Strategic Shrinkage Management

By Ric Kosiba, President, Bay Bridge Decision Technologies

The Importance of an Accurate Shrinkage Forecast

Anybody who has ever suffered through one of my SWPP Annual Conference presentations has probably seen the following two graphs:

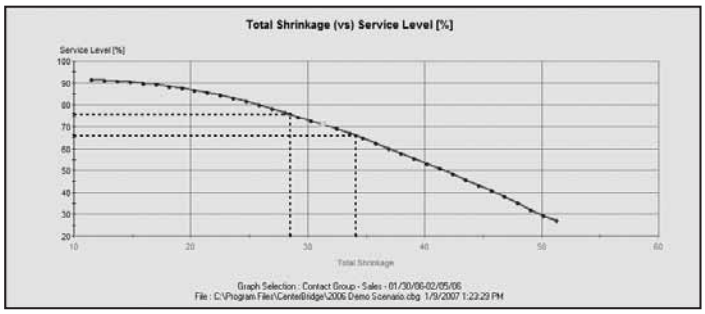

These two graphs – validated simulations of the same contact center and the same time period — show something very interesting. In the first graph, we determine what range of service we can expect if contact volumes were to vary by 5%. In the second graph we vary employee shrinkage by plus or minus 3 raw percentage points (between 28% and 34%) and get the very same service variance from shrinkage as we do from contact volume variance. This is very important.

In plain English, this strongly implies that your shrinkage prediction is as important to your plan as is your contact volume forecast.

Too many of us "flatline" our shrinkage forecast, or cut and paste last year's shrinkage percentage into our capacity plan.

How many of us measure shrinkage forecast error?

Hugging the Staffing Requirements Curve

I think most of us have been trained to think that our goal, when determining staff plans, is to get as close to a requirements curve as possible. I believe this too — but with caveats — as there are definite exceptions to this approach.

A staffing requirements curve assumes an awful lot. First, it assumes that there is a consistent service goal. While this assumption is common in our industry, there are obvious exceptions. If your contact center has significant seasonality of contact volumes, the cost versus service trade-off may be very different at different points in time due to the differences in your economies of scale. In these cases, it makes sense to have different service goals by time of year.

Similarly, if your contact center answers revenue producing contacts, and there are seasons where the sales margins per contact varies significantly (e.g., retailers at Christmas), then there may be seasonal differences in the profit versus service trade-off. In these cases (and others), a contact center analyst could find that it is optimal to change service standards by season. We've even seen one forward-looking reservations company that operates their workforce management process with different service goals by hour of day, because they have very different marginal revenue projections by hour. That is very good, forward-looking stuff.

Another standard assumption is that weekly staffing requirements can be accurately developed using Erlang or workload calculations. I think we've discussed the shortcomings of Erlang for strategic planning in past articles, and suffice it to say I do not believe Erlang is accurate enough for long term planning. For those folks who use a workload estimate, there is a flaw in the methodology that structurally puts plans at serious risk. It is this: workload calculations assume abandons and occupancy up front in order to determine your week-by-week staff plan. However, it is your staffing decisions that really determine the abandons and occupancy that your operation will achieve. Workload calculations assume the answer. Assume perfectly and your staffing requirement is OK. Assume incorrectly and it can be inaccurate and costly.

Obviously, if you've read any previous discussion articles, you know that I am a strong proponent of using discrete-event simulation modeling to determine requirements.

But assuming that we can develop an accurate weekly staffing requirement estimate, what's next?

The Standard Staff Planning Process

The typical process goes something like this: Volume and handle time forecasts are combined with a service goal and through a contact center model (e.g., simulation or Erlang) weekly staffing requirement forecasts are produced. These are most often combined with a schedule inflexibility factor (we goose the requirements up) to approximate how well (really how inefficiently) our agents are scheduled.

From there, we make our best guesses as to what shrinkage will occur and what attrition we can expect by week. Curiously, our industry does not typically utilize the same resources to forecast shrinkage or attrition as we do to forecast contact volumes.

From all of these forecasts and with a current headcount tally (often, more difficult to derive than laymen would expect), analysts can determine, week by week, whether there are enough agents on staff to answer the contacts within the allotted service standard (attrition is factored against this headcount tally). This staff assessment takes the form of a simple over/under spreadsheet.

From there, it's only the simple matter of determining when and where to hire your agents, when to offer overtime or undertime, when to schedule training, etc… Easy.

When and where to hire — in almost every spreadsheet we've seen — is a manual "guess and test" process, where hiring classes are placed in one of the 78 weeks or so (assuming an 18

Continued on page 7

Continued from page 6

month rolling plan) and the new over/under chart is generated in our spreadsheet. We then try adding another hiring class in another location and recalculate the over/under chart. We repeat this process until the "unders" are pretty much gone. We may, if we're particularly good, reduce hires through use of overtime. That is our hiring planning process.

If you are my age and have been around workforce management as long as I have been, this process might seem eerily familiar — it is exactly how we used to put together agent schedules before workforce management software became popular. If you recall, when we first went through the process of automating and optimizing agent schedules we saw a ridiculously fantastic return on our workforce management purchase. The reason is that computers and optimization algorithms do a much better job of determining solutions to difficult combinatoric problems than us mortal analysts could ever do. Determining where and when to hire and when and where to offer overtime and undertime in a multi-site, multi-skill, or multi-channel contact center network is exactly the same sort of tough problem that scheduling is.

By improving this strategic planning process, we can see ridiculously fantastic returns once again.

Optimally Planning and Actively Managing Strategic Shrinkage

There is, of course, much more to "hugging the requirements line" than hiring planning (even though that in itself is a difficult problem). Our degrees of freedom actually encompass much more. We have at our disposal the ability to plan overtime and undertime, to plan certain shrinkage categories, like training and possibly vacation, and to influence other shrinkage categories, like sick time, through our leave policy. By planning these well, we can move closer to a truly optimal plan.

By optimally planning and actively managing shrinkage, as though it is a line item in our budget (surprisingly, for many companies it is not), we can significantly reduce our costs, while producing both optimal and consistent service to our customers.

How do we do this? There are two types of shrinkage that shape how well we manage to our staff needs.

Forecasting Shrinkage

Some classes of shrinkage are beyond our direct control. Items like sick time, FMLA, and attrition (not always considered a shrinkage item), can be predicted and influenced, but cannot be directly managed. For these types of shrinkage, it is important that we apply the appropriate forecasting technologies, so we have an accurate prediction of these items.

We've looked at a lot of shrinkage data, and there are some universal truths associated with forecasting shrinkage:

1. Shrinkage is predictable: While many of us take shortcuts, like assuming an overall shrinkage number, shrinkage is very predictable, if the appropriate forecasting methods are applied.

2. Seasonality is important: Each shrinkage item has its own seasonality, and we should determine and account for this seasonal shrinkage in our plans.

3. Each center and staff group likely has a unique seasonal pattern: Different centers have different behaviors, and it will add to your plan's overall efficiency if these differences are considered.

4. Missing a shrinkage forecast is bad: Being over or under your shrinkage forecast leads to inefficient plans, over or under servicing, and likely another forecast miss down the road when expected shrinkage time-shifts.

Not Forecasting Shrinkage (Planning For It)

There are other classes of shrinkage that are often directly under our control, items like vacation (if it is allocated to agents), coaching, team meetings, training, etc…

Again, our goal is to make the appropriate service cost trade-off; this trade-off likely varies significantly by time or season. For instance, during peak weeks, the cost of meeting our standard requirement is much dearer than during valley weeks. Hence, you want to use your levers, overtime, undertime, and planned shrinkage at a time when it makes most financial sense.

There are technologies that utilize true mathematical optimization that can help us determine exactly when and where to hire, offer overtime, offer undertime, offer vacation, and plan for training, meetings, and other off phone shrinkage activities. Technologies like linear or integer programming can help determine the optimal allocation of hiring, overtime, shrinkage and undertime over our planning horizon.

Note that we are not talking about plans that incorporate just management of daily plans, but instead managing the seasonality inherent in our industry. There are great advantages to developing a strategic shrinkage and resource plan. Some of these include:

1. Planning shrinkage in advance will get you more flexibility: By publishing a schedule of times when you expect to be overstaffed, you'll plan for training and management activities, and grant vacation, VTO and AWOP, in time to actually get it from your agents and your management.

2. Holding management accountable: By publishing an undertime plan, and holding management accountable for getting the appropriate resources out the door as part of their budget, you'll see active management of shrinkage and undertime. You'll get more of it, when you need it.

3. Dictating shrinkage appropriately will lead to a better result: If you have the flexibility with your workforce to dictate when and how much vacation, for instance, is allowed, you should take advantage of that and plan it well.

Utilizing the mathematical tools and the discipline that our industry has developed for workforce management toward our strategic plans will pay huge dividends in lower costs and more consistent service. Managing your shrinkage, through a combination of forecasting, simulation, and staffing optimization technologies (and some good old fashioned common sense), will make these benefits real.

Ric Kosiba is a charter member of SWPP and co-founder and President of Bay Bridge Decision Technologies. He can be reached at EDK@BayBridgeTech.com or (410) 224-9883.