




ccPlanning

ccplanning.net

WORKFORCE PLANNING WHITE PAPER

The Resilient Plan

Planning through change, disruption and the unexpected



John Casey

Workforce Planning Expert · ccplanning.net

The Resilient Plan

Planning when the ground keeps moving — change, disruption and the unexpected

Most workforce planning is taught as if the world holds still: a forecast meets a fixed workforce, a roster covers a stable curve, and the only question is how cleverly you match the two. Real operations don't hold still. Systems get migrated, a big intake reaches its leaving cliff, a strike or a storm takes out half the floor, and the demand you forecast last week is already drifting from the demand arriving today. A plan that only works when nothing changes isn't a plan; it's a wish. This paper is about the other kind — the resilient plan, built to bend with change rather than break against it.

Resilience isn't a single technique. It's a habit of mind that shows up in four recurring situations every planner faces, and a handful of structural moves that make all four survivable. We'll take the four in turn — the productivity dip of a change, the cohort cliff of attrition, the judgement of when to re-plan, and the partial-workforce day — and then pull out the common thread: buffers, flex, playbooks and a planning rhythm that sees trouble coming.

Plan for the dip you will definitely have.

1. The change you choose: budgeting the go-live dip

Every major system migration — new telephony, a new CRM, a new WFM platform — is sold on the steady state it will eventually deliver and planned as if that steady state arrives on go-live morning. It never does. For weeks, sometimes months, agents are slower and less sure on the new tool while trainers and floor-walkers pull experienced people off the phones. Productivity follows a J-curve: handle time jumps, first contact resolution dips, shrinkage rises with training and hypercare, and only then does the curve climb back — settling, if the system is any good, above the old baseline.

The resilient response is to treat the dip as a near-certainty to be budgeted, not a risk to be hoped away. The two numbers that matter are the *depth* of the dip (how much worse, at the bottom) and its *length* (how many weeks to recover). Build both into the plan: a temporary handle-time uplift and extra shrinkage for a defined window, a lower expected occupancy while people learn, and a deliberate hypercare buffer in the first weeks — the cheapest insurance a migration has. Phase the cutover where you can so the whole operation doesn't hit the bottom of the curve at once, agree the recovery ramp with the project team, and track against it, so 'three weeks behind the ramp' is a fact you act on rather than a surprise you explain after the quarter.

2. The change that finds you: the attrition cohort cliff

Capacity plans usually treat attrition as a steady percentage — lose two per cent a month, hire to match. Real attrition is lumpier and more dangerous, because people leave on a tenure curve with a pronounced cliff in the first year. Hire a big cohort all at once — a peak intake, an acquisition, a class to cover a launch — and you have quietly scheduled a synchronised exit months later, when that whole group reaches the leaving cliff together and walks out of a hole a flat assumption never saw.

The defence is to stop modelling attrition as one number and start modelling it by cohort and tenure. Track survival curves for your intakes, project each cohort forward separately, and the synchronised dip appears in the plan months before it appears on the floor — exactly when you can still act. Then smooth the input: stagger intakes into smaller, more frequent classes so the exits spread out instead of stacking up. Watch early-tenure attrition as a leading indicator, and carry a deliberate buffer for the cliff you can't avoid. Attrition will always cost you; planning it as the wave it really is, rather than the trickle the spreadsheet prefers, is the difference between a managed gap and a service cliff.

3. The change in the numbers: when to re-plan, and when to hold

Every forecast is wrong; the only question is by how much and what you do about it. A week in, actuals are running above plan — do you re-forecast and re-plan, or hold your nerve and let it settle? Re-forecast too eagerly and you churn the plan on noise, exhausting the operation with constant change. Hold too long and you defend a number everyone can see is broken. Two questions decide it. First, is the variance signal or noise? A single odd day is almost always noise; a consistent, widening lean in the same direction is signal — forecasts are supposed to be wrong day to day around the right average. Second, is there still time to act? A re-forecast you can't resource is just paperwork.

So the rule of thumb is simple: re-forecast when the variance is genuine bias *and* there is still lead time to do something about it; absent either, hold. Distinguish the horizons, too — an intraday re-forecast feeds today's flex levers and can be made on lighter evidence because it's cheap to act on; a re-forecast of the weeks ahead changes hiring and rosters and demands a higher bar of proof. And separate updating the forecast from blaming it: a forecast revised on genuine new information is doing its job, not failing.

4. The change that hits you: the partial-workforce day

Most planning assumes the workforce turns up. Then a rail strike, a snowstorm, a flood or a power cut arrives and a large slice of your people simply can't — and the question stops being 'how do we hit service?' and becomes 'what do we do with the half we've got?' The instinct to ask everyone to work harder doesn't scale; you can't heroics your way out of losing forty per cent of the floor.

The workable answer is triage, decided in advance. Protect the contacts that genuinely matter — the vulnerable customer, the safety issue, the regulated deadline — and deliberately defer or deflect the rest, with honest customer communication so people aren't left waiting on a promise you can't keep. That means a prioritised routing plan, a reduced-scope service you can switch to, and a clear hierarchy of what gets answered when there isn't enough of anyone. Agree the tiers — what's protected, deferred and deflected at, say, 80%, 60% and 40% of normal capacity — pre-write the customer communications and routing changes, decide who can declare the disruption and flip the service, and rehearse it occasionally. You can't stop the snow; you can decide, while the sun's out, exactly how you'll run the day it lands.

Resilience is built in advance, in the calm, on purpose.

5. Building resilience in: the common thread

The four situations look different but share a cure. **Buffers**: a resilient plan runs with deliberate, costed slack where variation is known to bite — the hypercare weeks, the attrition cliff, the spike days — rather than planning to the average and being surprised by the distribution. **Flex**: the levers that close the gap on the day — overtime, shift swaps, banked hours, callbacks and deferrable work — are agreed and priced before they're needed, not improvised at 7am. **Playbooks**: the responses to the predictable shocks are written down, with decision rights named, so the operation reads a plan rather than inventing one under pressure. And a **planning rhythm**: nested loops from the live minute to the annual budget, each looking at its horizon on a cadence that matches how fast it can change, so trouble is somebody's job to spot before it arrives.

None of this is exotic, and none of it is free; resilience is a cost you choose to carry so that the bad day is survivable instead of catastrophic. The unresilient plan looks cheaper right up until the migration, the cohort cliff, the bias or the storm arrives — and then it pays for its thrift all at once, in missed service, burnt-out people and a planner explaining after the fact why the number that looked so precise was so wrong. The resilient plan costs a little more every week and saves you on the weeks that matter. The ground will keep moving. The only real choice is whether your plan was built to move with it.

Go deeper at ccplanning.net — the articles behind this paper: planning through a system migration, planning for an attrition wave, re-forecasting, and planning for disruption. Plus free calculators, the Academy, and the rest of the white-paper series.