



Setting the Budget for Maintenance Workload


By Richard G. Lamb, PE, CPA; Analytics4Strategy.com

As explained by the article, “The Secret is to Budget and Control Maintenance Opex Dimensionally,” activity-based budget and control for maintenance operations has become possible with the coming of data analytics. This has immense ramifications because budget and control is as fundamental to business operations as is accounting. And, like accounting, what qualifies as credible budget and control is very specific.

A core qualification of credibility is that a plant’s approach to budget and control must be built upon the drivers to total maintenance cost. The commonly observed cost-type-based maintenance budget and control process does not comply, even nearly, with the rule of drivers.

Activity-based budget and control is a build-up of drivers from workload (activity) to resources to unit costs. Therefore, the first stage of the workload-based budget procedure is to establish the number of jobs to be completed each month.

The historical obstacle to the stage’s objective is dramatically apparent in Figure 1. Good luck with drawing meaningful insight from and confidently making budget decisions upon the plotted 11 years of history. How many times, in Excel, have we drawn a straight trend line through such a plot as if it meant something?



The solution is to be had through data analytics. It is to decompose the plot into its constituent parts: level or trend and recurring cycles of season and calendar. Each is explored for insights leading to the final decision. In turn, through data analytics, the patterns of the component parts are reconstituted to forecast workload into the budget year. This article will explain the process as four steps.

A noteworthy feature of the explanation is that it is cast in the context of the [“R” software](#) for statistical computing and graphics. The point is that what is explained in this article is readily and immediately achievable. This is because R is an open source software which is free to anyone and any firm. Furthermore, “free” does not mean weak compared to the best of commercial offerings. Nor does the software give us a free “regular” version as bait to a for-fee “premium” version. R is not only extensively supported on the internet, but there are also free, high quality internet-based courses available to anyone.

Steps to Establish the Workload Budget

There are four steps to establish budget for monthly job count. They are to establish the workload groups, gather and prepare extended count history from the CMMS and associated systems, decompose and explore each groups’ time series, and set the budget for each group. Each heavily engages the methods of data analytics.

Step 1: Establish workload groups. The first step is to distinguish the workload groups upon which the budget will be built. At the least, they will be homogeneous as to work type and lead craft. Further yet, they will likely be variously grouped such as by asset type, etc. They will, of course, always be grouped by cost center. Regardless of the grouping on which the budget is built, they will roll up to match the structure of the plant’s monthly cost reports.

Although a topic for another article, a range of data analytic methods can be used to discover important, but invisible, groups according to strength of relationships between variables. The

possibilities are machining learning methods such as decision trees, clustering, text mining and other analytics.

For the discussion to come, we will speak generally as if the group of interest is reactive work for a specific asset group. Although the issues are generally the same for proactive work, reactive makes for a more interesting example. Major maintenance projects are budgeted based on project-type planning and control and, thus, are not established by the methods of this article.

Step 2: Gather and prepare extended count history. Because the massive amounts of data and variables captured in a plant's various systems are easy to tap into and work with, we should extract data as far back as there has been a CMMS to capture it. Change in workload may be too gradual to notice over just several years and patterns in the last several years may reflect decisions and events from long ago.

The core data of interest are the occurrences of work defined as a request of some type that is accepted and not later canceled. The procedure that identifies work meeting the definition varies from plant to plant. Therefore, the gathered data will include status history in order to identify which requests for work should be treated as occurring work.

We do not stop once we have the dataset of occurring work and status history. We will likely find it necessary to reach out for relevant data sets in the CMM, ERP, labor and other systems, and join them to the core dataset.

Once pulled together, the step will utilize analytics to look for outliers and missing data. Outliers will be investigated to discover which are cases of incorrect entry and which reveal something interesting. For missing data, analytics may be used to impute estimated replacements.

Analytics will also build new variables with the cleansed data. Some will be straight forward transformations from adding, subtracting, multiplying, dividing and categorizing variables. Advanced analytics may form new variables using methods such as binning, decision trees, clustering, principle components and text mining.

Much is learned in this step. The maintenance operation will gain insights it has never before had at its avail. It is especially noteworthy that some are textbook basic and essential to any maintenance operation, but have always been denied to us because they are only visible through data analytics. It is no surprise that we may return to the first step based on what is learned in this step.

Step 3: Decompose and explore the time series for each workload group. Now to extract meaning from the raw time series (Figure 1) for each work group. The article, "Explore What Did and May Happen With Time Series Question," explains the time series analytic tools for the task.

Please note that the data in this article are simulated rather than from an actual maintenance operation. The data has been simulated to show the range of perspectives that can be explored along the way to ultimately setting a monthly budget for workload.

The first action is to use time series analytics to reveal and remove the cyclical effect from the time series. In some cases, cyclical patterns may not be visually confirmable, but can distort our perspective if not removed. Thus, this action as it is executed allows us to confirm whether or not there is a cyclical influence.

In this case, it is obvious in that there is a seasonal pattern. However, we still need a sharper understanding of its shape. This is possible with the Cycle function and graphic capability of R as shown in Figure 2.

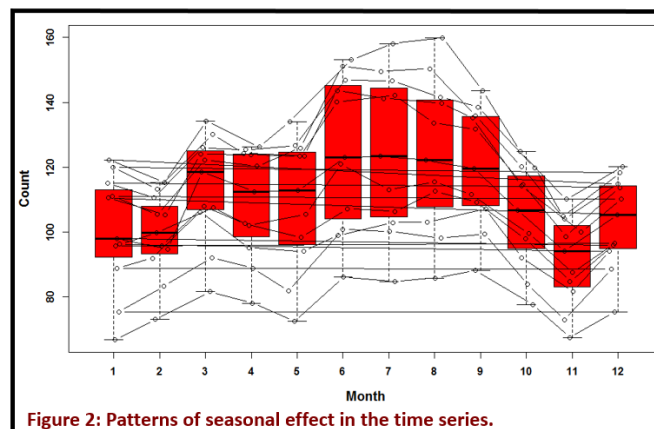


Figure 2: Patterns of seasonal effect in the time series.

Figure 2 is a many-dimensional, month-by-month perspective of the seasonal pattern along the 11 years of data. The vertical axis is the job count of occurring work and the horizontal axis is month.

The boxes and their vertical line extensions (whiskers) are boxplots. The heavy line across each box are the midpoints (median) of all cases. The boxes are the two middle quartiles (interquartile range, IQR) of the cases. Any outliers would be revealed as points falling beyond a whisker. The spread for each month is shown by the length of the IQR and whiskers.

The plots of point-lines show the cycle for each year. The horizontal cross over lines from the right-most to left-most points shows the sequence of years by connecting the end month of each year to the first month of the subsequent year. With them, we can see in which years, if any, there are breaks from the generalized year-to-year patterns. Notice that the points are “jittered” so that one is not blocked from view by another.

There is another issue for seasonal cycle for which we can get insight from Figure 3. It is whether or not the cycles are swinging wider with time (top and bottom panels). Just as for the previous visualization of annual cycles, the budget team will ponder the reason and message.




Figure 3: Decomposition of time series into level, trend and seasonal.

Figure 3 is the graphic output of the Holt Winters model. It is available from the R software, as well as, commercial software such as SAS. Its innards run iterations of decomposition until it finds the closest fit between what its algorithm predicted one period ahead and what actually occurred. The result is to decompose the plot of occurred work taken from the CMMS into the level and seasonal components.

The fit is further evaluated through the autocorrelation of its residuals (not to be explained in this article). If not, we would seek other better fitting models.

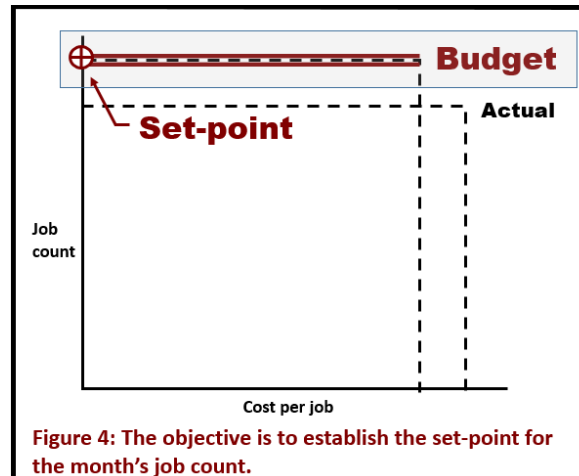
The algorithms extract season from the total time series to generate the **level** plot. Level shows the pattern of occurring work over the years undistorted by repetitive cycles.

The plot allows us to see the long-term pattern of workload. Is there a bathtub shape? Are the intended effects of particular maintenance strategies apparent? Has the pattern changed in ways we have not before realized? The messages and questions of the patterns are local. It falls to those who know the maintenance operation and its history to draw meaning from them and, in turn, be instigated to seek additional insights and questions.

The slope plot, called **trend** in this case, deals with the change from one period to the next. The nature of trends is to be sustained for short periods at the most. The plot shows the count of occurring orders that are generated by a trend rather than reflecting the longer history of workload count. Once again the meaning is local to be interpreted locally.

In the case of Figure 3, the level is multiplied by the plotted **seasonal** factor. This is because we can see in the top and bottom panels that the seasonal element is swinging wider with time. If that were not the case, the model would plot a count to be added to level and trend.

Step 4: Set the budget. The article “The Secret is to Budget and Control Maintenance Opex Dimensionally” explained activity-based budget and control using the concept shown in Figure 4. The monthly budgeted and actual costs are the intersection of job count and cost per job. There are variances whenever the rectangles do not fully overlay. This step will establish the set point for the top side of the budget rectangle.



To begin, the budget team must make a certain decision for each workload group. It is whether or not the component plots of Figure 3 reflect assumptions that will hold up through the budget year. If yes, we will use the Holt Winters model to forecast the year. If no, we may set the forecast “manually.”

For the article, we will continue as if the answer is “yes.” If the answer were “no” we would extract the data table behind the components of the forecasted plot and conduct ad hoc analytics to set budget upon adjusted assumptions.

Figure 5 shows the past forecasted into the budget year. It is modeled based on extending and reconstituting the components of Figure 3. In reality, the forecast would be approximately 18 months because budgets are usually formed before the current year has concluded.




Figure 6 shows the profile at which jobs are forecasted to occur month-by-month. The horizontal line is the forecasted occurrences as a smoothed workload. The plant will likely set the smoothed workload as the set point to the monthly budget rectangle in Figure 4.




Figure 6: Month-to-month and smoothed workload.

However, the set point may be shifted to reflect operational policies and strategies. An example is if the plant, through event history analytics, concludes that the mean time to return to readiness of particular assets can be shorter than it has historically been. The budget set-point for some months would be temporarily set higher in order to adjust backlog to match the new retention policy.

Immediate Value Propositions

Determining the occurring workload is the foundation to workload-based predictive budget and control. Although not expanded upon in this article, there are immediate opportunities to harvest tremendous value from having done the work.

The most obvious is that as the steps unfold, we will find things we never knew we did not know. What is most exciting is that many discoveries will have significant, actionable ramifications for essential aspects of maintenance excellence.

A specific value proposition is to act upon the indisputable phenomenon that workload smoothing is a primary determinant of craft productivity—creating “automated” daily craft loading. Another is to take the guess out of the meaning of backlog and no longer being forced into trying to control maintenance operations through backlog—akin to wagging the dog. The article, “Size Maintenance Capacity on Forecasts, Not Backlog.”

Another yet, is the potential of finding strong lead-lag relationships between some workload groups and production losses. This is known as cross-correlation analytics in time series.

A final note. So much maintenance intelligence, such as hours, dollars and events, are presented to us as raw time series. We have all struggled to find meaning in them. Time series analytics as described in this article allows us to break through the fog and the dark of raw intelligence.

Related papers: [ROACE: Financial North Star to Maintenance and Reliability Operations](#) | [The Secret Is to Budget and Control Maintenance Opex Dimensionally](#) | [Size Maintenance Craft Capacity on Forecasts, Not Backlog](#)



This work is licensed by Richard G. Lamb under a [Creative Commons Attribution 4.0 International License \(CC BY 4.0\)](#).