Support Center UNIVERSITY www.SupportCenterU.com

Considerations When Benchmarking Customer Support Expenses in the Software Industry

Introduction

What is a reasonable and appropriate level of expense for a software company to spend providing "support" for its customers?

This is a question that has plagued managers for many years…for as long as there has been software and customers to support. It is also a question that cannot easily be answered. This paper will attempt to provide insight into the most common methods of measuring support costs. I'll discuss the 'pros & cons' of each method. I'll also discuss the potential for each as a means of benchmarking.

Overview

Over time, a few methods of measuring and/or benchmarking support costs have become more popular. Each of these methods has strengths and weaknesses (pros & cons), that are discussed in more depth below. There are numerous other financial measures that companies can utilize to evaluate support performance. However, any effort to compare or 'benchmark' support costs must inevitably depend on one or more of these common methods.

The four primary methods of measuring and/or benchmarking support costs:

- ✓ Cost per call (and/or incident)
- ✓ Cost per minute (several variations)
- ✓ Cost in relation to overall revenue
- ✓ Cost in relation to support revenues

Benchmarking of support costs is the topic of numerous industry studies and reports. The primary reason for this continuing effort is 'market demand'. As previously noted, there is a continuous demand from senior management of software companies to understand if they are investing the right level in customer support. Also, there is no single source of credible information or even multiple market leaders in providing this data. And so, a hungry market continues to search for reliable, credible, and appropriate data to use in benchmarking efforts.

Note: When utilizing any method, different functions that may or may not be present or similar at other companies must be 'separated out'. Examples include Professional Services and Training.

The “Cost per Call” Method (a.k.a. Cost per Incident)

Definition: This method requires collection of all applicable costs associated with delivering customer support and then dividing that cost by the number of calls or incidents handled. The result is an average ‘cost per call’ or ‘cost per incident’.

This is undoubtedly the most common way for companies to measure and report support costs. It is also one of the most popular methods used by industry studies/reports. It’s a “chicken and egg” situation. Is the method popular because so many industry studies report this way or do industry studies report this way because of the popularity of the method? Either way, it has become the most common method in use…despite its flaws.

A relatively new and valuable variation is ‘cost per solution provided’ method, in which the costs associated with the different means of providing support solutions, such as telephone, email, and web self-service, are separated. Because this method often proves that certain methods (i.e. web self-service) is more cost efficient, it has become one of the primary selling tools for vendors of those tools. It is also an extremely valuable management tool for justifying additional investment in those areas that are proven to be more efficient (cost in relation to effectiveness).

The first and biggest challenge in utilizing this method is determining what expenses to include and, consequently, which not to include. You should always include all of those items that are carried in the Support budget and are directly related to delivering customer support. This includes all of the staff/personnel expenses, such as salaries/payroll, payroll taxes, benefits, and bonuses. It should also include all costs that are directly attributable to the personnel, such as training and other expenses. These expenses are usually calculated on a ‘per head’ basis or as a percentage of the payroll. Support expenses also include related expenses such as phones, mail, supplies, etc.. But this is where it starts to get complex.

In some companies, items such as phone equipment, phone lines/usage, and other necessary expenses are purchased by the Support Department and the expense can be easily and accurately captured. In other companies, these expenses are not separated, but are instead part of a corporate overhead. It may or may not be easy to allocate the appropriate portion back to Support. We find the same challenge when companies allocate General and Administrative (G&A) Expenses.

The second challenge is finding ‘good’ comparisons for benchmarking purposes. Aside from the challenge of ensuring that the companies you compare to include the same costs, we have a second challenge. There are a number of variables that come into play and a huge range in the level of effort required to provide customer support. For example:

- Percent of support provided to Customers: Some companies provide 100% of required support directly to their customers. Others expect the channel to provide customer support and the company’s responsibility is to support the channel. Some do both and there are numerous variations along the scale.
- Complexity of the product: There are many factors that affect this and the range is huge. Some companies can hire support reps with nothing more than a high-school education and they solve 90% of the calls in less than 10 minutes. Other companies need to hire PhD’s and they can only resolve 50% of the calls in less than one hour and the remaining 50% require hours of research and follow up work. Many of the industry studies have made attempts to identify the factors that impact complexity and then sort the data by those ‘categories’

Examples of Categories that may affect complexity (and therefore support costs)

Category	Logic for Sorting	Sample "Sorts"
Platform/OS	The platform may indicate the complexity of the product (i.e. Enterprise software requires a different level of support than desktop applications)	PC/Desktop, Client/Server, Mini/Mainframe
Product Category	The product type may indicate the complexity of the product	Vertical Market Applications, Business Products (i.e. Spreadsheets & Word Processors), Home/Entertainment Products, Network/Enterprise Software, Utility/Communication Products
Cost of Product	There may be a relationship between the cost of the product and the complexity of the product	Under $100, $101-500, $501 1,000, $1001-10,000, Over $10,000
Size of Company	There may be economies of scale and so companies want to compare themselves to similar-sized businesses	Can be by number of employees and/or sales revenue

Measuring cost per incident and isolating the cost of handling phone calls versus emails, etc. is a powerful management tool. It is very useful as a means of measuring improvements or deterioration of overall efficiency (cost per incident). It is also extremely valuable as a means of identifying which methods of providing support are the most cost effective, presumably with a goal of shifting customers toward that method.

The method has limited value for benchmarking purposes. There are just too many variables in the costs included for measurement and, more importantly, the range of the support complexities and requirements from one company to another to make this a valid means of benchmarking.

The "Cost per Minute" Method

Definition: This method requires collection of all applicable costs associated with delivering customer support and then dividing that cost by the number of minutes of support provided. The result is an average 'cost per minute of support provided'.

This method has both pros and cons when compared to the previous method (cost per call). We still have the issue of 'which costs are included and which aren't'. Also, there is still a fairly wide range of infrastructure requirements dependent upon the complexity of the product and the requirements of the particular market. More important is the question of "which minutes do you include in the calculation"? Do you include just the phone time or do you include the follow up (research) time?

The goal, and one of the positive aspects of this approach, is that it eliminates the issue of one company/product requires twice as much effort as another company. Utilizing this approach, we are theoretically just comparing the 'cost per operating minute'. However, the big flaw with this method is that it does not allow you to see changes in efficiency (i.e. increase or decrease in average handle time).

Example: *Let's take the example of a company that is getting less and less effective in handling customer calls. Over a period of time, the handle time per call could increase dramatically, even double for instance. And yet, as long as salaries remained the same and staff utilization remained the same, and cost to train new hires remained the same…then the 'cost per minute' would remain constant, even though headcount had increased at a rate greater than call volume (minutes and therefore cost per incident have increased).*

Extend the example and imagine the scenario where an organization had suffered high turnover and/or call volume was increasing. As a result, the organization hired a large number of new employees. These new employees, while costing much less than the senior staff, are not nearly as efficient and it will take them quite a while to become fully proficient (their handle time per call/incident is greater). However, the cost per minute would potentially seem to improve…lower cost of staff plus longer time per call probably equals lower cost per minute. While the efficiency of the organization is degrading, the metric may actually say things are improving (cost per minute is decreasing).

The method has limited value for benchmarking purposes and, arguably, limited value as a management metric.

The "Cost in Relation to Company Revenue" Method

Definition: This method requires collection of all applicable costs associated with delivering customer support and then dividing that cost by the total company revenue. The result is a percentage.

This method has some benefits versus the previously described approaches, but suffers some of the same problems. On the positive side, the measure is in relation to revenue so you are truly measuring the impact to the company. For instance, support costs may rise, but the relationship or percentage of revenue may remain the same (or go down) because of rising revenue.

This metric is relatively common in industry benchmark studies. The challenge continues to be 'are we comparing apples to apples'? The answer is, "probably not." And so, utilizing this metric and the data provided by industry reports is of minimal value. If your expenses vary dramatically from companies that appear to be comparable, then further research is certainly indicated. However, using industry benchmark data to set goals and/or as a primary measure of performance is unwise.

This metric, when tracked consistently and compared to previous periods, is extremely valuable. It will provide a gauge of company performance (not just support) and will indicate when further investigation is warranted.

“Cost in Relation to Support Revenue” Method

Definition: This method requires collection of all applicable costs associated with delivering customer support and then dividing that cost by the total service/support revenue. The result is a percentage, which can also be translated into ‘margin’ (the percentage of profit on support revenue).

Note: This method assumes that your support operation generates revenue (charges separately for support) as opposed to providing this as a ‘free’ service to customers (bundling with the product).

This method has some benefits versus the previously described approaches, but suffers some of the same problems. Obviously, we have the same issue of determining which costs to include (but that’s an issue with every method). On the positive side, the measure is in relation to revenue so you are truly measuring the impact to the company. For instance, support costs may rise, but the relationship or percentage of revenue may remain the same (or go down) because of rising revenue.

Allocating Maintenance Revenue: A challenge that many companies have is determining the amount of service revenue to allocate to the service department. The issue stems from the fact that the Service/Maintenance Contract includes product updates in addition the telephone support and other support features. This is not an issue for those rare companies that generate product updates or maintenance releases from within the support department. However, the vast majority of companies rely upon Product Development to produce these maintenance releases. This inevitably results in a battle between the two departments over how to allocate the revenue dollars.

I’ve seen three approaches commonly used. The simple and most common approach is an arbitrary allocation based on management decision (negotiation). The second, less common, approach is to base the allocation of the level of effort or costs of the two departments’ efforts. This takes a fair amount of effort in itself and so is not often used. Lastly, is basing the allocation on the customers’ perceived value of the different features in the support package. This is the best and most appropriate method. However, it takes a substantial effort and is therefore very rarely used.

Warranty Costs: An issue that frequently comes up when using this method is the cost of ‘warranty’. Depending on your warranty policy and your particular product requirements, this cost can be substantial. It is common that customers need more help during the first few months of product use than at any other time. It is also common that this period of heavy use is covered by the product warranty or ‘free support’ period. As a result, there are substantial support costs and no offsetting support revenues.

The typical view is that the warranty is part of the product…part of what a customer is paying for when they buy your product. That would lead to two potential approaches for dealing with the cost of providing that warranty. One approach is to collect the costs associated with providing warranty and then ‘charging’ those costs as ‘cost of goods sold’ (COGS) the same as manufacturing and other direct costs. This approach has been adopted from the hardware industry where the practice is commonplace.

The second approach is to allocate a portion of the product revenue to the support department. This essentially has Product Management (or the appropriate department at your company) ‘paying’ the support department to perform this function.

Either way, the charging of costs or the paying of revenue typically leads to scrutinization of costs. This often leads to discussions about 'alternatives', including outsourcing all or a portion of support. This obviously puts support under pressure to perform comparable to the alternatives and is likely a healthy situation for the company overall.

Conclusion

There is no single 'right way' to measure support expense. Each method has its own unique weakness and potential for misleading management. Relying on any one method alone could cause management to miss important facts or trends. Therefore, a combination of at least two methods is required.

I recommend using the Cost per Incident Method (broken down by the various support delivery methods) combined with either the Cost in Relation to Revenue or Cost in Relation to Support Revenue. Together, these metrics will give management a relatively complete picture.

It is important that management thinks through which cost are to be included and how they will be calculated. The same goes for revenue allocation or support 'charge back'. Equally important is consistency. The majority of value to be gained as a result of measuring and reporting these numbers is to recognize trends and opportunities and then make the appropriate adjustments. Therefore, consistency is the key. It is not so important as to which costs are included as it is important that we measure the same way every time.

Attempting to benchmark your costs versus other companies is of minimal value. The benchmarking data that is available in industry reports is too general and the sources too questionable to make serious comparisons. The cost of performing a customized research project is too great in relation to the value received (for most companies). For most companies, you can use the available studies to validate that you are 'in the range'.

You can and should study your own data and analyze the cause of any trends or fluctuations. If costs appear to be out of line, or are moving in the wrong direction, further investigation or an outside review may help isolate the cause and define corrective action.

Note: Another variation is the "Cost per Customer" method. This method has not been addressed by this paper, but will be discussed in a future revision.

About the author: Dave Brown is an industry consultant, teacher, and author and is considered an expert in the area of process improvement and change management. Dave teaches management-training programs for Support Center University (www.SupportCenterU.com) and consults with selected clients to establish world class service operations. Reach Dave at his office in Boulder, Colorado (303-494-4932) or by email dave.brown@SupportCenterU.com.