

Save money, raise asset productivity: Why maintenance staffing matters

The battle for higher asset productivity is won or lost at the front line. Staffing decisions are crucial. The best way to make them turns out not to be top-down, but bottom-up.

April 2018

by Matt Gentzel, Brian Green and Drew Horah

Manufacturing and other asset-intensive industries face a continual need to reduce costs while increasing operational reliability. These objectives frequently come into conflict when companies make decisions about maintenance staffing levels. Maintenance staffing is easy to get wrong: cut too deep or too fast and reliability suffers. And mistakes are tricky to fix: in many parts of the world, tight labor markets, coupled with declines in vocational-training programs, are making recruiting, hiring and training maintenance staff significantly more difficult.

The way many companies plan their current and future maintenance requirements isn't helping. At worst, current practices can be downright harmful. Some organizations try to calibrate their maintenance hiring practices, but only according to high-level estimates of future maintenance requirements. Others avoid planning altogether and instead set hiring goals at or below natural attrition rates. And many plant managers have their hands tied by financial constraints imposed from above, forcing them to lay off staff or scale back recruitment efforts.

Too often, the result exacerbates shortages of skilled personnel, since cutting short-term costs can seem more urgent than ensuring adequate long-term maintenance resources. And the leaders closest to the problem—site level maintenance and operations managers—often struggle to communicate the downside of short-sighted decisions, or to provide alternative analysis to back up their arguments.

Even approaches that are more rigorous can be prone to failure. Some companies make detailed wrench-time measurements to determine their current maintenance workload. Then they project future requirements based on productivity increases they hope to capture through the use of lean techniques and new, digitally-enabled approaches. As they implement their plans, however, workforce cuts tend to happen faster than the productivity improvements needed to fill the gaps.

The outcome of these failures is all too predictable. Without enough staff, maintenance backlogs get longer and unplanned failures rise. Because unplanned work is more difficult and time-consuming to complete, staff availability declines even further, a spiral that eventually leads to excessive downtime, rising costs, and lost sales: exactly the situation companies want to avoid.

A better way

Companies need a better solution for planning maintenance workforces. Such an approach would make effective use of the available data to estimate demand for maintenance craft hours. It would combine long- and short-term planning, with year-by-year forecasts for at least five

years. It would be sufficiently transparent to enable robust discussions, especially about the assumptions underlying supply and demand estimates—while also generating ideas to get supply and demand to match. Finally, and most importantly, it would solicit input from those closest to the issue: maintenance and reliability managers are the best people to drive the analysis, drawing on input from other stakeholders including operations, finance and HR.

Such an approach need not be excessively arduous. We recommend that companies begin with two three-step processes, one to agree on future maintenance demand, the other to understand current and future level of supply available to meet that demand.

Sizing maintenance workforce demand

The first step is to define the baseload of work performed by both internal maintenance workforce and contractors. This should include all preventative maintenance and typical levels of planned and reactive repairs. Additional work performed by maintenance personnel, such as planning, scheduling, or related activities should also be included. Most organizations can obtain this information readily from their maintenance records.

Next, they should develop assumptions about how the baseload will change over time. This will include efficiency-improvement projections by work type, based on available data and realistic assumptions around the speed at which the site can reach best-in-class. It will also consider the impact of other improvement efforts, such as shifting work from unplanned to planned, increasing preventative maintenance compliance, and major projects likely to reduce unplanned maintenance demand, such as equipment replacement or the introduction of condition-monitoring technologies.

Finally, companies should estimate the impact of above-and-beyond projects. This part of the demand calculation incorporates any maintenance labor needs for one-time events, such as unusual capex spend, implementation of maintenance programs such as reliability centered maintenance (RCM), or surge programs such as backlog reduction.

Understanding workforce supply

To understand how their maintenance workforce is likely to evolve over the five-year planning period, companies should begin by projecting the size of the existing workforce into the future, based on historical rates of attrition. Next, they should look at the demographics of their existing workforce. If a large number of personnel are approaching retirement age, the company will need to consider how much time will be required to recruit and train their replacements. Finally, the company should think about other supply-improvement levers. These can include increasing the number of maintenance tasks allocated to equipment operators, or the transfer of personnel between roles.

Closing the gap

The outcome of this initial analysis phase will be a clear picture of the likely future gap between maintenance supply and demand, whether that gap is a shortage or a surplus. Crucially, because the data and assumptions that feed into the analysis have been agreed by site

maintenance leaders, they are also prepared to take ownership of the results.

With site-level support, organizations can then begin a process of brainstorming with local maintenance leaders to generate potential solutions to the identified gaps and assess their feasibility. Once everyone agrees on the best way forward, action plans can be assigned to appropriate teams. Naturally, a regular cadence of reviews is essential to ensure everything is progressing as it should, and to adapt to the plan as the needs of the organization evolve.

This approach allows for a well-controlled, well-managed evolution. Over a five-year period, surpluses can often be addressed by natural attrition, or by training and transferring personnel to new roles. With time to act on upcoming shortages, companies can put appropriate recruitment, training or contracting strategies in place.

□ □ □

Planning and delivering the right maintenance workforce will always be challenging, especially when costs are under pressure. The bottom-up approach described here gives companies better chance of achieving a plan that actually works. Local data leads to local projections, which lead to local recommendations, which lead in turn to local action plans and ownership ■

Matt Gentzel is a partner in McKinsey's Pittsburgh office, ***Brian Green*** is a senior implementation coach in the Houston office, and ***Drew Horah*** is an associate partner in the Atlanta office.

Copyright © 2018 McKinsey & Company, Inc. All rights reserved.