



SHOP SCHEDULING

Which scheduler does your shop actually need?

Infinite, finite, or something smarter in between. A practical look at scheduling for make-to-order shops.

The choice you shouldn't have to make

Ask two manufacturers what they want from a scheduler and you'll often hear the same thing: a plan the floor can actually hit, without a year of setup to get there. Most scheduling software makes you choose between those two, and neither choice is comfortable.

This guide walks through the two classic approaches, why each one disappoints make-to-order shops, and a third path that gets you most of the benefit for a fraction of the work. It finishes with what an AI-native ERP changes about the whole problem.

Option 1: Infinite-capacity scheduling

How it works

Infinite-capacity scheduling works backwards from due dates and assumes every machine is always available.

Why it's appealing

It is simple to understand and quick to stand up. Everyone can read the plan, and it gives you a date for every job straight away.

Where it breaks

It ignores reality. Because nothing is ever fully loaded in its world, it hands you dates that are physically impossible to hit. The floor quietly learns not to trust the plan, and you are back to promising on gut feel.

Option 2: Finite-capacity scheduling

How it works

Finite-capacity scheduling respects real machine and labor limits, so the plan it produces is genuinely achievable.

Why it's appealing

On paper it is the right answer: dates that account for capacity, not wishful thinking.

The catch

It only works if your routings, setup times and cycle times are accurate. Getting there is usually a six-month data project, and the schedule is not worth trusting until it is done. Shops that start with full finite scheduling often end up back on the whiteboard while they wait for the data to catch up.

The middle path: bottleneck-first scheduling

For most make-to-order shops, the right answer sits between the two. Three moves get you there.

1. Schedule the whole floor on lead times

Keep the overall plan simple and readable. Everyone can see where a job sits and where it is going next, without a specialist to interpret it.

2. Watch capacity as a load view

Instead of a hard finite schedule everywhere, show load against capacity so overloads surface as exceptions rather than silent broken promises.

3. Schedule tightly only where it matters

Apply real finite scheduling to the one or two work centers that genuinely constrain you: the press brake, the CNC cell, the paint booth. That is where the accuracy actually pays off.

THE PAYOFF Most of the accuracy of a full finite model, for a fraction of the effort.

It also matches how your team already thinks about the floor. They know where the constraint is. Bottleneck-first scheduling simply makes the system agree with them.

Keep a person in control

Whichever approach you take, keep a human in the loop. The schedule should show you the conflict and let you decide, not quietly resequence the whole shop overnight.

Automation that hides its reasoning erodes trust fast. Automation that surfaces the problem and leaves the call to you builds it. The goal is a scheduler that makes your planner faster and more confident, not one that takes the wheel when no one is looking.

What an AI-native ERP changes

The hardest part of finite scheduling has always been the data. This is where an AI-native ERP changes the equation.

The data problem, solved over time

Tangle learns real setup and run times from what actually happens on the floor, and keeps correcting itself. You get capacity-aware dates you can commit to, without the traditional six-month implementation project up front. The schedule gets more accurate as you use it, rather than demanding perfection before it earns its keep.

A scheduler that fits your shop

Because Milo, Tangle's native AI engineer, builds inside Tangle, the scheduler is not a fixed product you have to adapt to. If you schedule by crew rather than machine, run two shifts on one cell, or need a shared outside-processing step to appear in the sequence, Milo can shape the scheduling views and rules around how your shop really runs. You describe it in plain English, and preview the change before anything goes live.

So which do you need?

A quick way to place yourself:

- Still on spreadsheets or the whiteboard? Start with lead-time scheduling and a load view. It is a big step up for very little effort.
- Know your constraint? Add tight, finite scheduling at that one or two work centers, and leave the rest on lead times.
- Tempted by full finite everywhere? Only worth it if your routing and time data is genuinely clean, or you have a system that will learn it for you.
- Buying new? Ask each vendor to schedule one of your real, awkward jobs, and to show how the plan changes when a machine goes down.

See it on your own floor.

Bring us a job that is hard to schedule today, and we'll show you how Tangle plans it, and how Milo shapes it around your shop.

Book a demo at tangle.io/demo