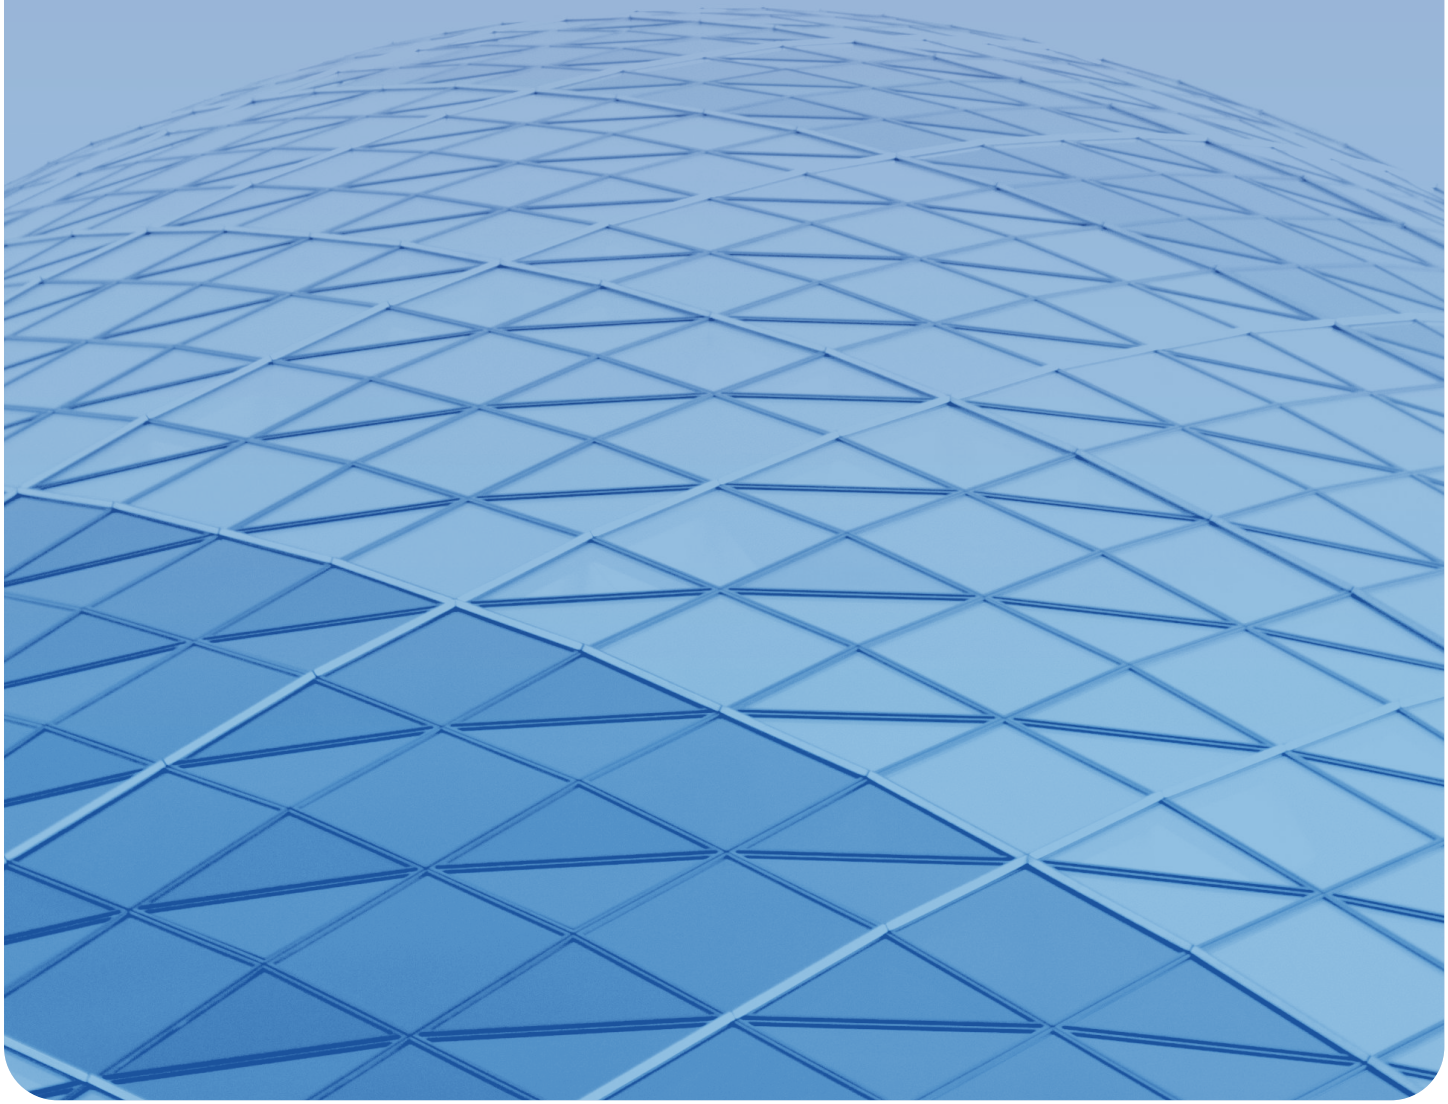
metrikus

Nine checks to run before your occupancy number goes infront of leadership

Most workplace teams can produce an occupancy figure. Far fewercan defend one. These are the nine things worth checking beforeyours is the number a decision gets made on.



What's in here

04 Where did the number actually come from?

05 What's the denominator?

06 Peak or average?

07 Which days are in it?

08 How many weeks?

09 Are ghost bookings in the number?

10 Does every floor behave the same?

11 Does the number match what people say?

12 Could you rebuild it next month, the same way?

14 What a defensible number looks like

15 Three things to do this month

16 Difficult to answer these? Let's talk.



The number and the follow-up question

There is a particular kind of meeting most Heads of Workplace have sat in. You have brought a number — utilisation, occupancy, however your organisation phrases it.Someone senior looks at it, pauses, and asks where it came from.

The honest answer, the one you do not give in the room, is that it came from an export, which came from a system, which measures something slightly different from the other system that also produces a number. You know roughly how it was built.You would not want to defend every part of it line by line.

This is not a competence problem, and it is not unusual. It is structural. Workplace teams are rarely short of building data. What they are short of is a settled, repeatable way to turn that data into a figure that survives questioning — because the data sits in systems bought at different times, by different people, for different reasons, each describing the same building in its own vocabulary.

The cost of that shows up in three places. Time, because every number has to be rebuilt by hand. Trust, because two systems rarely agree and people notice. And decisions, because a figure nobody is confident in tends to get discussed rather than acted on.

This audit is not a methodology and it is not a maturity model. It is nine questions worth asking about a number before anyone else asks them for you. Run them against a figure you have already produced. Most teams find at least two that change the answer — and the ones that do not change the answer still make it easier to defend.

> “
>
> The challenge today isn't collecting more data. Most organisations already have plenty of it. The difficult part is making sense of that data altogether.
>
> **Nate Colle,** Metrikus

Where did the number actually come from?

Badge data, booking data and sensor data measure three different things: who entered the building, what people intended, and what actually got used. A badge figure counts the 20-minute coffee visit the same as a full day at a desk.

Three sources, three different questions. Access control answers who came through the door. Your booking system answers what people intended to do. Sensors answer what was actually used. Most organisations hold at least two of the three and report whichever exports most easily.

The gaps between them are not small. Most access control only records entry, not exit, so what you get is a daily total rather than a picture of the day — the person who came in for a twenty-minute coffee counts exactly the same as the person who sat at a desk for eight hours.

It misses anyone who came through a door behind a colleague, and it can double-count anyone who stepped out at lunch and badged back in. Booking data fails the other way: it is a record of intention, and intention is a weak proxy for behaviour.

None of this makes any one source wrong. Each is a good answer to its own question. The risk is using one to answer another's question — and then sizing an estate against it.

Do this

- [ ] Write the source next to the number, every time it is presented.
- [ ] Pull the same period from two different sources and compare them. The size of the gap is yourworking margin of error.
- [ ] If all you have is badge data, call the output attendance rather than utilisation.

Three sources. Three different questions.

ACCESS CONTROL ANSWERS	BOOKING SYSTEM	SENSORS
Answer How many came through the door	**Answer** What people intended to do	**Answer** What was actually used
Doesn't See Exit, length of stay, tailgating	**Doesn't See** Whether they turned up	**Doesn't See** Who it was, and why

What's the denominator?

“68% occupied” against what? Desks, seats, square metres, or headcount? A utilisation figure without its denominator named beside it isn't a metric.

Percentages feel precise. They are only ever as precise as the thing they divide by.

68%

“68% occupied” could mean 68% of desks had someone at them at some point in the day, or 68% of desks occupied simultaneously at the peak moment, or 68% of usable floor area in active use, or 68% of headcount present in the building. Four numbers. Four quite different decisions.

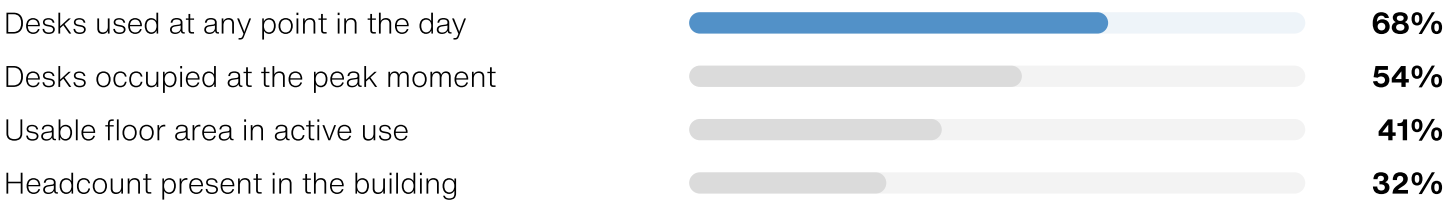
The choice is not neutral, either. An organisation with a generous desk ratio will score low againstdesk count however busy it feels, and an organisation that divides by total headcount will scorelow if that headcount includes fully remote staff. Neither figure is dishonest, but each pushes theconclusion in a particular direction.

The practical failure is more mundane. Two teams in the same organisation produce differentfigures from identical underlying data purely because they divided by different things, and thenspend the meeting arguing about whose number is correct rather than what to do.

Do this

- [ ] Name the denominator in the same sentence as the number, not in a footnote.
- [ ] Agree one organisation-wide definition of utilisation and write it down where both your team andfinance can see it.
- [ ] Where you report more than one measure, label them distinctly — desk utilisation and buildingattendance are not interchangeable.

One day. Four denominators. Four headlines.



Peak or average?

An average hides the day that hurts. 42% average with 91% on a Tuesday isn't underuse — it's bad distribution, and it needs a completely different fix.

An average and a peak argue for opposite decisions.

42%

Run at 42% on average with 91% on a Tuesday and you do not have an underused building. You have a badly distributed one. The average makes the case for handing space back. The peak makes the case that the building fails on the day most people experience it. Both figures are accurate. Only one of them matches what people remember.

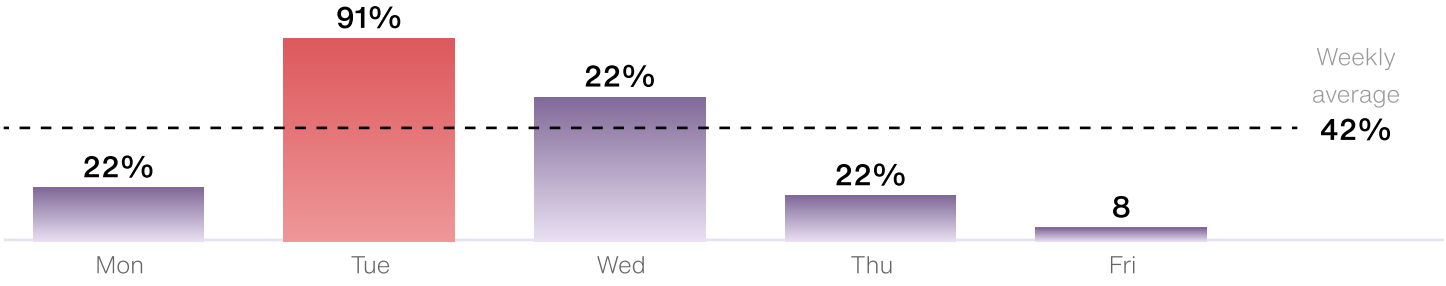
That matters beyond the reporting. Peak-day experience is what drives the complaints, the "there is never anywhere to sit" comments, and the general sense of whether the office works. An organisation that designs to the mean will reliably produce a workplace that feels worse than its numbers suggest.

Most estates are better served by sizing against a defined percentile than against the average or the absolute peak — busy enough to be realistic, not so busy that you are building for the once-a-quarter all-hands.

Do this

- [ ] Report average and peak side by side as standard, never the average alone.
- [ ] Pick the percentile you design against — many organisations settle between the 80th and 90th —and state it explicitly.
- [ ] Check whether your peak is one freak day or a reliable weekly pattern. The response is different.

One week. Two defensible numbers.



42% on average, **91%** on a Tuesday. One figure says give space back; the other says the building fails on the day people remember.

Which days are in it?

A five-day average across a three-day-attendance workforce will always look like under use. Report attendance-day occupancy separately, or you'll argue for the wrong decision.

Hybrid working has made the shape of the week the single most important feature of occupancy data. A five-day denominator buries it.

If your organisation has settled into Tuesday to Thursday, a Monday-to-Friday average builds two guaranteed-quiet days into every figure you produce. The building will look permanently underused no matter how busy it is on the days people were actually asked to attend. You will then be invited to give space back on the strength of it.

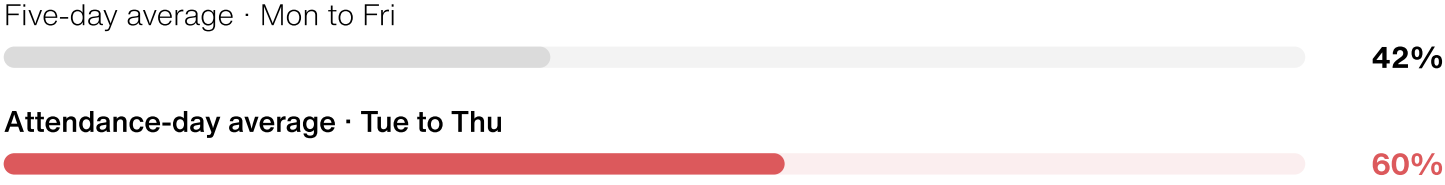
The same week, measured two ways, can move a headline figure by close to twenty points — and only one of those figures is answering the question your policy actually asked.

There is a second thing the five-day average hides: whether people are coordinating. Two organisations can report an identical weekly average, one with demand spread evenly and one with everything stacked on a single day. Those two organisations need completely different interventions, and the headline number cannot tell them apart.

Do this

- [ ] Report attendance-day occupancy separately from the five-day average, and lead with which ever matches your policy.
- [ ] Show a day-of-week breakdown alongside every headline figure.
- [ ] If your policy names specific days, measure against those days — otherwise you are grading the building on a test it was never set.

The same week, two denominators.



How many weeks?

One quiet week is not a trend. One busy Tuesday is not consistent utilisation. Four weeks is the minimum you'd want to draw a conclusion from, and a full quarter is what holds up when someone challenges it.

A single month contains a surprising number of distortions: a bank holiday, a half-term, an all-hands, a major client visit, a week where the lifts were out.

A full quarter smooths most of the noise and is long enough for seasonality to become visible rather than mysterious. The shorter the window, the more confidently you will be wrong —and short windows have a habit of being chosen because they happened to be available rather than because they were representative.

There is also a direction-of-travel point worth making. A single window tells you the level.Two consecutive windows tell you the trend, and the trend is almost always what a leadership team actually wants to know. **“We are at 46%”** invites an argument. **“We have moved from 38% to 46% over two quarters”** invites a decision.

Do this

- [ ] Never present a window shorter than four weeks, and state the window alongside the number.
- [ ] Keep a rolling quarter so you can always show direction rather than just level.
- [ ] Note any known distortions in the period — a fortnight of building works is worth a footnote. The Occupancy



Are ghost bookings in the number?

Desks and rooms booked and never used inflate demand and hide real availability. Check booked-versus-attended for both before you conclude you are short of space.

Your booking system records intention. Your sensors record behaviour. The gap between them is your ghost booking rate, and in most organisations nobody has ever measured it —for desks or for rooms.

The causes are familiar and they run across both. Standing desk bookings held by people whose pattern changed months ago. Recurring meetings that outlived their purpose.Defensive block-booking by anyone who has turned up and found nowhere to sit. And no-shows that nobody cancelled, because cancelling takes effort and not cancelling costs nothing.

Desk bookings matter most, because they are usually what the occupancy figure is built from.

If your utilisation number comes from a booking system rather than from sensors, every unreleased desk is counted as a person. The building reads fuller than it is, and the number gets inflated in exactly the direction that makes the estate look fine.

Rooms fail the other way, and the effect is double-edged. The room estate looks oversubscribed, so the organisation concludes it needs more rooms — while the people who genuinely need one cannot find anything free. You end up paying to solve a problem you do not have, and the real problem gets worse.

Recovering ghost bookings is almost always the cheapest capacity available to a workplace team. It costs a process change rather than a building.

Do this

- [ ] Compare booked hours against sensed-occupied hours across four weeks — desks and rooms separately, not combined.
- [ ] If more than 20% of booked hours go unused, treat it as a booking-behaviour problem rather than a capacity problem. Below that, look at capacity.
- [ ] Turn on auto-release for desk and room no-shows, then re-measure, before you add asingle desk or room.



Does every floor behave the same?

A building-level average with one floor at 85% and three at 30% comes out around44% — a number that describes none of the four floors.

Building-level averages are where almost all reporting happens and where almost no decision gets better. They are comfortable precisely because they flatten the thing that would have been useful. The variation is the signal.

And there is always a reason for it. A team anchors on that floor. The layout suits focused work, or it does not. It is the floor nearest the good coffee, or the one that runs cold in winter, or the one with the only decent meeting rooms.

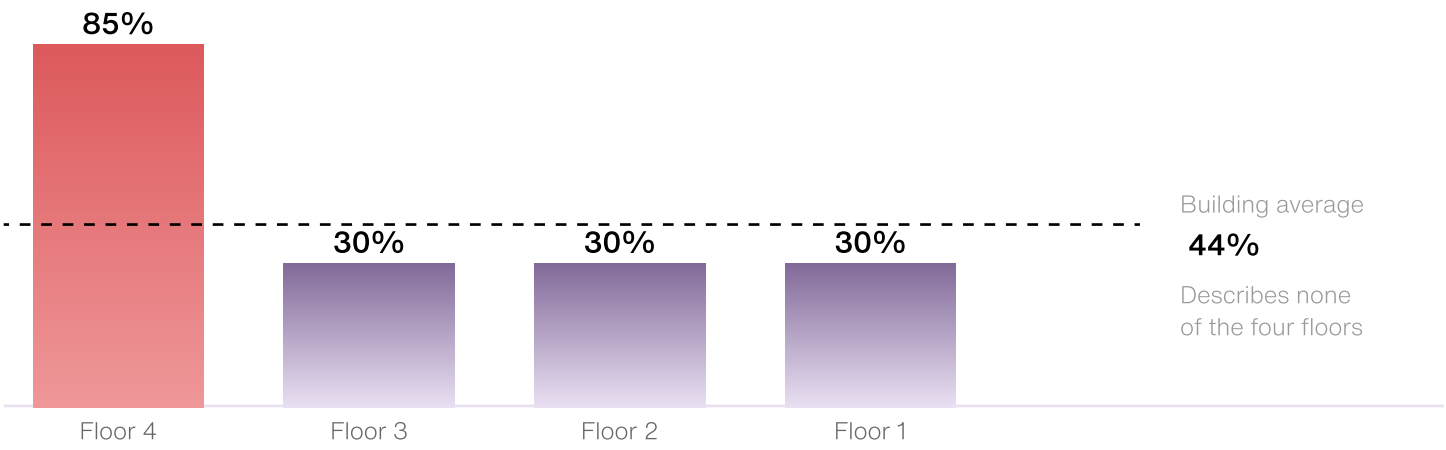
Most of those explanations come with an intervention that costs considerably less than a lease decision.

There is a direction to the risk, too. A building that averages 44% looks like a building with space to give back — so the space goes. And because the decision was made on the average, it usually comes off the floor that was already at 85%.

Do this

- [ ] Break every headline number down by floor before you present it.
- [ ] For any floor more than fifteen percentage points away from the building average, find out why before drawing a conclusion.
- [ ] Check whether the variation holds week to week or was a one-off. Stable variation is a design finding; unstable variation is usually a temporary project team.

One week. Two defensible numbers.



Action ⟶ Rank your floors before you present anything. The outlier is where the decision is.

Does the number match what people say?

If the data says 40% and employees say it's rammed, the data isn't wrong — it's answering a different question, usually about rooms rather than desks.

This is the credibility check, and it is the one most often skipped. When the data says 40% and the people in the building say it is rammed, the instinct is to defend the data. Usually both are right and they are describing different things. Desks get measured; meeting rooms get experienced. Averages get measured; peaks get experienced. The whole building gets measured; one floor gets experienced.

That gap matters commercially, not just socially. A number that contradicts what everybody in the building believes will not change anyone's mind, however well sourced it is. It will simply mark the data — and by extension your function — as out of touch, which is expensive to recover from. Reconciling the two is far more persuasive than either on its own. “The building runs at 40%, and the reason it feels full is that our meeting rooms run at 88% on Tuesdays” is a sentence that gets a budget.

Do this

- [ ] Sense-check any significant figure with two or three people who use the space daily before you present it.
- [ ] Where data and perception diverge, present both and explain the difference rather than picking aside.
- [ ] Treat a persistent mismatch as a gap in what you are measuring, not a problem with how employees perceive things.

Measured and experienced are not the same thing.

WHAT GETS MEASURED	WHAT GETS EXPERIENCED
Desks	Meeting rooms
The average day	The busiest day
The whole building	One floor, one team

Action ⟶ Check the right-hand column against two or three people who use the space daily.

Could you rebuild it next month, the same way?

If it took three days of spreadsheet work, you don't have a metric. You have a one-off report that will quietly disagree with the next one.

Reproducibility is what separates a number you can track from a number you can only present. If the method lives in one person's head and in the particular way they filtered an export on a Thursday afternoon, the next version will differ — and that difference will be read as a change in the building rather than a change in method.

The time cost compounds as well. The days each month spent reconciling exports are the hidden tax on every workplace team that has never written its method down, and that is time taken directly from acting on what the data says.

Writing the method down is unglamorous, but it is the single highest-return thing on this list.

Do this

- [ ] Document the method on one page: sources, filters, denominator, window, exclusions.
- [ ] Time how long reproduction takes. If it is more than half a day, that is the thing to fix before anything else.
- [ ] Have a second person reproduce the number once, working only from your written method.



“

The hidden cost that you never really see is time, trust, and wasted effort.

Nate Colle, Metrikus





What a defensible number looks like

A defensible occupancy figure is not a more accurate one. It is one where you can answer every follow-up question without hesitating — where it came from, what it divides by, how long a window it covers, and what it deliberately excludes.

If you cleared all nine checks, your number is in better shape than most. Write the method down while it is fresh, so the next one matches and the comparison means something.

If two or three gave you pause, that is the normal result and a more useful outcome than a clean pass.

The point was never a perfect figure. It is knowing which parts of your number you would defend and which you would caveat, before anyone else forces the issue.

Three checks cause the most trouble in practice. The first, because the source is usually whichever system exported most easily rather than the one that answers the question. The fifth, because short windows get chosen for availability rather than representativeness. And the seventh, because building-level averages are comfortable and almost always conceal the thing worth acting on.

None of those require new hardware to fix. They require a decision about method, written down once and then kept to.

Three things to do this month

1

Write the method down

One page: sources, denominator, window, exclusions. This is the fix that makesevery subsequent number comparable, and it costs an hour.

Owner _____ By When _____

2

Measure your ghost booking rate

Four weeks of booked hours against sensed-occupied hours. It is the fastest route torecovered capacity that does not involve a property decision.

Owner _____ By When _____

3

Break your headline figure down by floor

Then go and find out why the outlier behaves differently. That answer is usually themost actionable thing in your whole dataset.

Owner _____ By When _____

metrikus

Difficult to answer these? Let's talk.

Metrikus connects the building data you already have — sensors, booking systems, access control, BMS and environmental monitors —into one view, so questions like these take minutes rather than days.

Book a 30-minute walkthrough →



Metrikus is a building intelligence platform for workplace and real estate teams.metrikus.io