



Better planning: three principles for reliable delivery commitments

21 August 2026, 16:30

Pascal Pollet

Why good planning means more than just making plans

Despite a great deal of planning, on average only 65 per cent of orders are delivered as promised. Discover how safety time, buffer time for urgent work and a simple capacity model lead to more reliable and less stressful planning.

Manufacturers spend a great deal of time on planning. Research carried out by Sirris in industry shows that, on average, each production worker spends 15 to 20 minutes a day on planning and follow-up tasks. In a company with a hundred production staff, that represents 25 to 33 hours a day. Despite all this work, the average delivery reliability rate stands at just 65 per cent. In other words, more than one in three orders is not delivered as promised.

That's not just a nuisance for the customer. A late delivery often causes fresh problems further down the supply chain. The customer has to adjust their own schedule, staff and machinery are left waiting, other orders are rescheduled and new urgent requests arise. In this way, companies pass the chaos on to one another.

What is remarkable is that some companies do manage to achieve a high level of delivery reliability, with far less planning effort. They don't necessarily have smarter planners or more advanced software. Above all, they have a different approach to planning.

We are therefore starting **this series of articles on better planning**. In this first article, we look at how companies can determine reliable delivery dates. In doing so, we start from **three basic principles**.

A truck unloading in loading bay

Principle 1: Make delivery reliability the starting point

In many companies, a delivery date is estimated first. Only afterwards does the company measure how many orders were actually delivered on time. Delivery reliability thus becomes a result of the planning. That's completely the wrong way round.

A company should determine in advance the level of delivery reliability it wishes to achieve. Only then can the necessary delivery time be calculated. It's not the promised delivery dates that should determine delivery reliability; rather, the desired level of delivery reliability should determine the promised delivery dates.

That calls for a different way of thinking about lead times. Due to disruptions, quality issues, material shortages, staff absences and varying processing times, an order never has an exactly predictable lead time. There is always some variation. Some orders are processed more quickly than expected, while others take longer. The figure below shows the typical distribution of lead times.

Variation in lead time, with average lead time and safety time.

A reliable delivery commitment therefore consists of two parts:

$$\text{Promised delivery time} = \text{expected lead time} + \text{safety margin}$$

The expected or average lead time describes how long an order normally takes. The safety margin protects the customer against the inevitable deviations from that expectation.

The required safety time depends on the desired delivery reliability and on the variability of the production process. If you want greater delivery reliability, it's best to build in more protection. If you make the processes more stable, you will need less protection.

Safety time is therefore not just an arbitrary margin that a planner adds 'just to be on the safe side': it's a control knob that allows the desired delivery reliability to be set as required.

Principle 2: Plan for urgent work rather than reacting to it

Virtually every manufacturer has to deal with rush orders. These range from a customer's urgent requests to rework, a quality issue, a missing part or an unexpected repair.

Despite this, these tasks tend to be treated as exceptions time and again. As a planner, you then have to decide which order to bring forward and which other orders will therefore be held up. But when urgent work is a regular occurrence, it isn't an exception: it's an integral part of demand and must be planned for accordingly.

An urgent order can only reduce its own waiting time by increasing the waiting time for other orders. As a result, lead times become more varied and it becomes more and more difficult to make reliable delivery commitments.

The most robust solution is therefore to set aside time in advance for urgent and unexpected work. This means that the schedule must not be fully booked. When every available minute has been allocated to an order in advance, any disruption, absence or urgent job automatically results in other orders being delayed.

A schedule that allocates only part of the theoretically available capacity in advance has space to cope with disruptions. The amount of space required depends on the variability and the volume of urgent work. A planned workload limit of 80 to 85 per cent may be a useful starting point, but it is not a universal standard.

That doesn't mean that employees do nothing for the rest of the time. The reserved space is generally used for rush orders, rework, unexpected problems, maintenance or improvement activities. If there are fewer disruptions than expected, orders that have already been scheduled can be carried out earlier. The aim is not to leave capacity unused: it is to prevent any disruption from throwing the entire schedule into disarray.

Principle 3: Choose the simplest capacity model that works

A capacity model is a simplified representation of how much work a production system can handle over a given period. It determines which resources are included, the unit in which the workload is expressed (such as items or hours), and the time scale used for planning (such as days or weeks).

Greater detail may seem more accurate, but it also increases the need for correct and up-to-date information. When this is lacking, the main result will be a false sense of precision. A simple model that reflects the key capacity constraints well is therefore often more robust than a detailed model that only partially describes reality.

Sometimes it is enough simply to plan for a major bottleneck. In other situations, a classification by product family works better. Even simple units, such as the number of orders, items, kilograms or installation days, can be sufficiently accurate.

The rule of thumb is simple: choose the roughest capacity model that still provides reliable delivery dates. Complexity is only sensible when simplicity is demonstrably inadequate.

Different ways of determining a delivery date

In the following articles in this series on planning, we will apply these principles to various practical methods for committing to reliable delivery dates. We will always look for the simplest approach that is suited to the production environment and delivers the desired level of delivery reliability.

Ready to improve delivery reliability in your business?

Would you like to make your production planning more reliable? The experts at Sirris will be happy to help you analyse and improve your planning approach. Together, we'll look at which method is best for your production environment and business objectives.

[Contact our expert](#)

Take the QRM Silver daytime or evening course

Want to shorten lead times and improve delivery reliability? Choose the QRM Silver course that best fits your schedule and learn from experienced QRM experts how to apply QRM in your own production environment.

[Register for the daytime course](#) [Register for the evening course](#)

The course is taught in Dutch.

Authors



Pascal Pollet