



ccPlanning

ccplanning.net

OPINION WHITE PAPER

The Occupancy Trap

The efficiency metric that is really a burnout dial — and the spiral it drives

John Casey

Workforce Planning Expert · ccPlanning

A opinion piece in the ccPlanning white-paper series · read the rest at ccplanning.net

Contents

- 1. What occupancy actually is..... 3
- 2. It is a consequence, not a lever..... 3
- 3. The trade-off you are really making..... 3
- 4. The spiral..... 4
- 5. What to do with it instead..... 4

1. What occupancy actually is

Occupancy is the proportion of an agent's logged-in, available time spent handling contacts rather than waiting for the next one. It is not utilisation, it is not adherence, and it is emphatically not a number to maximise.

People conflate it constantly. Utilisation usually folds in shrinkage; adherence is about being on the rostered activity; occupancy is specifically the busy-versus-waiting ratio while available. At 83% occupancy an agent spends roughly five minutes of every six on contacts and one waiting — and that waiting is not idleness to be eliminated. It is the slack that keeps the queue short.

2. It is a consequence, not a lever

Here is the point most operations miss: occupancy is an **output** of your staffing and your service-level target, not an input you set. For a given offered load, the maths of queueing fixes the relationship — staff to hit 80% in 20 seconds and occupancy lands wherever the maths puts it. That figure rises with size, because large queues pool randomness more efficiently, so a 20-seat team and a 500-seat operation hitting the identical service target can sit at very different occupancies. Which is why a single occupancy target applied across teams of different sizes is meaningless — it asks small teams to do something the maths forbids.

Occupancy has a healthy band — and a danger zone



Above the low-90s isn't efficiency — it's a warning, and a self-reinforcing one.

Every leaver pushes occupancy higher for those who remain, which drives the next leaver.

You cannot set occupancy. You can only set staffing and a service target — occupancy is what falls out. Targeting it directly is targeting a symptom.

3. The trade-off you are really making

Occupancy and service level pull against each other. To answer quickly you need agents waiting and ready, which means some idle time — lower occupancy. Squeeze that idle time out to lift occupancy and you remove the very buffer that absorbs random arrival spikes, so service becomes volatile and abandonment climbs. The slack that looks like waste on an efficiency dashboard is the thing keeping your service level stable. There is real idle time worth recovering through better scheduling and honest shrinkage control — but beyond a point you are trading customer experience for a vanity number.

4. The spiral

Sustained high occupancy is a wellbeing problem before it is a service one. The seconds between contacts are how agents reset after a difficult call, take a breath, finish a thought. Strip them away and you get rising errors, longer-term sickness, disengagement and attrition — and attrition is enormously expensive. Worse, it is self-reinforcing: every leaver pushes occupancy higher for those who remain, which drives the next round of leavers. Plenty of operations have run a team into the ground chasing a 95% target without ever connecting it to the attrition line right next to it. It is a spiral, not a dial.

5. What to do with it instead

Treat occupancy as a health check, not a target. Watch it by team and interval, set a sensible ceiling for sustained periods, and when it runs hot read it as a signal that you are under-staffed or your shrinkage assumptions are optimistic — not as proof your people could simply try harder. Recover genuine idle time through scheduling that matches the curve and through honest shrinkage planning, but protect the breathing room that keeps both your service level and your people stable.

The most efficient operation is not the one with the highest occupancy. It is the one that hits service, holds its people, and runs occupancy at the level the maths and the humans can both sustain.

Occupancy above the low-90s is not a sign of efficiency. It is a warning — read it as one.

About the author

John Casey has spent more than 30 years in contact-centre workforce planning, including roles as a workforce-planning manager and operations director, and now writes and teaches at ccplanning.net. The views in this paper are his own. It represents general professional experience and is not affiliated with, nor does it represent, any current or former employer.

ccplanning.net · Free guides, calculators and the rest of the white-paper series.