

## **MEASURING PERFORMANCE – IDEAS FOR THE SERVICE SECTOR**

Faced with stiffening competition, increasingly demanding customers, high labor costs, and, in some markets, slowing growth, service businesses around the globe are trying to boost their productivity. But whereas manufacturing businesses can raise it by monitoring and reducing waste and variance in their relatively homogeneous production and distribution processes, service businesses find that improving performance is trickier: their customers, activities, and deals vary too widely. Moreover, services are highly customizable, and people—the basic unit of productivity in services—bring unpredictable differences in experience, skills, and motivation to the job.

Such seemingly uncontrollable factors cause many executives to accept a high level of variance—and a great deal of waste and inefficiency—in service costs. Executives may be hiring more staff than they need to support the widest degree of variance and also forgoing opportunities to write and price service contracts more effectively and to deliver services more productively.

As with any task or operation, to improve the productivity of services, you must apply the lessons of experience. Consequently, measuring and monitoring performance (and its variance) is a fundamental prerequisite for identifying efficiencies and best practices and for spreading them throughout the organization. Although some variance in services is inescapable, much of what executives consider unmanageable can be controlled if companies properly account for differences in the size and type of customers they serve and in the service agreements they reach with those customers and then define and collect data uniformly across different service environments. To do so, it is necessary to bear in mind a few essential principles of service measurement.

- First, service companies need to compare themselves against their own performance rather than against poorly defined external measures. Using external benchmarks only compounds the difficulties that service companies face in getting comparable measurements from different parts of the organization.
- Service companies must look deeper than their financial costs in order to discover and monitor the root causes of those expenses. This point may seem self-evident, yet many companies fail to understand these causes fully.
- Finally, service companies must set up broad cost-measurement systems to report and compare all expenses across the functional silos common to service delivery organizations. The goal is to improve the service companies' grasp of the cross-functional trade-offs that must be made to rein in total costs.

None of these principles is easy to implement. Top executives are likely to face resistance from managers and frontline personnel who insist that services are inherently random and that service situations are unique. Managers who have grown used to the protection that lax measurement affords may be reluctant to view their operations through a more powerful lens. But only by adopting these principles and implementing

rigorous measurement systems throughout the organization can service executives begin to identify reducible variance and take the first steps toward bringing down costs and improving the pricing and delivery of services.

### **Why variance is difficult to measure**

Executives who launch variance-measurement programs in a service business are often surprised at the level of difference they discover among similar sites and groups within their own organization, let alone when they compare one company with another. In general, a company's metrics are not uniform across its business units, so that, for example, one group in a call center may regard all calls on a given issue as a single case, while another logs every call separately. A top executive with a background in consumer goods (where items are similar and thus comparable) assumed control of a service business and was shocked to find that the variance of key metrics among similar sites ranged from a factor of 2 to 30. Site managers explained this vast range by asserting that every site was different—and, according to their metrics, they were right.

### **Services are different**

To make meaningful comparisons, companies have to identify the sources of difference in their businesses and devise metrics that compare these businesses meaningfully. The considerations that show up frequently include the most obvious differences among jobs and groups, such as regional variations in labor costs, local geographies and difficulties in reaching accounts, the workload mix (for example, repairs versus installation), and differences in the use of capital (whether equipment is owned or leased by the company or owned by the customer). Several other major issues come into play as well.

**Service-level agreements.** The more types of services a business offers, the more variability it can expect in its agreements. The metrics for a help desk that provides customer support for 5,000 users in a 9-to-5 office are very different from those for a help desk that supports logistics in a round-the-clock industrial environment. Even when offerings are similar, variance can be introduced locally through the way contracts are interpreted. In one IT-outsourcing company, two desktop support accounts with service-level agreements that specified an eight-hour response time had very different cost metrics. When asked why, the manager of the poorly performing account said that, despite the contract's limits, "If we don't answer within the hour, our client goes ballistic." The written service-level agreement had been trumped by an unwritten one that was costing real money.

**Environment, equipment, and infrastructure.** Each customer's environment has unique aspects that are difficult to measure. A logistics provider will see huge differences between managing a big, automated warehouse and a small, simple one. Field services that support industrial systems must contend with many generations of equipment and upgrades at customer sites. Some clients have their own on-site staff to support service, while others may be difficult even to reach. Given the range of possibilities, it's usually not very helpful simply to measure the average cost of a service call.

Work volume. Size is a major reason for the wide variance among accounts and business units. Interestingly, managers of both small and large accounts claim that size makes their particular metrics worse. Both have a point: large accounts should benefit from scale, but in general they are also more complex, and that drives costs back up. Volume needs to be considered, but only in tandem with other patterns (including scale benefits and the breadth of work) that help explain costs.

### **The data problem**

Underlying all of these problems is an inability to identify what must be measured and how to normalize data across different environments. Even when companies know what to measure, they struggle to achieve accuracy. Data are rarely defined or collected uniformly across an organization's environments. A service call involving the installation of two elevators, for example, could be measured as a single installation in one part of a company and as two in another.

Contributing to this ambiguity is the fact that data collection is usually driven by the requirements of financial cost reporting, which often fails to shed light on ways of boosting performance. Accountants for an IT services company may need to know the cost of each server, for instance, but an executive looking to reduce variance would also need to know the number of service incidents by server type and the time spent on each incident. Variance in demand drivers is also important: did the number of calls to a help desk rise because more users bought a product, for example, or because it changed? Financial metrics might fail to detect this important distinction.

### **Principles of service measurement**

Many executives don't understand how to measure and manage what appear to be unique activities, and they confuse correctable performance variance with irreducible environmental variance. Embracing three principles that identify variance and allow for meaningful comparisons can help executives overcome these difficulties.

### **Use internal benchmarks**

While a company must know what its peers are achieving, it's a mistake to measure its performance against the competition: these benchmarks are typically just samples of data with little explanation behind them. Companies that use external benchmarks are often frustrated to find themselves off by a factor of five to ten, positively or negatively.

Using external benchmarks compounds the internal difficulties that service companies face in normalizing activities and the data that define them. Consider a measure such as costs per unit of information processed: some companies include allocated costs, such as corporate overhead and salaries; others don't.

Internal benchmarks deliver more detailed metrics, allowing a company to find its own best practices and to see where and how they are achieved. It can then have access to

all relevant information to assess differences among business units and accounts. In defining internal benchmarks, for example, a company can determine which costs are included or how asset costs are allocated—details that get lost in external benchmarking. A company can see what's really possible within the organization by using its own benchmarks.

A cost tree with detailed metrics is an important tool to help companies define internal benchmarks. External metrics might deliver numbers on the top level of the tree, but only by developing internal trees for each service line can a company begin to understand its true cost drivers. A tree allows a manager to compare the performance of different accounts against similar metrics and also to calculate which improvements will have the most impact on the top-level figure. Once a team has gathered cost data throughout the tree, for example, it could target opportunities to cut costs and calculate which efforts would have the most impact on the bottom line. Creating cost trees can also help companies write better service agreements that exclude unprofitable activities or generate more revenue where service costs warrant it.

### **Measure cost drivers**

Even after service companies begin to define and capture the detail that lies beneath the top level of the cost tree, they still need to discover the underlying cause of each expense. Measuring only the cost of repair calls, for instance, probably wouldn't reveal whether they all stem from a single poorly built product, which could be improved or sourced differently for less than the cost of the repairs, or from factors such as variability in the performance of repair teams. Better measurements look at cost drivers, such as cost per employee (a resource metric), incidents per employee per day (a productivity metric), or—in a product-based service business—the number of incidents per product (a volume metric).

Of course, companies must also omit allocated costs, which can confuse the issue. A business unit's support infrastructure, for example, could include human resources, physical plant, and product engineering, all of which must be considered from a financial point of view. But such costs do little to determine productivity and are something of an obstruction when companies try to spot variance and waste. Once these obstacles are removed, managers can stop trying to cut costs that may be beyond their control and instead address the drivers they can improve. Before measuring the financial costs, it's often helpful to measure the items and events that drive costs, such as people, machines, incidents, service calls, and change orders.

### **Measure deep and broad**

When service companies try to measure only their selected costs—rather than taking a comprehensive approach—they are often surprised to see that their teams hit every budget target while still losing money. That's because services are fungible, and it's easy to measure the wrong things or to shift costs, intentionally or not, to unmeasured areas.

Consider the case of a cable company that was trying to reduce the resolution times of its help desk and service calls. After setting goals, managers saw resolution times shrink, but total service costs were rising. In this case, help desk representatives, eager to meet their goals, spent less time trying to resolve problems remotely. After asking only a few questions, these employees referred cases to field service reps, who were happy to have a series of fast and easy calls to boost their own metrics. Unfortunately, the number of field service calls, which are far more expensive than help desk calls, rose dramatically. To resolve this problem, management combined call centers and field services into a single cost tree and monitored the percentage of calls passed from the one to the other, as well as the time spent on each type of call. Managers then encouraged the call center reps to spend more time trying to resolve difficult calls before passing them along to field services, thereby increasing the average call time but helping to reduce total costs. Thus a critical purpose of any cost tree is to yield insights about how better (or worse) performance in one area of the tree might affect another.

### **Setting up measurement systems**

With these principles in mind, executives can begin to define their metrics, collect data, and implement processes that will drive their efforts forward.

### **Build the tree and choose your metrics**

Cost trees should be detailed enough to spot efficiency problems and broad enough to be comparable across operating units. Once companies have identified the allocated costs and cost drivers, they can begin to build the cost tree. Broad input from the field (line managers, engineers, field and service reps) is vital, along with input from senior executives, who are generally better able to focus on the total costs required to deliver a service to customers. In this way, the tree captures all the costs of (and details on) the most important cost drivers. The tree should also be constructed to compare key metrics across a range of environments—for example, all call centers, whether they operate 24 hours a day or 9.

As the data arrive, management will want to monitor the top level of the tree as well as the key metrics below. In most cases, we find, three to five metrics monitor 80 percent of the variance in costs.

### **Collect with care**

Without clearly defined metrics and knowledgeable people to support the gathering of data throughout the organization, companies can spend too much time cleaning up messy data. Training and improved processes can alleviate this problem.

Managers should review the data collection rules and templates with the people who develop them—usually employees from different regions or accounts. Even with new procedures in place, however, there will be much room for interpretation. It's therefore helpful to show not only how data should be collected and entered but also how users

occasionally misinterpreted these processes in the past—an approach that sheds light on gray areas the rules might not address. Guidelines for identifying problems early on can save time later. It's also important to establish boundaries beyond which suspect metrics should be investigated. One service company, whose teams handled from two to five service calls a day, wondered why one of its teams was reporting an average of only a single call. It found a good reason: the account belonged to a prison system, where security procedures made each visit a daylong affair.

Reviewing data collection in the early stages of implementation can help to ensure that procedures are followed. Equally, sharing reports with regional and account leaders gives them an early view of their standing and can help identify unusual patterns in the data.

### **Institutionalize measurement**

Managers accustomed to tracking costs in accordance with accounting needs will have to understand these new metrics and make them consistent throughout all levels of the organization. Periodic reviews, whose frequency should be based on the availability and shelf life of data, are essential for individuals and work groups. Visible interest from senior management—such as sending an executive vice president to attend a regional metrics review—promotes a strong message to everyone that a company is intent on identifying variance and improving service performance. Compensation should be tied to these metrics.

### **What's measured can be managed**

Once executives have learned to measure the variance inherent in service companies, they can begin to manage processes to eliminate waste, to improve the delivery of services, to price services more accurately, and to write better contracts. Although a company can do many things to control the variance of its service delivery, most of them fall into three main areas: managing demand, standardizing environments, and applying appropriate resources to tasks.

Managing demand offers the biggest potential for improvement. Cost trees help managers identify the sources of demand for services—sources that might include faulty products, poorly performing service units, or any number of other causes. Some fixes must be made within the organization (better training, better products, automated-response systems); others depend on shaping the behavior of customers (for instance, by offering tools and guidance to help them resolve problems themselves).

Standardizing operating environments requires the most discipline, since salespeople are strongly tempted to sell as much customization as a client wants. Standardization can yield enormous results: in addition to raising productivity, it helps the workforce become more flexible because people can transfer with less retraining. Where possible, companies should standardize not only service product lines and tasks but also the work environments of employees and the equipment they use to deliver services. Scripted

routines help eliminate errors and allow employees to emulate high performers. Furthermore, clearly defined programs limit overdelivery, a common problem in service companies.

What's more, identifying cost variances can help companies allocate their human resources more effectively. In general, it's more productive to handle problems with the least expensive resources that can resolve them: calling in experts or sending out field technicians increases costs and slows response times—and therefore makes customers less satisfied. Metrics on costs per call or device demonstrate the benefits of using less expensive labor, thus encouraging companies to keep requests upstream and to place first responders (often a call center) in less costly regions to further increase savings and productivity.

Finally, companies that have a better picture of where costs are incurred can price services more accurately to avoid losing revenue on unprofitable activities. They can write better contracts that take into account cost drivers hitherto written off as inescapable variance.

As services become an ever larger part of the global economy, managers are rightly looking for ways to improve productivity and efficiency. Services may be more difficult to measure and standardize than the manufacture of products, but executives should not abandon hope. Adopting the principles set forth in this article will help companies improve the delivery, pricing, and sales and marketing of services.