




ccplanning.net
workforce planning, done well



A C C P L A N N I N G . N E T W H I T E P A P E R · O P I N I

Kill Your Dashboard

The contact-centre metrics we should stop worshipping — and what to measure instead.



February 2027 · ccplanning.net

Contents



Right-click and choose "Update Field" to populate page numbers in Word.

A manifesto, not a masterclass

The previous papers in this series were even-handed. This one is not. It is an argument — that the contact-centre industry measures the wrong things, defends them out of habit, and quietly lets a handful of sacred metrics do real damage to service, cost, and people. The dashboard you check every morning is, in several places, lying to you, and the longer you trust it the worse the decisions it drives.

This is deliberately provocative, because the metrics under attack are so entrenched that a measured tone would let them survive unexamined. None of what follows is anti-measurement — the opposite. It is an argument for measuring what matters instead of what is traditional, and for having the courage to retire the numbers that have outlived their usefulness. Disagree freely; that is the point of a manifesto. But check your own pack against it before you do.

The thesis in one paragraph

Most contact-centre dashboards are built from metrics that became targets, got gamed, and now optimise the measurable over the meaningful. Five sacred cows do most of the damage: the service-level target that hides the distribution, the occupancy figure that burns agents out, the AHT target that rewards speed over resolution, adherence used as a stick, and the composite quality score that hides the drift. Each was once useful and is now mostly a habit. Kill them — or at least demote them — and replace them with measures of what the operation is actually for: resolving customer contacts well, at a sustainable cost, with people who stay.

1. Why your dashboard lies to you

The root problem is older than contact centres. Goodhart's law: when a measure becomes a target, it ceases to be a good measure. The moment a number is something people are managed against, they optimise the number — and the number detaches from the reality it was supposed to represent. Every metric on your dashboard that carries a target is subject to this, and the longer it has carried one, the further it has drifted from the truth.

Why your dashboard lies: Goodhart's law in four steps

When a measure becomes a target, it stops being a good measure. The contact centre wires the loop tighter than anywhere.



The tighter you wire a proxy to pay, ranking, or status, the faster it rots.

This is not because anyone is dishonest. It is the rational, predictable response of people managed against a number. Tell a team that average handle time must come down and it will come down — by rushing customers, skimming on notes, and pushing the unresolved contact to a callback that does not count against today's figure. The metric improves; the operation gets worse. The dashboard shows green while the thing the dashboard was meant to protect quietly degrades. A planning function that does not understand this is steering by instruments that have been bent by the act of watching them.

When a measure becomes a target, it stops being a good measure. Every number on your dashboard that carries a target has drifted from the truth in proportion to how long it has carried one.

The reason this matters more in contact centres than almost anywhere is that the metrics are unusually visible and unusually tied to consequences. An agent can see their adherence and AHT in near real time; a team leader is ranked on their team's service level; a centre head reports occupancy upward monthly. The feedback loop between number and behaviour is tight and constant, which means the gaming is fast and pervasive — not occasional gaming by bad actors, but continuous, rational adjustment by everyone, all the time. The tighter you wire a proxy metric to pay, ranking, or status, the faster it rots. And the contact centre wires them tighter than any other function.

2. The five sacred cows

Five metrics do most of the damage. None is worthless; each was once a reasonable answer to a real question. But each has hardened into a target that now distorts more than it informs, and each deserves either retirement or a sharp demotion from the headline of your pack.

The five sacred cows

Each was once useful. Each is now a target that distorts more than it informs. Retire or demote.


- 1 · Service-level target (80/20)**
 arbitrary, gamed, blind to the distribution it hides
- 2 · Occupancy**
 a burnout machine sold as an efficiency virtue
- 3 · AHT as an agent target**
 rewards speed over resolution; manufactures repeat contacts
- 4 · Adherence as a stick**
 punishes the agent for the planner's schedule error
- 5 · The composite quality score**
 hides the drift it was built to summarise

Sacred cow #1 — the service-level target

Eighty percent of contacts answered in twenty seconds. It is the most sacred number in the industry, and it is arbitrary, gamed, and blind. Arbitrary, because almost no operation chose 80/20 deliberately — they inherited it from a vendor template or a previous job, and it bears no relationship to what their customers actually need. Gamed, because a half-second answer-then-hold, or a strategic understaffing of the low-volume intervals, can flatter it. And blind, because a single percentage hides the distribution entirely: an operation hitting 80/20 can still be leaving its most vulnerable customers waiting ten minutes, because the average is carried by the easy contacts. Service level is not useless — but as a headline target it conceals more than it reveals, and it should be demoted to one input among several, reported alongside the distribution it hides.

Same service level, very different experience

Two operations both hit 80% in 20s. The average hides who is waiting. The distribution does not.




Sacred cow #2 — occupancy

Occupancy — the share of logged-in time agents spend handling contacts — is treated as an efficiency virtue: higher is better, drive it up. This is actively harmful. Occupancy above roughly 85% sustained is not efficiency; it is a burnout machine. It means agents have almost no recovery time between contacts, and the cost arrives later as attrition, absence, and quality decline — costs that dwarf the marginal labour saving the high occupancy delivered. Worse, occupancy and service level are in tension: small operations cannot have high occupancy and good service at once, because the maths of

queuing forbids it. An operation chasing occupancy as a target is optimising the one efficiency metric most directly opposed to keeping its people — and paying for it in the attrition line it never connects back to the dashboard.

Occupancy isn't free: the burnout curve

Push occupancy past the sweet spot and the saving is repaid, with interest, in attrition and absence.



The marginal labour saving from high occupancy is dwarfed by the attrition cost it creates — booked to a different line.

Sacred cow #3 — the AHT target

Average handle time is a vital planning input and a destructive performance target. As an input, it sizes the requirement; forecast it well and the plan is sound. As a target imposed on agents — "get your AHT down" — it rewards exactly the wrong behaviour: speed over resolution. The agent who resolves a complex issue in twelve minutes is penalised against the one who closes three calls in the same time and generates two callbacks. AHT targets manufacture repeat contacts, the most expensive failure in any operation, while showing an improving number. Measure AHT, forecast AHT, investigate AHT outliers — but never target it at the agent level, because the behaviour it drives costs more than the time it saves.

Sacred cow #4 — adherence as a stick

Adherence measures whether an agent was where the schedule said, when it said. It is a useful coverage signal and a terrible disciplinary tool, yet it is overwhelmingly used as the latter. The problem: adherence scores the agent against a schedule the planning team built, so when the schedule is wrong — optimistic shrinkage, an aged demand curve, breaks placed in the peak — adherence punishes the agent for the planner's error. The coaching conversation lands on the wrong person, the planning team never learns its schedule was the problem, and trust burns in three places at once. Adherence belongs on the planning team's diagnostic dashboard, not in the agent's performance review.

Sacred cow #5 — the composite quality score

The single quality number — 87% this month, steady — is the most comforting and least useful metric on the pack. It is built to summarise, and in summarising it hides the very thing it should surface: the drift underneath. Accuracy can be collapsing while compliance compensates, and the composite sits unchanged for months until the operation is visibly in trouble. The composite is a glance-and-go number for executives who will not read the components, and it fails at the one job a quality metric has — to flag the underlying movement early. Report the components alongside it, every time, or do not report it at all.

3. The pattern behind all five

Step back and the five share a structure. Each measures something easy to count rather than something that matters. Service level counts speed-of-answer because it is easy, not customer outcome because it is hard. Occupancy counts utilisation because it is on the dashboard, not sustainability because it is not. AHT counts seconds because seconds are countable, not resolution because resolution is messy. Adherence counts schedule-conformance because the system produces it, not whether the schedule was right. The composite counts an average because an average fits on a slide, not the drift because the drift needs a chart.

Every sacred cow measures something easy to count instead of something that matters. The dashboard is full of proxies that everyone has forgotten are proxies.

They are all proxies — reasonable stand-ins, once, for things genuinely hard to measure. The damage is not that they are proxies; it is that everyone has forgotten they are proxies and now manages the stand-in as if it were the real thing. The fix is not more metrics. It is the discipline to keep asking, of every number on the pack: what is this actually a proxy for, and would I be better off measuring that?

4. What to measure instead

Killing metrics without replacing them is just blindness. The replacements are harder to measure — which is exactly why the proxies took over — but they are what the operation is actually for, and the effort of measuring them is repaid in decisions that improve rather than game.

What to measure instead

Outcomes over activity. Distributions over averages. Leading over lagging. The things you'd care about unwatched.

Stop targeting	Start measuring
Service level (the average %)	the wait distribution, by customer segment
Occupancy (higher = better)	cost-to-serve + sustainability (attrition, absence)
AHT (agent target)	first-contact resolution + repeat-contact rate
Adherence (as a stick)	coverage realised (planning diagnostic, not appraisal)
Composite quality score	the components + their trend

For **service**, measure the distribution, not the average — the share of customers waiting beyond a threshold that actually matters to them, segmented for the vulnerable and the high-value, rather than a single percentage that the easy contacts carry. For **efficiency**, measure cost-to-serve and the sustainability of the operation (attrition, absence, the occupancy that produces them) rather than occupancy as a virtue in itself. For

handling, measure first-contact resolution and repeat-contact rate — the outcomes AHT was a poor proxy for — so the agent who genuinely resolves is rewarded, not penalised. For **quality**, report the components and the trend, not the composite. And for the **planning function itself**, measure forecast accuracy honestly (WAPE and bias) and the coverage the schedule designed in, so the function is judged on what it controls.

A fair objection: outcome measures are slower and noisier than the proxies, and an operation needs something to watch today, not a resolution figure that resolves over weeks. True — which is why the answer is not to throw away the operational signals but to stop targeting them. Keep service level and AHT and occupancy on the planning team’s instrument panel as diagnostic inputs, watched and investigated; just take them off the performance scorecard people are managed against, and put the outcome measures there instead. The distinction between a diagnostic instrument and a managed target is the whole game. The same number that is invaluable as a planner’s gauge becomes corrosive the moment it is wired to an agent’s appraisal or a team leader’s ranking.

The unifying principle: measure outcomes over activity, distributions over averages, leading indicators over lagging ones, and the things you would still care about if no one were watching the number. A dashboard built on those is harder to assemble and impossible to game in the corrosive way the sacred cows are — because the only way to improve a true outcome measure is to improve the outcome.

5. How to actually kill a metric

Knowing a metric should die is easy; killing one that the business has watched for a decade is not. Sacred cows are sacred precisely because senior people are attached to them, targets are written into contracts and bonuses around them, and "we’ve always measured it" carries more weight than it should. A manifesto that stops at "these are bad" is useless. Here is how to retire one in practice.

How to actually kill a metric

Don't start by deleting it. Prove the replacement, show the lie, demote, fix the incentives, bring people with you.



A metric that no longer drives the pack but still drives pay will not actually die.

First, **do not start by removing it** — start by adding its replacement alongside it. Report first-contact resolution next to AHT, the wait distribution next to service level, the components next to the composite. Let the better measure prove itself in parallel for a quarter or two. Second, **show the divergence**: find and present the cases where the sacred cow said green and the real measure said red — the AHT that improved while repeat contacts climbed, the service level that held while vulnerable customers waited. Nothing kills a metric faster than a concrete example of it lying. Third, **demote before you delete**: move the sacred cow off the headline and into the appendix before removing

it entirely, so the change is gradual and defensible. Fourth, **fix the contracts and bonuses** the metric is wired into, because a metric that no longer drives the pack but still drives pay will not actually die. And throughout, **bring people with you** — the goal is not to win an argument about metrics but to change what the operation optimises, which only sticks if the people managed by the numbers understand why.

Conclusion: measure what matters, even when it is hard

The contact-centre dashboard is a museum of metrics that were once useful and are now mostly habit. Service level, occupancy, AHT targets, adherence-as-stick, the composite score — each survives not because it earns its place but because killing it requires effort and courage, and green numbers are comfortable. The operations that pull ahead over the next decade will be the ones willing to retire the comfortable proxies and measure the harder, truer things underneath: resolution, sustainability, the distribution, the trend, the outcome.

Green numbers are comfortable. The operations that pull ahead will be the ones with the courage to retire the comfortable proxies and measure the harder, truer things underneath.

You do not have to agree with every cow on the list. But take the argument seriously enough to audit your own pack: for each headline number, ask what it is a proxy for, whether it has been gamed, and what you would measure if you were building the dashboard today with nothing to defend. The numbers that cannot survive that question are the ones to kill. Your operation is not for hitting 80/20 or running at 90% occupancy; it is for resolving your customers' contacts well, at a cost you can sustain, with people who want to stay. Measure that.

About ccplanning.net

ccplanning.net is an opinionated, practitioner-focused resource for contact centre workforce planning — forecasting, scheduling, real-time management, capacity planning, MI, and the leadership of the planning function. It publishes free articles, browser-based planning calculators, and a fortnightly newsletter for working planners.

This is the seventh paper in the series and the first written as an open opinion piece. Earlier papers covered AI, the business case for planning, building a planning function, forecasting, scheduling, and real-time management — most measured and practical. This one is meant to start an argument. Tell us where you think it is wrong: the conversation happens on LinkedIn, and the best disagreements will shape what we write next.

Argue with this

Read it, check it against your own dashboard, and tell us on LinkedIn which cow you would defend and which you would kill. Pair it with the articles on composite metrics, whether service level is a dead KPI, the cost of perfect adherence, and leading vs lagging indicators — all free at ccplanning.net.