

Setting the Standard:

A Guide for Determining Optimal Service Goals for Inbound Sales Centers

By Ric Kosiba, Ph. D., Bay Bridge Decision Technologies (www.baybridgetech.com)

“Tradition!” Despite most companies’ stated focus on profitability, the single most frequent response we get when we ask call center professionals about how they set service goals is that: It’s a tradition! Other top responses include “our industry uses 80/20” and “the boss said so.”

Setting goals for service level, ASA, abandons, and other quality measures at optimal levels is critical for inbound sales centers. Goals drive staffing costs, agent schedules, center sizing, technology purchases, call routing rules, performance reviews, executive bonuses... and, most importantly, revenue and profit. What determines the “best” service goals for inbound sales centers? In our view, *service goals should be set at levels that maximize profit.*

However, our experience is that service goals are most often set with little reference to profit, the single most important metric to the company. It does not have to be this way. This article outlines an eight step process using discrete-event simulation that you can use to make routine the important task of setting profit maximizing service standards.

Why is Profit-Based Analysis So Rare?

We find that many managers would like to take a more profit oriented approach to setting service goals, but they lack the tools to do so. The modeling required is complex and, unfortunately, Erlang C based methods alone are insufficient.

Setting optimal service goals requires a robust model of your call center(s). This model must accurately predict your revenue and service quality performance based on caller behavior, handle times, staffing plans, and cost drivers for each month or week of your planning horizon. The model must be able to tell you the benefit, in marginal revenue and profit, of hiring one more phone agent (i.e., if I add one more agent, what is the increase in expected sales volume, and how does that compare to the agent’s cost?).

Unfortunately, our old standby call

center model, Erlang C, will not help us here. To illustrate, think about the impact of revenue per call on target average speed of answer. Clearly, you will want a faster ASA if your revenue per call is \$1,000 per call than if it is \$1. Each call that is abandoned is much more revenue lost. So, to set an optimal ASA goal, you must have a center model that predicts how abandoned rates will vary with staffing, so you can determine what abandon rate maximizes profit. Since one of Erlang C’s core assumptions is that no callers abandon, obviously Erlang C is not much help in solving this type of problem.

An Approach to Determining Optimal Service Goals

Remember that our objective is to determine profit maximizing service goals (e.g., service level, abandon rate, blocked rate, average speed of answer, or whatever metrics are important to your call center). In order to determine the costs of changing service goals, you must decide what lever you will use. The most straightforward method would be to alter your staffing to achieve different service goals, but you could also, for example, implement training or purchase new technologies that increase customer patience or shorten handle times. For this discussion, we will assume that you’ve decided that you will manage your service goals by altering the number of employees on the phone.

In other words, you want to answer the question: “For a given call volume and call value, how does net profit vary as I change my staffing levels, and what service standards result from each of those staffing levels?” In practice this means that we want to find the staffing level that maximizes profit, and then observe the resulting service measures at that profit-maximizing staffing level.

Of course, profit is not always the only consideration – for example, your company may have non-negotiable service targets (“always answer in three rings”), contractual or regulatory considerations, or particular competitors you must match. But profit should be your

primary measure, and policies should be evaluated in terms of their cost in lost profit vs. their benefit in other areas. In practice, we frequently find that sales centers are in fact *understaffed*, and could benefit from more available agents. Their managers, however, are unable to make a good business case for adding staff because they lack the supporting analysis.

Simulation Modeling

The modeling tool we recommend is discrete-event simulation. A simulation model is a detailed representation in computer software of how your center actually operates. The model is created using information regarding call arrival, customer patience, handle times, staffing levels, revenue, and cost derived from historical experience. The simulation is then “run” by virtually generating every call that arrives at your center(s) over the forecast period, processing each one according to your business rules, and tabulating the resulting cost, service quality, and profit performance. The model can be used to determine the profit performance of your center as staffing changes, and thus highlight for you the staffing level at which profit is maximized.

There are a number of excellent resources that can guide you in developing center simulation models and commercially available software that will make the modeling and analysis process easier. Also, keep in mind that once built, your simulation model can be used for many other types of “what-if” analysis in addition to those detailed here.

As an aside, another method that can be used for this type of center performance modeling is regression modeling. A regression model is a set of predictive equations (e.g., ASA as a function of staffing level and caller volume) that is developed using historical call center data. Since regression predicts the future by extrapolating from the past, however, it can be inaccurate when used to model scenarios outside of the call center’s experience. Because simulation modeling is more accurate and flexible than


Continued on back

regression, the following work steps assume you are using a simulation model.

Here are the key simulation steps:

- 1. Forecast Call Volumes and Arrival Patterns.** You probably already have a call volume forecast, and historical call arrival distributions that will form the basis of your arrival forecasts (by hour or half hour interval of each week.) These are generally available from your ACD database or workforce management system.
- 2. Determine the Customer Patience Distributions.** Key to your model is estimating of how long your customers are willing to wait on hold before abandoning. Figure 1 shows a cumulative “customer patience” curve, where each point on the curve represents the time on average a given percentage of your callers are willing to wait before abandoning. This curve can generally be derived from ACD data.

Figure 1: Cumulative Probability of Abandon by Time



- 3. Determine your Handle Time Distribution.** Next, you will need to understand your distribution of call handle times around their average. This data is easiest to derive from call level detail, but can be approximated using interval data if needed. Your simulation model will need this information to assign a predicted handle time to each call as it simulates its processing by your call center.
- 4. Determine your Staffing Patterns.** Your model will need a method for forecasting how your workforce management system will schedule the available staff pool as you grow or shrink your staffing to reach a profit-maximizing level. A good approach is that you will add or subtract staff, but maintain their average distribution during the week (i.e., you will staff the same percentage of the weekly total during an interval, say 11:00 to 11:30 am on Tuesdays, as the number of staff rises or falls). Use your

ACD or workforce management reports to derive a percentage staffing distribution (similar to a call arrival distribution by interval), which is simply the number of agents available for that interval as a percentage of agents available for the week. This approach works well because, like call distributions, staff distributions tend to be quite consistent, even though the absolute level of calls or staff may change significantly.

- 5. Determine the Value Per Call.** The next step is to determine the net contribution value per call (i.e., marginal value to the company of a call answered, prior to subtracting variable call center costs) and how this value may change seasonally (e.g., larger than average order sizes in October and November). Take into consideration factors like the value of products sold, cross sell opportunities, retention opportunities, and the probability of a sale. For an inbound call, what is the average value to the company of answering it?
- 6. Determine your Servicing Costs.** Since our concern is marginal value per call, we must consider marginal cost per call as well. Identify which of your costs truly vary with staff (e.g., recruiting, wages, taxes, benefits, supervisory, workstations, but excluding items such as corporate allocations and overhead) and per call (e.g., network costs). If your analysis determines that major changes are needed in service quality targets, you may also need to consider the second order effect of these large changes on your “fixed” costs such as facilities, management, and technology.
- 7. Develop your Simulation Model.** With your assumptions about customer behavior, staffing, and costs, you can now build the simulation model of your center. The details of this process will depend on the specific software that you use. A critical step in this modeling process is to validate your simulation models by comparing actual center performance to predicted performance and ensuring the model results match to reality.
- 8. Run the Model and Interpret the Results.** Once you are comfortable that your model is valid, you can use it to iteratively develop your optimal staffing levels. To do this, vary your staffing while holding all other vari-

ables constant, and view the resulting profit performance.

For example, if you currently believe you need 100 agents to handle your forecast call volume (once you’ve accounted for shrinkage) you can run simulations to determine the profit with 75 agents, then 90 agents, and so on, to, say, 125 agents. Then, plot staffing vs. profit to determine the staffing level associated with the highest profit. This level is your optimal staffing for that call volume.

The forecast service quality measures (e.g., ASA, abandon rate, service level, blocked rate) associated with this call volume and staffing combination are the profit maximizing levels for these measures.

Another way of looking at this is to directly graph profit against a service measure, for example, service level (e.g., percentage of calls answered in 20 seconds). Each service level corresponds to a different staff level for a particular call volume. The high point on the curve is your profit maximizing service level. The staffing level required to achieve this service level is the profit maximizing staffing level.

It is a good idea to graph these relationships for all service measures important to you, as well as for typical ranges of call volume and call value to get a good feel for the key drivers of profit and how they vary.

Summary

Developing a discrete-event simulation model of your call center and conducting this service goal optimization analysis can take a good analyst a few weeks, but the results are well worth the effort. Discrete-event simulation is a powerful tool, both for determining your optimal service levels, and for answering many other difficult “what-if” questions.

Using simulation to determine your profit maximizing staffing levels and service quality targets will be revealing. Right away you will see if your current quality goals, by leading you to either overstaff or understaff, are costing you profit. Using this information, you can make your business case that your staffing levels and quality goals are too high, too low, or “just right.”

Ric is a charter member of SWPP. You can reach him at EDK@baybridgetech.com.