics, and soliciting feedback—in short, following all the rules for solid metric development—it still may not succeed. In fact, there is a danger that metrics teams will shoot for perfection and fall prey to analysis paralysis. In reality, KPI teams can get only 80 percent of the way to an effective set of metrics; the last 20 percent comes from deploying the metrics, seeing how they impact behavior and performance, and then adjusting them accordingly.

"Only when you put the metrics out there, do you really understand what behaviors you are driving," says John Lochrie, former senior vice president of Direct Energy Essential Home Services.

Understanding Metrics

Types of Metrics

Metrics versus Indicators. A metric is a measurement of business activity, such as a "number of new customers," "average mean time between repair," or "total sales." But in a performance management system, we want to do more than just measure business activity; we want to measure how well we are executing business strategy.

To measure business strategy, we compare business activity to a goal defined in a plan designed to achieve the strategy. In the examples just listed, goals might be "25 new customers per month," "10,000 hours between repairs," or "$10 million in sales this quarter," respectively.

A metric that measures business activity against a goal is called a performance indicator. If designed properly, a performance indicator embodies the organization's or group's strategy. It measures what is important and compares performance to time-based targets. Performance indicators show how closely the output of our business activity is tracking to plan and "indicates" whether we are on course to achieve our strategic objectives.

There are two major types of performance indicators: outcome metrics and driver metrics. There are also diagnostic or activity metrics that don't have goals attached but are nevertheless important to measure. In addition, there are risk indicators that measure risk instead of performance. Finally, there are key performance indicators (KPIs), which are few in number (i.e., "key") but have widespread impact on the business.

Outcome Metrics. Outcome metrics—sometimes known as lagging indicators—measure the output of business activity that a strategy is designed to achieve. So, if the organization's strategy calls for a 12 percent return rate, the outcome metric might be "return rate per month." Outcome

metrics are generally backward looking. They measure past activity that has already happened and cannot be changed.

A majority of outcome metrics are financial in nature—revenues, profits, return on equity—because most commercial organizations define their strategy and success in financial terms. But this is not always the case. Cisco, for example, displays its overall customer satisfaction metric alongside its stock price on the company's intranet home page, reflecting how deeply the company believes customer satisfaction and loyalty are the key to the company's success. Nonprofit and government agencies have many nonfinancial outcome metrics because their mission is to serve the public good in some way. For instance, a hospital might have "patient safety" as an outcome metric while a public works department might have "quality of roads."

Driver Metrics. Driver metrics—sometimes known as leading indicators—measure business activity that influences the results of the outcome KPIs. Driver metrics are tactical in nature: They measure activity that happens between the periods in which outcomes are measured. So if an outcome metric is monthly sales, a typical driver metric might be weekly or daily sales. The purpose of driver metrics is to track current activity so people can make necessary adjustments to meet or exceed outcome goals for the period. Driver metrics foster proactivity.

Actionable. The best types of driver metrics are actionable. They give workers enough time to make necessary adjustments to affect the desired outcomes. For example, if monthly sales drops, executives might realign vacation schedules to ensure there are enough salespeople working at all times. Or they can ask salespeople to schedule more meetings with clients and prospects during the next two weeks.

Predictive. Often, organizations use regression algorithms to predict future outcomes based on current levels of activity so workers can see whether they are going to meet targets by the end of the period. For example, an online mortgage company identified two driver metrics that correlate with the ability of its salespeople to meet daily quotas: (1) the amount of time they spend on the phone with customers and (2) the number of clients they speak with each day. The company now displays these two "current-state" metrics prominently on its dashboards. It also created a third driver metric based on the other two that projects every 15 minutes whether salespeople are on track to meet their daily quotas.

Brainstorming Drivers. It behooves organizations to spend time brainstorming activities and behaviors that correlate with outcome drivers. The more an organization aligns drivers and outcomes, the better it can adapt to changes in the marketplace and achieve its goals. So, in our last example, an organization might determine that the best driver metric for monthly sales is the number of salespeople on the street. Or they might get more granular

and determine that the best driver is the number of face-to-face meetings that salespeople have scheduled with customers in the next two weeks.

Five Whys. There are two techniques for brainstorming driver metrics. One is the "Five Whys." Here a metrics team starts by identifying key outcomes aligned with strategy that it wants to achieve. Then a facilitator selects one outcome and asks the team, "What drives, causes, or contributes to that outcome?" The team brainstorms potential drivers and selects one. The facilitator then asks, "What drives that metric?" The process repeats itself three more times until the team comes up with a few good candidate metrics.

Sensory Perceptions. Stacey Barr, a performance management specialist in Australia, suggests that a group take a more "sensory" approach. She says a team should identify an outcome and then describe how it would "look, feel, and sound" if it were happening. Then the team could brainstorm ways to count whether the outcome was happening. These counts become the driver metrics. For example, a manufacturing organization that wants to reduce injuries might envision a limber, flexible employee who participates in daily stretching exercises. To count such behavior, the team decides the only reasonable approach is to have the floor supervisor check a daily form to indicate that workers have stretched. Spotlight 11.1 provides examples of how to correlate outcome and driver metrics.

Other Metrics

Activity Metrics. Activity metrics measure business activity related to performance indicators but don't have goals associated with them. They provide additional context about performance that helps businesspeople make informed decisions. For instance, activity metrics include top 10 lists, such as top 10 customers by revenue or bottom 10 suppliers by on-time delivery performance. Activity metrics might also include 360-degree views of customers or suppliers as well as counts that indicate demand for some resource, such as number of cell phones issued, number of shuttle bus runs, or number of procurements processed.

Risk Indicators. Unlike performance indicators, which measure how well something is being done, risk indicators measure the riskiness of a business activity or how it might adversely affect operations. A risk indicator measure provide an early warning sign to identify events that may harm the continuity of existing processes. Like performance indicators, risk indicators have goals associated with them. For instance, the Ministry of Works in Bahrain profiled in Chapter 9 has numerous risk objectives and associated metrics.

Key Performance Indicator. A KPI is a performance indicator that has a profound impact on the business. David Parmenter, author of the book *Key Performance Indicators: Developing, Implementing, and Using*

Spotlight 11.1 Cause-Effect Relationships

It is easy to define outcome metrics, but it takes imagination to identify driver metrics. One must follow the trail backward from results measured by an outcome metric to a first-mover driver. Because each outcome metric has numerous drivers, the key to defining effective drivers is to find the one or two that have the greatest impact on results desired by executives. Here are a few examples of driver metrics and the outcome metrics indicators they influence.

Driver Metric	Outcome Metric
Complex repairs completed successfully during the first call or visit	Customer satisfaction
Number of signed, positive employee suggestions each week or ratio of positive to negative comments	Employee satisfaction
Number of part for which orders exceed forecasts within 30 days of scheduled delivery	Per unit manufacturing costs
Number of days with lowest prices for comparable products	Market share
Number of customers who are delinquent paying their first bill	Customer churn
Number of loyalty rewards cashed in each month	Customer loyalty

Eyes of the Beholder. Truth be told, often there isn't a lot of difference between an outcome metric and a driver metric. "One man's outcome measure is another man's value driver," says Neal Williams, founder of Corda, a dashboard solutions provider. "An outcome KPI in one dashboard could be a driver KPI in another." For example, a business intelligence (BI) team's outcome metric is "number of active BI users per month." But this is simply a driver metric for the overall information technology (IT) organization, whose outcome metric might be "return on assets" or "customer satisfaction" or "cost efficiency."

Winning KPIs (John Wiley & Sons, 2007), says a KPI is a nonfinancial metric that is measured frequently and affects most of the company's critical success factors and performance metrics. Most important, he says, a KPI is monitored directly by the chief executive officer (CEO).

Parmenter recounts the story of Lord King, chairman of British Airways, who reportedly turned around the ailing airline in the 1980s using a single

KPI: the timely arrival and departure of airplanes. "[Lord King] was notified, wherever he was in the world, when a British Airways plane was delayed more than two hours. The British Airways airport manager at the relevant airport knew that if a plane was delayed beyond this threshold, he or she would receive a personal call from the chairman. It was not long before British Airways planes had a reputation for leaving on time," says Parmenter.

Of course, British Airways, or any airline for that matter, needs to make sure it doesn't value on-time arrivals and departures at the expense of safety and customer satisfaction. Employees ranging from flight attendants, pilots, mechanics, and bag handlers, who are under pressure to achieve time-based metrics, may be tempted to cut corners in ways that could have serious consequences. If bags are lost, customers harassed, pilots fatigued or ill, or safety compromised, the airline's brand could be seriously tarnished. Therefore, KPIs must be balanced by a host of performance and risk indicators to ensure the business maintains an even keel.

A well-designed KPI triggers a chain reaction of process improvements throughout the organization. In the British Airways example, late planes increase costs because airlines have to accommodate passengers who miss connecting flights; customer satisfaction declines because customers dislike missing flights; worker morale slips because workers have to deal with unruly customers; and supplier relationships are strained because missed flights disrupt service schedules and lower quality. When Lord King focused on late arrivals and departures, managers and staff figured out ways to change business processes and behaviors so they didn't receive a career-limiting call from the CEO. As a result, a single operational KPI monitored by the CEO created a ripple effect of process improvements that produced rapid and dramatic gains in performance.

Metric Components

Six Attributes. A performance indicator—whether an outcome or driver metric—has six primary attributes. Some of these attributes are visually depicted in Exhibit 11.1.

1. **Value.** A performance metric contains an actual value or number that represents the measurement of performance. The bold line in Exhibit 11.1 represents total sales each month. Ideally, if a user moved the cursor over a point in the line, the chart would pop up the actual numeric value for sales during that period.
2. **Time frame.** Performance metrics have a time frame by which they are measured. A time frame includes the end date when the target must be achieved and interval dates that are used to gauge progress along the way. The chart in Exhibit 11.1 measures performance monthly but does not indicate an end date.

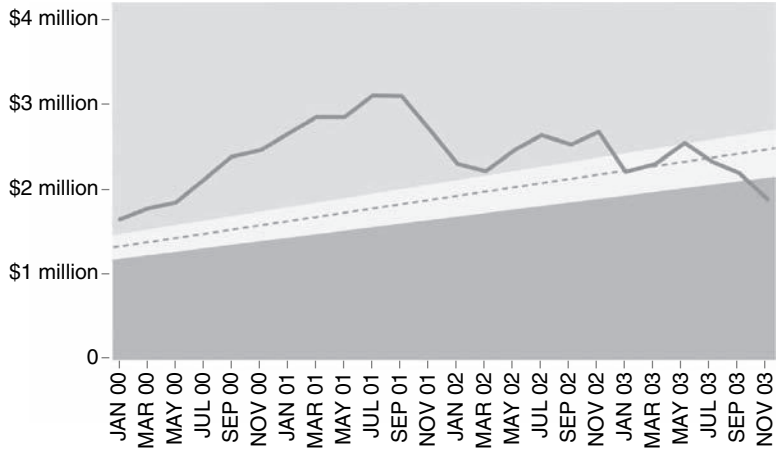

EXHIBIT 11.1 Sales Performance Chart

3. **Benchmark.** A performance metric compares actual values to some benchmark or baseline. A most common benchmark is the result from the same period last year. The benchmark could also be an arbitrary number, such as a budgeted or forecasted number. Or a metric could have multiple benchmarks (i.e., prior year, plan, and forecast), which is often the case with financial numbers. Benchmarks can also be external measures, such as the performance of an industry leader.
4. **Targets.** Each metric is associated with a target defined by executives in a planning or budgeting session. The target represents a tangible goal in context of the benchmark. For example: "We want sales to exceed last year's number by 10 percent this period." There are end targets, which typically are measured annually, and interim targets, which are measured in smaller intervals, such as daily, weekly, or monthly. The chart graphically depicts how the interim target increases at a steady rate each month. Presumably, the end target is defined annually.
5. **Ranges.** Targets are divided into ranges of performance, usually "above target," "on target," or "below target." Ranges commonly are based on percentages, although more complex rules can be used as well. For example, performance that is 5 percent above target might be considered "above target," while performance within 5 percent of target is considered "on target," and performance below 5 percent of target is deemed "below target."

   Most metrics teams apply three or four ranges to each metric, although some apply more. For example, the operations department of a U.S. railroad had 16 ranges for each metric. Organizational culture seems to be the biggest factor in defining the number of ranges. Most

executives and managers are fine with three or four, while engineers and operations staff prefer more.

6. **Visual encodings.** Finally, a dashboard visually encodes ranges so users can quickly grasp how close or far away they are from achieving the target. Most metrics encode ranges using traffic lights: green, yellow, and red. The chart shows graphically how targets and ranges change by a steady percentage each period.

Targets

Performance metrics derive much of their impact from the targets associated with them. The visual encoding of those targets into ranges causes users to "jump" when they see a red light. Ranga Bodla, senior director of enterprise performance software at SAP, defines five types of targets.

1. **Achievement.** With an achievement target, performance should reach or exceed the target in an upward direction. Anything over the target is desired but not required. An example is revenue, where you want to grow sales to at least 100 percent of the target, but more revenue is better as long as you maintain quality and customer satisfaction.
2. **Reduction.** A reduction target is the opposite of an achievement target. Here, performance should reach or exceed the target in a downward direction. Anything less than the target is valuable but not required. An example is overtime, where you need a certain amount of overtime to get work done, but less overtime is desirable as long as you don't sacrifice other goals, such as on-time deliveries.
3. **Zero.** A zero target is an extreme form of a reduction target in which the goal is zero. Basically, zero means that there is no circumstance in which a positive number is a good thing. Examples are number of injured employees or product defects.
4. **Absolute.** With an absolute target, the goal is to come as close as possible to the target value without exceeding it in either direction. For example, a manufacturing plant wants suppliers to ship goods as close as possible to an agreed-on delivery time. If a part ships too soon, the plant has to store it as inventory at considerable cost; if the part ships too late, the assembly line has to shut down until the part arrives; if it ships on time, the plant achieves maximum efficiency.
5. **Min/Max.** Min/max is like an absolute target but with more leeway in what is acceptable. Instead of shooting for a fixed number, performance should exist within a range of values. Anything above or below the range is not good. An example is mean time between repairs in which a manufacturer wants to deliver products that require repair only every 24 to 36 months.

Setting Targets. Business users should be able to define targets and thresholds for metrics using Boolean rules (i.e., if, then, or else) or by inserting values. The rules are applied to the metrics whenever the underlying data is updated. Many commercial dashboards, especially strategic dashboards, now enable authorized business users to define targets and thresholds via a point-and-click graphical interface.

Setting Alerts. The system should also let business users define personalized alerts when a metric exceeds a certain threshold. The alerts may be one and the same as the traffic lighting system (see Rohm and Haas example in Chapter 8) or a special icon or ticker tape animation that scrolls across the screen when the alert is triggered. The best alerts notify users through a preferred channel (e.g., wireless phone, e-mail, or land-line phone) so they don't have to log on to the dashboard to see significant events. In complex event processing systems, intelligent alerts can even trigger predefined actions, such as updating a database or sending e-mail messages. (See Chapter 6.)

Displaying Metrics

Conveying the Full Story. Performance dashboards display multiple attributes of a performance metric to give users a comprehensive view of performance. The number of attributes displayed depends on the application and user requirements. The goal is to provide the key information in a quick glance.

The chart in Exhibit 11.2, which comes from a balanced scorecard, shows six metric attributes: name, trend, status, target value, actual value, and variance percentage. Notice that it uses symbols to indicate trend and status so that color-blind people can interpret the encodings correctly.

The metric name usually is hyperlinked so users can click on it to drill down and view a time-series chart. Status measures performance against the target and usually is shown with a colored stoplight. Trend measures performance against the prior interval or another time period and often is

Name	Trend	Status	Target	Actual	Variance %
Finance					
Decrease European Discounts	+	✓	16.00%	15.53%	-3%
Keep APAC discounts < 17%	−	✗	17.00%	24.77%	31%
Keep N.A. discounts < 17%	−	✓	17.00%	15.66%	-9%

EXHIBIT 11.2 Displaying Performance Metrics

displayed using sparklines (a miniature time-series trend line without a scale), arrows, or in this case plus/minus signs. The actual and target values are self-explanatory and usually are displayed as text. Variance (not depicted) measures the gap between actual and target values and usually is displayed using text or a bullet chart (or a variance chart if users drill to a separate page). Variance percentage divides the variance against the target so users can interpret variance quickly.

Interpreting Variances. Variance, by itself, is tricky to interpret. Its meaning depends on the type of target. For example, most people automatically assume that a positive variance is a good thing. But if the metric uses a reduction target, a positive variance could be bad. For example, a positive 5 percent variance in overtime means workers are taking more overtime than desired, which dramatically increases costs. The difficulty of interpreting variance is why most metrics designers display a combination of attributes for each metric, including status, trend, and variance. The combination of attributes usually gives viewers an accurate impression of overall performance.

However, sometimes the attributes give mixed signals. For example, what does it mean if status is red but the trend is green? Should people take action to address the red status, or should they relax because someone has already addressed the problem and progress is being made? A dashboard that enables users to annotate metrics with comments would help in these situations. It would also behoove organizations to train employees how to interpret metrics properly, either through formal classroom instruction or an apprenticeship or mentorship.

Buried Scores. However, sometimes a performance dashboard disguises the true state of performance. This occurs when a compound metric displays a positive status even though one underlying metric is negative. The negative score of the single metric is outweighed by the positive scores of the other metrics. So, unless users dig below the surface of the compound metric, they will never see that a major problem is brewing.

The problem arises because most compound metrics today calculate overall status by averaging the percentage variance from target of underlying metrics. Some designers are circumventing the law of averages by giving added weighting to negative scores so they have more impact on overall status. Others will display an alert next to the compound metric that indicates that there is problem with one of its underlying metrics.

Number of Metrics

No Hard-and-Fast Rules. Many people ask, "How many performance metrics should we display on a dashboard?" In general, the fewer metrics the better, but there is no hard-and-fast rule about how many to display

on a single dashboard screen. Ultimately, you are trying to help businesspeople manage processes and solve problems. The less work they have to do to get the information they need to do their jobs effectively, the more useful they will find the dashboard.

There is a natural tendency among organizations to keep adding metrics and never delete any. Over time, the metrics lose their power to grab the attention of employees and focus their behavior on key value-added activities. "When people have too many metrics to track, the message gets blurred," says the director of customer management at an energy services provider.

The number of performance metrics depends on many factors: the type of users, the nature of the data, the complexity of the processes being monitored, and the skill and maturity of the users. In general, business executives and managers work effectively with about 5 to 7 metrics displayed as charts or more if the data are arranged in tables. The executive dashboard at Rohm and Haas (see Chapter 7) displays 12 metrics in a table with two associated line charts.

Engineers and service managers prefer densely packed dashboards with lots of metrics and text. Front-line workers, such as salespeople and field technicians, need simpler displays with a handful of easy-to-understand metrics as well as charts to help them calculate their incentive pay. 1-800 CONTACTS, for example, provides a simple, one-page dashboard for its call center salespeople. (See Chapter 7.)

Hiding Complexity. Monitoring complex processes from end to end may take dozens or more metrics. For example, a repair call resolution metric might require five submetrics to capture performance accurately at each stage of the repair process, from taking an order, scheduling the repair, and validating the repair to receiving customer payment. One metric will not shed enough insight to help managers know what part of the process is experiencing problems.

Obviously, it is difficult to display this many metrics on a single dashboard page without overwhelming users. One option is to group metrics by category and place them on different tabs (i.e., pages or folders) within the same dashboard. Another option is to house related metrics on different dashboards and then interlink the dashboards so users can navigate easily from one to the next.

For instance, if users want to examine a single departmental process from top to bottom, you can place strategic, tactical, and operational metrics on different dashboards and link them together. Or if the process crosses departmental boundaries—such as order to cash or procure to pay—you can link different departmental dashboards or create a single dashboard for the entire process, nesting metrics using tabs or radio buttons if needed.

If you aren't sure how many performance metrics to create, err on the high side. What does not get measured does not get done, and what does not get done can hurt the organization. The key to selecting metrics judiciously is to validate that they are aligned with strategic objectives and distribute them to performance dashboards at the appropriate level in the organization. Not all metrics need to appear on the top-level scorecard; most, in fact, should be delegated to lower-level ones.

According to research from TDWI, most organizations adhere to the less is more rule regarding KPIs. On average, organizations deploy 12 metrics on the top-level screen of their dashboards and 159 metrics overall. Operational dashboards have fewer top-level metrics (11) and fewer total metrics (141) than the average; tactical dashboards have the same number of top-level metrics as the average (12) but fewer overall metrics (110); and strategic dashboards have more top-level metrics (14) and total metrics (183) than the average. (See Exhibit 6.3.)

Metric Ecosystems

Balancing Metrics. When creating performance metrics, it's important to make sure they are aligned. Typically, one metric is not enough to drive performance in the right direction. It usually takes an ecosystem of metrics to ensure proper balance.

For instance, a call center may want to increase the productivity of its customer service representatives by measuring how many customers they talk to per hour. The unintended consequence of this measure is that it will discourage service representatives from spending enough time with customers to solve their problems. To avoid a precipitous drop in customer

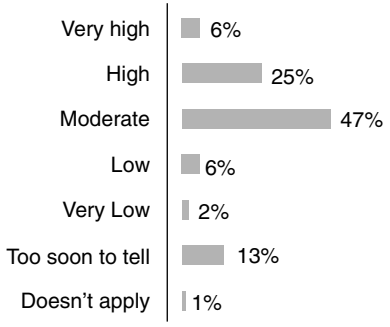

EXHIBIT 11.3 To What Degree Have Your KPIs Improved Performance in Your Organization?

satisfaction, the call center needs to balance its productivity metric with one focused on quality or customer satisfaction.

"We've seen our staff take unexpected actions to boost a metric that undermined other areas," says one director of BI.

Ironically, a balanced measurement ecosystem often consists of contradictory metrics. In our call center example, a balanced ecosystem might consist of a productivity metric (e.g., "Make as many calls per hour as possible.") and a quality metric (e.g., "Take as much time as you need to solve customer problems."). Humans have the capacity to manage this tension between opposites to some degree. But the metrics inevitably will spawn new procedures that resolve the tension more effectively. For example, the call center might divide agents into two groups, one that is incented to handle lots of easy calls quickly and forward harder calls to a second group, which is encouraged to spend more time with each customer.

Characteristics of Effective Performance Metrics

Organizations create many performance metrics, and most may exhibit the attributes we've discussed to this point. However, a well-constructed performance metric doesn't mean that it's an effective one. Exhibit 11.3 shows that a third of performance metrics have a "very high" or "high" impact on performance, while almost half (47 percent) have a moderate impact.

As I said at the outset, creating performance metrics is as much art as science. To guide you in your quest, here are 12 characteristics of effective performance metrics:

1. **Strategic.** To create effective performance metrics, you must start at the end point—with the goals, objectives, or outcomes you want to achieve—and then work backward. A good performance metric embodies a strategic objective. It is designed to help the organization monitor whether it is on track to achieve its goals. The sum of all performance metrics in an organization (along with the objectives they support) tells the story of the organization's strategy.
2. **Simple.** Performance metrics must be understandable. Employees must know what is being measured, how it is calculated, what the targets are, how incentives work, and, more important, what they can do to affect the outcome in a positive direction. Complex KPIs that consist of indexes, ratios, or multiple calculations are difficult to understand and, more important, not clearly actionable.

   "We hold forums where we show field technicians how our repeat call metric works and how it might impact them. We then have the

best technicians meet with others to discuss strategy and techniques that they use to positively influence the metric," says a director of customer management at an energy services provider.

3. **Owned.** Every performance metric needs an owner who is held accountable for its outcome. Some companies assign two or more owners to a metric to engender teamwork. Companies often embed these metrics into job descriptions and performance reviews. Without accountability, measures are meaningless.
4. **Actionable.** Metrics should be actionable. That is, if a metric trends downward, employees should know what corrective actions to take to improve performance. There is no purpose in measuring activity if users cannot change the outcome. Showing that sales are falling isn't very actionable; showing that sales to a specific segment of customers is falling compared to others is more actionable.

   Actionable metrics require employees who are empowered to take action. Managers must delegate sufficient authority to subordinates so they can make decisions on their own about how to address situations as they arise. This seems obvious, but many organizations hamstring workers by circumscribing the actions they can take to meet goals. Companies with hierarchical cultures often have difficulty here, especially when dealing with front-line workers whose actions they have historically scripted. These companies need to replace scripts with guidelines that give users more leeway to solve problems in their own novel ways.
5. **Timely.** Actionable metrics require timely data. Performance metrics must be updated frequently enough so the accountable individual or team can intervene to improve performance before it is too late. Some people argue that executives do not need actionable or timely information because they primarily make strategic decisions for which monthly updates are good enough. However, the most powerful change agent in an organization is a top executive armed with an actionable KPI.
6. **Referenceable.** For users to trust a performance metric, they must understand its origins. This means every metric should give users the option to view its metadata, including the name of the owner, the time the metric was last updated, how it was calculated, systems of origin, and so on. Most BI professionals have learned the hard way that if users don't trust the data, they won't use them. The same is true for performance metrics.
7. **Accurate.** It is difficult to create performance metrics that accurately measure an activity. Part of this problem stems from the underlying data, which often need to be scanned for defects, standardized, deduped, and integrated before displaying to users. Poor systems data create lousy performance metrics that users won't trust. Garbage in,

garbage out. Companies should avoid creating metrics when the condition of source data is suspect.

Accuracy is also hard to achieve because of the way metrics are calculated. For example, a company may see a jump in worker productivity, but the increase is due more to an uptick in inflation than internal performance improvements. This is because the company calculates worker productivity by dividing revenues by the total number of workers. Thus, a rise in the inflation rate, which artificially boosts revenues—which is the numerator in the metric—increases worker productivity even though workers did not become more efficient.

Also, it is easy to create metrics that do not accurately measure the intended objective. For example, many organizations struggle to find a metric to measure employee satisfaction or dissatisfaction. Some might ask users in surveys, but it's unclear whether employees will answer questions truthfully. Others might use the absenteeism rate, but this might be skewed by employees who miss work to attend a funeral, care for sick family members, or stay home when child care is unavailable.

8. **Correlated.** Performance metrics are designed to drive desired outcomes. Many organizations create performance metrics but never calculate the degree to which they influence the behaviors or outcomes they want. Companies must refresh performance metrics continually to ensure they drive the desired outcomes.
9. **Game-proof.** Organizations need to test all performance metrics to ensure that workers can't circumvent them out of laziness or greed or go through the motions to make a red light turn green without making substantive changes. "Users always look for loopholes in your metrics," says one BI manager. To prevent users from fudging customer satisfaction numbers, one company hires a market research firm to audit customer surveys.
10. **Aligned.** It's important that performance metrics are aligned with corporate objectives and don't unintentionally undermine each other, a phenomenon called suboptimization. To align metrics, you need to devise them together in the context of an entire ecosystem designed to drive certain behaviors and avoid others.
11. **Standardized.** A big challenge in creating performance metrics is getting people to agree on the definitions of terms, such as sales, profits, or customer, that comprise most of the metrics. Standardizing terms is critical if organizations are going to distribute performance dashboards to different groups at multiple levels of the organization and roll up the results. Without standards, the organization risks spinning off multiple, inconsistent performance dashboards whose information cannot be easily reconciled.

12. **Relevant.** A performance metric has a natural life cycle. When first introduced, the performance metric energizes the workforce and performance improves. Over time, the metric loses its impact and must be refreshed, revised, or discarded.

    "We usually see a tremendous upswing in performance when we first implement a scorecard application," says a program manager at a major high-tech company. "But after a while, performance trails off. In the end you can't control people, so you have to continually reeducate them about the importance of the processes that the metrics are measuring or you have to change the processes."

Designing Effective Metrics

Gathering Requirements

Chapter 10 discussed how to sell, manage, and sustain a dashboard project and examined the different teams required to design, develop, and manage the project. This section provides additional detail about how the KPI team gathers and prioritizes requirements for creating performance metrics and creates a blueprint for development that the technical team can implement.

KPI teams use a variety of methods to capture requirements for dashboard metrics. Chief among these is to "review and prioritize metrics used in existing reports" followed closely by "interviews and joint design sessions" selected by 55 percent and 51 percent of respondents respectively to a 2008 TDWI survey. (See Exhibit 11.4.)

EXHIBIT 11.4 Approaches to Gathering Metric Requirements

Leverage Existing Metrics. Rick Sherman, president of Athena IT solutions, a Boston-based BI consultancy, says, "In the discovery process, you learn that there are almost no green fields. If an executive has an idea for a metric, the company is probably already measuring it somewhere in a spreadsheet or report." Sherman first interviews executives and managers and then the analysts who create existing reports. He identifies gaps between what the business wants and existing metrics deliver, and then works with the data to ensure it meets user requirements for timeliness, accuracy, and cleanliness, among other things.

Start with Objectives. Robert Kaplan, cofounder of the balanced scorecard methodology, takes a different tack. He says executives should take all their existing measures and "lock the [metrics] in a room" while they hash out the strategic objectives that they want to drive the organization. Once they've completed their strategy and mapped out interdependencies among objectives, then they can go "unlock the room" and see which existing metrics can be reused. This process ensures that strategy drives metrics, not the reverse.

"Don't go hunt for the right KPIs," says Kaplan. "That is not the answer. A successful balanced scorecard project is not a metrics project. It's a strategy and change management process."

Bill Barberg, president of Insightformation, a strategy management consultancy and software provider, estimates that half of the metrics in a typical balanced scorecard solution already exist and half are designed from scratch. He says companies have already defined most of their executive-level and operational metrics. The fertile ground for new metrics comes from the process of creating cross-functional strategic themes. Themes are collections of related objectives in a strategy map and represent an organization's overall strategy.

Interviews. Regardless of the approach, most KPI teams use interviews to gather requirements. Interviews usually are done by business analysts who ask a series of predefined questions to top executives about the business strategy, objectives, goals, and expectations for the project, among other things. The analysts then gather additional detail by interviewing mid-level managers and subject matter experts who can fill in the details of specific processes, identify data sources, and point out metrics used in current reports, what those metrics mean, and how they are calculated.

Some KPI teams like to create use cases to discover metrics and test their validity. These use cases can be solicited by asking users to create day-in-the-life-of stories about how they manage a particular process. For example, a use case might define a series of steps that a business user goes through to manage the sales pipeline, for example. Use cases are good ways to understand the core processes performance metrics are designed to measure.

Group Sessions. Some KPI teams supplement individual interviews with facilitated group sessions. The group sessions can help teams achieve consensus on the key purpose and goals of the dashboard project and define key metrics. The group might use brainstorming techniques to generate driver metrics (see "The Five Whys" section) or work through multiple use cases to flush out core metrics. Neal Williams has developed a 9-step facilitated workshop to help KPI teams create a metrics blueprint that developers can use to build dashboards. (See Spotlight 11.2.)

Spotlight 11.2 KPI Design Workshop

Neal Williams, president of Corda, developed a nine-step workshop to accelerate the deployment of performance dashboards. Before then, groups would purchase Corda's dashboard software and start implementing the tool before they knew what they wanted to measure. Says Williams, "Many teams get a mandate to implement a dashboard and then struggle through many product iterations before they figure out what to measure and how to display it."

The two- to three-day workshop enables a KPI team comprised of three to five businesspeople and an IT professional to create a detailed blueprint that comprises a complete set of specifications. Williams says the blueprint reduces the number of product development iterations by 80 percent to 90 percent and significantly accelerates dashboard delivery times.

The KPI team writes the output of each step *in English sentences*. There is no coding during the workshop, and prototypes are done on a whiteboard or sketch pad. Once the workshop is finished, the team hands the blueprint, including prototype sketches, to the technical team, which uses it to develop the dashboard and metrics.

The nine steps are:

1. **Frame.** Write a framing statement that defines one or more strategic objectives that the performance dashboard is designed to support. Each dashboard project tackles a maximum of three to four objectives, each of which will generate multiple metrics.
2. **Elaborate.** Write business questions for each strategic objective that you want the dashboard to answer.
3. **Define metrics.** Define measures and dimensions that will help you answer each question; for example, show sales by product, region, and channel.
4. **Define targets.** Describe benchmarks and targets for each metric.

5. **Diagram drill paths.** Diagram the drill paths needed to answer each question. What information should appear on the top page? How might users navigate from a summary to detailed views? These diagrams can be created with pencil and paper if desired.
6. **Define behaviors.** Define type of indicator (outcome or driver) for each metric and describe the behaviors that each is supposed to drive. Define the actions users should take to achieve the objective the metric supports.
7. **Check data.** Check the availability and condition of data for each performance metric.
8. **Check for compliance and balance.** Check to make sure metrics adhere to standard definitions and rules in use across the organization. Make sure each metric doesn't undermine key objectives or conflict with other metrics, and that it cannot be easily circumvented by users.
9. **Assign owners.** Assign a business owner and data steward for each metric. The business owner is responsible for the definition and meaning of the metric while the data steward certifies that the data populating each metric is accurate.

Requirements Forms. To guide business analysts during interviews, most KPI teams create a requirements form to capture responses from business managers in a standardized way. This ensures that analysts ask a consistent set of questions and gather a comprehensive set of information that is easily synthesized and standardized.

One high-tech company, for example, uses two forms to define metrics for its strategic dashboard: one to gather business requirements and another to define technical specifications. The business requirements form asks for a general description of the metric, how it aligns with corporate strategy, the name of the metric, its owner, its target and stretch goals, and how the metric is calculated, among other things. (See Exhibit 11.5.)

The technical specification document provides technical details for each proposed metric. For example, it asks for data sources and formats, extraction logic, scorecard layouts, target specifications, analytical layouts (including columns, rows, data types, formats, and formulas), chart views, and security requirements. Most important, the form asks for the names of the business and technical owners of the metrics so project team members can follow up with additional questions, if needed.

Surveys. Some KPI teams use surveys to gather user requirements. This method is a more efficient way to gather input from large numbers of people. It also gives users more time to provide thoughtful responses

EXHIBIT 11.5 Request Form

	**Metric Request Form** **Submitted By: _____ Date: _____**
Business or Function	*What business or function do you request a metric for?*
Region/Country Scope	*What is the geographical scope of the metric?*
Metric Perspective	*What balanced scorecard perspective does the metric fit in: Customer, Financial, Internal, Learning?*
Metric Title	*Give a brief name to the metric (fewer than 20 characters)*
Metric Description	*Describe the metric in business terms.*
Business Justification & Strategic Importance	*Specify whether the metric is a strategic metric or an operational excellence measure, and justify your selection. How does the metric measure progress toward strategy execution?*
Metric Business Owner, Subject Matter Expert, Business IM Owner	*Specify the owners of the metric, either from the business or the function that will be measured on the results (can be name or job title). Also, who collects, reviews, approves, and reports the data?*
Metric Goals	*Specify both the target and stretch goals for the metric (indicate over what time period)* *Also, how is the goal selected and who approves the goal?* **Target: Stretch:** **Goals-setting process and approval from:**
Definition, Calculation, and Criteria	*How is the metric calculated? What criteria are used? Identify any differences between worldwide or sub-region definitions.*
Data Source and Availability	*What is the data source for the actual results, and how are data collected? When is the data available (i.e., which workday, every six months, annually, etc.)?*

Supporting Reports	*What detailed reports are available to support the metric results?*
Related Metrics	*List upstream metrics influenced by this metric. List downstream metrics that have influence on this metric.*
Additional Information	*Specify additional information related to the metric.*
Status	*Status of the metric request from the PMMS worldwide program office team (approved, pending additional info), targeted implementation date, etc.*

to questions compared to a one-on-one interview. Users can skim through the survey questions and collect their thoughts before submitting an answer. In contrast, a one-on-one interview puts users on the spot, unless they have received the questions in advance and spent time reviewing them prior to the interview.

Prioritizing Metrics. Once the KPI team has gathered requirements, it needs to cull the number of performance metrics to a manageable number. The easiest way to do this is to create a list of criteria to evaluate each proposed metric. For example, the team might want to use the 12 characteristics of effective performance metrics defined earlier or some other relevant criteria. Cisco, for example, uses 5 criteria to evaluate potential metrics: scalability, normalization, accountability, goalable, and actionable. (See Chapter 9.) The team should weight each criterion on a scale from 0 to 1 and then score each metric on a scale of 1 to 5 with 5 equal to "closely meets the criterion." Exhibit 11.6 is an example of template used to evaluate performance metrics.

Validating Metrics

Once a KPI team has defined and prioritized its performance metrics, it should present them to the individuals whose performance will be measured to get their feedback. This is a critical part of the process that shouldn't be shortchanged.

Nuances. Sometimes it is difficult to define a metric that accurately captures the nuances of a business process. For example, it took Direct Energy Essential Home Services considerable time to create a repeat call metric. The metric tracks the efficiency of field service technicians who repair residential furnaces and air conditioning systems and is designed to

EXHIBIT 11.6 Culling Performance Metrics

	Criteria	Linked to Strategy	Drives Behavior	Action-oriented	Easy to Understand	Condition of Data	Standard Definition	TOTAL
Measures	Weight	1.0	1.0	0.8	0.7	0.5	0.7	
Financial perspective								
Measure 1 Measure 2 …								
Customer perspective								
Measure 1 Measure 2 …								
Internal Operations								
Measure 1 Measure 2 …								
Learning and Growth								
Measure 1 Measure 2 …								

Instructions: Rate each measure on a scale from 1 to 5, with 5 = "Closely meets requirement."

Source: Adapted from Karen Degner.

encourage technicians to fix problems on an initial site visit instead of making repeated trips, which occurs if technicians don't stock the right parts in their repair vehicles.

When the KPI team presented the metric to service technicians, the team discovered that they not only fix equipment but are encouraged to sell it as well. The team then struggled to define when a repeat call is a good or bad thing and how to measure that. If the technician makes a repeat trip to bring literature about a replacement system, is that good or bad? What if a homeowner only permits a technician to make minor repairs to an aging system, but then the system breaks shortly afterward and the technician returns to fix or replace it?

To flesh out these nuances, KPI teams need to get feedback from the people who are closest to the process the metrics measure. These people know best whether a performance metric accurately represents a process as well as how easy it might be for employees to circumvent or game the metric. A KPI team that implements performance metrics without getting feedback and buy-in from front-line staff risks stirring up a hornet's nest of trouble, especially if incentive pay is tied to metric results.

Standardizing Terms

Gaining Consensus. The hardest part about creating metrics is defining the meaning of common terms that they are based on. It is ironic that the terms used most frequently in an organization are the hardest to pin down. For example, most organizations have difficulty defining exactly what a customer is and is not. Is a customer someone who has purchased product sometime in the past or just within the last 12 months? Is a customer someone who has registered on a Web site and downloaded documents but has never purchased anything? In the same way, every industry has difficulty defining core terms. An airline often can't get agreement on the definition of a flight segment and a manufacturing firm has difficulty achieving consensus on the meaning of the word *part*.

The challenge in standardizing terms increases with the scope of the project and the number of departments supported by a performance dashboard. It's easy to standardize terms and rules within a single department because everyone views the business in much the same way. But when you cross departmental boundaries, achieving consensus on a common term, such as *sales*, is surprisingly difficult. For example, a "sale" to the finance department occurs when payment is received and deposited in the bank; but the sales department sees it differently: A "sale" occurs when a customer signs a purchase order. And marketing believes a "sale" happens when a customer pays for an item with a credit card or check.

CEO Intervention. The more groups and people, the more divergence there will be in the definitions of terms, rules, and calculations that compose a metric. Sometimes the only way to resolve these differences is for the CEO to call a meeting of top executives and force them to hash out standards with which they all can live.

"We have two distinct businesses, commercial and government, and the measurements each uses are very different, which makes it very challenging to develop corporate-wide standards," says a senior manager of reporting at a major U.S. consultancy. "We've had strong backing from our CEO to make this work and we've made a lot of progress. But even with that, it takes a lot of time and our partners recently decided to postpone trying to resolve the issues until a later time."

A high-tech company faced a similar situation. Says the program manager:

> *We wanted a worldwide metric for cost reduction, and we discovered that the operation and finance people had 32 ways to measure cost reduction. Some of these were duplicates, others measured different facets of costs. The project team arranged a meeting between two top financial executives, and they agreed to standardize on 6 metrics for cost reduction.*

Collecting Data

Poor-Quality Data. Sometimes the data to support a metric simply do not exist, or they are in poor condition. The most well-defined performance metrics are irrelevant if there are no data to populate them. To get a handle on data issues early in the process, the KPI team needs to appoint a systems analyst to participate on the team and scout out data sources for proposed performance metrics. If the analyst reports that the data don't exist or are full of missing or invalid values, duplicate records, or inconsistencies that might take weeks or months to clean up, the KPI team has a decision to make: If the metric is important enough, the team can kick off a major data reconditioning project or build a new system to automate data collection. If the proposed metric isn't that critical, the team can revise or scrap it. The key here is that the KPI team builds metrics in context of available data so that they don't waste time and money defining metrics that can't be supported.

Distributed Data. Another common problem is that the data required to populate a metric are spread across multiple systems that capture and format data differently. Even if the distributed data are in good condition, which they usually are not, the project team must expend significant effort to integrate the data in a consistent fashion. Thus, the dashboard project

turns into a data warehousing project. Obviously, dashboards that leverage metrics in an existing data warehouse are much easier to build and deploy than dashboards that must acquire data natively.

"Data integration is critically important but it is often overlooked, especially by the business side of the house," says a manager of performance management. "Businesspeople often don't know there is a problem until the technical team reports back that it can't deliver all the relevant performance metrics. The larger the organization, the bigger the data integration challenge."

Setting Targets

One of the toughest challenges of a KPI team is to set realistic targets. Targets should not be so challenging that they discourage workers, nor should they be too easy, which leads to complacency. Often setting targets is a matter of trial and error. However, it is best to identify realistic targets up front to avoid problems later on. Many times this requires an organization to establish formal policies and procedures for setting targets.

Get Buy-in. It is important not to set targets in a vacuum. Executives and managers often are tempted to set targets based on their own experience and knowledge of the business. But such unilateral goal setting does not engender goodwill among the staff responsible for achieving the goals. It is critical that the KPI team gathers input from employees to understand what targets are reasonable and gains their buy-in to the project. Ultimately, employees are doing the work and should feel that the goals are reasonable.

Cisco, for example, has a Corporate Quality team that negotiates performance targets for customer satisfaction scores with every department in the company. The targets are then reviewed by an internal cross-functional board and sent to the executive team for approval. Such elaborate negotiation is necessary when a portion of employees' incentive pay is based on achieving the targets, which is the case at Cisco. (See Chapter 9 for more details.)

Calculating Metrics

A final thing that metrics designers need to consider is in what layer of the architecture they should write metric calculations. Some metrics may be quite complex, requiring aggregating data across multiple dimensions and levels in an organizational hierarchy or combining one or more metrics in a ratio or compound measure. In addition, most dashboards display metrics over time and calculate variance from period to period. Such time-series calculations can consume a lot of processing power and slow query performance significantly.

When physically creating metrics, dashboard designers have three choices. They can define and calculate metrics in (1) the dashboard, (2) a BI semantic layer, or (3) the database.

The easiest and most natural approach for developers is to create metrics within dashboard applications since most dashboard products support calculation and rules engines. The downside is that embedding rules inside a dashboard creates a spreadmart (dashmart?) since every developer is likely to define the same metrics differently. It also creates a maintenance nightmare: When a developer wants to revise a metric, he or she needs to update the metric in every dashboard that uses it. Finally, dashboards that calculate metrics on the fly in the application layer will suffer poor performance, especially when they display time-series charts for complex metrics.

The second option—a BI semantic layer—abstracts the metric from the dashboard application and makes it available to any application that can read the semantic layer. Designers write the metric once and deploy it many times. The fact that the metric is maintained centrally ensures that any revisions are automatically available to downstream applications. The downside here is that many dashboard applications don't support BI semantic layers and performance still can be an issue, especially if queries to populate the metric are executed dynamically.

The third option involves executing the metric inside a database. This requires data designers to incorporate the metric inside the data model and then create transformations to populate the model with relevant data and apply the appropriate calculations. Like a BI semantic layer, this method manages the metric centrally, which ensures consistency. It also provides better performance because designers can calculate the metric in advance as data are transferred from source systems to the database and store the time-series data locally. The downside is that changes to the metric require designers to change the data model, which can be tricky and time consuming.

Summary

Performance metrics are powerful agents of organizational change, but creating them is challenging. To create effective performance metrics, organizations need to understand the components of a performance metric and create a KPI team that gathers requirements, prioritizes, validates, and standardizes metrics and sets realistic targets after getting buy-in from individuals whose performance will be measured.