

Workforce Management Series - sequel 1:

Forecasting Fundamentals: The Art and Science of Predicting Call Center Workload

The basis of any good staffing plan is an accurate workload forecast. Without a precise forecast of the work to be expected, the most sophisticated effort to calculate staff numbers and create intricate schedule plans is wasted effort. The old adage of “garbage in, garbage out” is especially true when applied to call center workforce management. An accurate forecast is the most important step of the process.

The purpose of the forecast is to predict workload so that we can get the right number of staff in place to handle it. And there are many different situations in the call center environment that require a forecast to be done. The most common scenario for which we forecast is simply normal, day-to-day operations. But you may also require a forecast for special situations such as planning for new call type(s), opening a new center, a merger or acquisition, or a change in operating hours. Or you may be implementing a new technology that will affect your call volume or pattern and need to determine what the resulting change means to staff workload. Whatever the reason, it’s important to understand the basic principles behind workload forecasting and how to apply them to accurately plan call center resources.

The forecasting process is both an art and a science. It’s an art because we are, after all, predicting the future. And the accuracy of your forecast will be due in some part to your judgment and experience. But it’s also a science - a step-by-step mathematical process that takes past history and uses it to predict future events. A working knowledge of these specialized statistical techniques, along with a pencil, paper, and calculator will get you through the process. And for those of you that have workforce management software in place that automates the forecasting process, don’t think that you’re off the hook! It’s just as critical for you to understand these calculations as it is for someone that’s doing them by hand. It’s important you understand the numbers coming from the software tool to verify accuracy of results and perhaps more importantly, explain the numbers to management. So even if you have tools to help, learning the fundamentals of forecasting is worthwhile.

Step 1: Gathering the Data

The first step in the forecasting process is gathering representative historical data. We assume that past history is the best predictor of the future in most call centers, so gathering this history is the first task. The most obvious source of this information will be historical reports from the ACD -- specifically the number of calls offered and handle time information by half hour.

If you’re wondering about how far back to delve into your historical reports, we like to have two years worth of past history if it’s available and if it’s relevant. Less than two years worth may suffice, but won’t give you the most accurate tracking of trends and monthly/seasonal patterns that 24 months will clearly show.

It’s important to note that we typically assume the NCO (number of calls offered) accurately portrays the workload for which we need to staff. This assumption is valid

as long as “all calls are getting in” and that none are blocked at the network level by insufficient telephone trunks. It’s always a good idea to validate this assumption by requesting periodic “busy studies” from your local and long distance carriers.

Another critical step of the data gathering process is to eyeball your information to make sure there are no data aberrations. You’ll want to look for any abnormally low or high numbers as well as missing information. When you identify something out of the ordinary, you should first determine the reason for the anomaly, and then decide if it needs to be adjusted or not. As an example, let’s look at a previous year’s daily call volumes for July.

S	M	T	W	T	F	S
			5281	4212	3610	0
209	5200	5531	5407	5488	5420	1110
910	5892	5587	3785	5512	5536	1212
951	5932	5590	5467	5541	5598	1234
933	6031	5655	5512	5593	5699	

You’ll see several aberrations in the historical information. One is related to the 4th of July holiday weekend. Call volumes begin to drop on Thursday, are significantly lower on Friday, are zero on the actual holiday and following Sunday, as well as the Monday that follows. What should you do about the aberrations?

Since the reason for the anomaly is a holiday that will repeat, we’ll want to account for the holiday as we predict what volumes we’ll receive next July. However, the actual day of week of July 4th changes from year to year, so the pattern will not be exactly the same. If the 4th shifts to a Monday, we might expect the Tuesday following the holiday to be much lower while the Thursday and Friday prior might not be significantly affected. This is where the “art” comes in – using your intuition and judgment as part of the forecasting process.

The other aberration happens on the third Wednesday of the month. You’ll see that call volumes are 30% lower than the previous Wednesday. There could be several explanations for this discrepancy. It might just mean that the ACD didn’t record calls that hour due to a power outage. Or perhaps there was a compelling news event that afternoon and call volume dropped significantly. In either event, you’d want to “normalize” the data back to a realistic number before including the data in your forecasting calculations.

On the other hand, there might be an event that happens the third Wednesday of each month that really does cause call volume to drop. Assume this data represented the calls to an internal help desk, and that on the third Wednesday of every month, there was a two-hour company-wide meeting. In that case the numbers on the report accurately reflect the volume that day and would also be an accurate number to use to forecast future numbers.

The key in dealing with a data aberration is to first determine the reason it occurred. Then, if it’s a one-time incident, or an event that might occur again but you can’t predict when (like a storm), you’ll want to normalize the numbers up or down to reflect realistic volumes. On the other hand, if it’s a repeatable, predictable event, these numbers need to stay in the data so that the forecast reflects the event in the future. (Hint: It’s important to note in the data why each aberration occurred so you’ll remember it for future planning purposes!)

Once you've analyzed and adjusted the historical information, then we're ready for the next step...

Step 2: Predicting Monthly Calls

The next step in the process takes us from raw data to a prediction of what's coming for a future month. There are several approaches to get us to this future forecast:

Point Estimate. This is the simplest approach and assumes that any point in the future will match the corresponding point in the past. (i.e., the first Monday in April next year will be the same as the first Monday in August of this year). This approach has obvious shortcomings in that it does not account for any upward or downward trends in calling patterns. It's also dangerous in that the forecast can be dramatically different if the original data was atypical.

Averaging Approaches. There are a variety of methods that incorporate simple mathematical averaging, ranging from a simple average of several past numbers, to a moving average where older data is dropped out when new numbers are available. The most accurate averaging approach involves weighted averaging, where more recent events are given more weight or significance than older events. So if the call volumes on the first Monday of April for the past three years have been 2400, 2500, and 2600 calls: the simple average would be 2500 calls, the moving average might be 2550 calls (dropping out the oldest data). In a weighted average approach we might assign an 80% weight to the most recent number, with only a 10% weight assigned to each of the prior years giving us a prediction of 2570. But while the weighted average approach is probably the closest to what an actual forecast would be, it still misses the upward trend in the data that simply can't be identified and incorporated by averaging together old numbers.

Time Series. The recommended approach for call center forecasting involves a process called time series analysis. This approach takes historical information and allows the isolation of the effects of trend (the rate of change) as well as seasonal or monthly differences. It is the approach used in most call centers and serves as the basis for most of the automated workforce management forecasting models. The basic assumption is that call volume is influenced by a variety of factors over time and that each of the factors can be isolated and used to predict the future.

The first step in a time series approach is to isolate the effect of trend. Trend is basically just the rate of change in the calls. While that trend can be upward or downward, in most call centers, trend simply means the growth rate. It is important to determine this rate as an annual trend rate as well as a month-to-month change.

Once the trend rate has been determined, the next factor to isolate is the effect of seasonality or month-to-month variances. This process is fairly tricky, since you can't really determine monthly or seasonal factors just by looking at the most recent twelve months of data. In looking at the first column of monthly call volumes below, is December really a "busy" month compared to May, or is December's volume higher because we've been experiencing a large upward trend and has just simply had seven more months to grow?

	Monthly Volume	Detrended Volume	Seasonal Pattern
January	9,350	13,944	1.048
February	10,450	15,028	1.129

March	11,560	16,031	1.205
April	11,140	14,898	1.119
May	10,000	12,896	.969
June	8,490	10,558	.794
July	9,680	11,608	.873
August	10,540	12,189	.916
September	12,880	14,363	1.080
October	12,670	13,625	1.024
November	13,170	13,657	1.027
December	10,850	10,850	.816

To determine the effects of seasonality, it's important to "detrrend" the most recent twelve months of data – in other words, bring each month up to current levels by factoring in the month-by-month trend rate. After detrending, we can do an "apples to apples" comparison. The months of the year can be compared against one another to determine what are actually busier than average or slower than average months. In the example above, we see that May is actually "busier" than December based on calling patterns, with March and April actually being our peak times of year.

The trend rates and seasonal patterns identified using time series analysis are then used to pinpoint specific future monthly forecasts. The time series process is the recommended approach to forecasting future workload and if done precisely, can generally create forecasts with 95% or higher forecasting accuracy.

(Note: The process of time series analysis including trend isolation, detrending analysis, and seasonal pattern identification is a fairly complicated one and the step-by-step process is beyond the scope of this article. For more information on the steps, contact The Call Center School at 615-812-8400)

Step 3: Creating Daily and Half-Hourly Forecasts

Once monthly forecasts are in place, the next step involves breaking down the monthly forecast into a daily prediction, then further down into an hourly or half-hourly numbers. To predict daily workload, you must first calculate day-of-week factors. Most call centers have a busier day on Monday than other days of week and it's important to know what percentage of the week's workload this day and others represent.

The good news is that it's not necessary to go back and analyze two years worth of information to determine these factors. Typically evaluating the last few weeks worth of daily call volume data is sufficient to identify daily patterns. Just select several "clear" weeks of data (those without holidays or other major events that might skew the proportions) and see what the total Monday volume is compared to the weekly total. Then repeat for the other days of week. These percentages reflect your day-of-week patterns.

Once the daily forecast is in place, it's time to repeat the process for time-of-day patterns. It would be nice and easy to schedule staff if the calls came in evenly

throughout the day, but since that's not reality, it's critical to know when the peaks, valleys, and average times are. Again, gather several "clear" weeks of data and evaluate the Mondays to look at how each half-hour of the day compares to the daily total to create your Monday half-hourly patterns. Then repeat for the other days of the week. The result will be 24 hourly or 48 half-hourly percentages that represent intra-day call patterns and you'll have one for each day of week.

We've now broken down our historical data and past trends to develop a monthly, then daily, then half-hourly forecast of workload. Keep in mind that this forecast must include not only call *volume* predictions, but should include a prediction about *handle time* as well. To calculate workload and predict staffing and schedule requirements later, we want the total picture of workload, which is number of calls multiplied by average handle time. Make sure your handle time predictions accurately reflect the time of year, day of week, and time of day since call length may vary for a number of reasons having to do with business variations as well as caller behavior.

Step 4: Adjusting for Other Business Influences

The final step in the forecasting process is an important one. There are many factors that influence the call center's workload and the smart workforce planner will have a process in place that considers all these factors in the forecasting process.

Think about all the different areas of your organization that influence the calls you receive. The most obvious one is the marketing department who has tremendous impact on your work based on the sales and marketing promotions they do. Hopefully you have a formal communications process in place to hear about marketing plans well ahead of the actual event so they can be built into the forecasting assumptions.

Make sure you consider all the other pertinent areas as well. Will the billing department's new invoice format cause a flood of calls? How about sales forecasts from the Sales VP that can help you plan staff based on the new customer account base a year from now? Is the fulfillment area changing the way they package and ship products that may cause an increase (or decrease!) in your call volume? It's critical that you communicate regularly with all these influencers of call center workload as you prepare and fine-tune the forecast.

Once the forecast is in place, then you're ready for the next step – calculating staff requirements to meet service goals. Stay tuned for the detailed steps in our next article: *The Math of Contact Center Staffing: How to Calculate Staff Numbers for Incoming Calls and Multi-Media Contacts*.

About the Author....

Penny Reynolds is a Founding Partner of **The Call Center School**, a company that provides a wide range of educational offerings for call center professionals. Penny is a popular industry speaker and is the author of numerous call center management books, including *Call Center Staffing: The Complete, Practical Guide to Workforce Management* and *Call Center Supervision: The Complete Guide for Managing Frontline Staff*. She can be reached at 615-812-8410 or by email at: penny.reynolds@thecallcenterschool.com.