

FORECASTING

Forecasting you can verify.

Three tiers from statistical baseline to AI-corrected. Every forecast is committed by a human, reasoned through by the AI when enabled, and measured monthly against actuals. No black box. No promises that can't be checked.

THE TRUST PROBLEM

Most hotel forecasts are unverifiable.

Two options dominate the independent-hotel market: a manually built forecast in Excel, anchored to gut feel and last year's spreadsheet, or an AI forecasting product that promises accuracy but doesn't show its work.

The manual forecast at least has a human in the loop — but it lives in a workbook nobody else opens, gets stale within a week, and gets re-built from scratch every quarter. The AI forecast is a black box. The hotel team can't see why a number changed, can't override cleanly, and can't answer the owner's question "how accurate was the last forecast" with anything more specific than "we think it's pretty good."

A forecast you can't measure is a guess
in a nicer-looking spreadsheet.

THE PEAQPLUS APPROACH

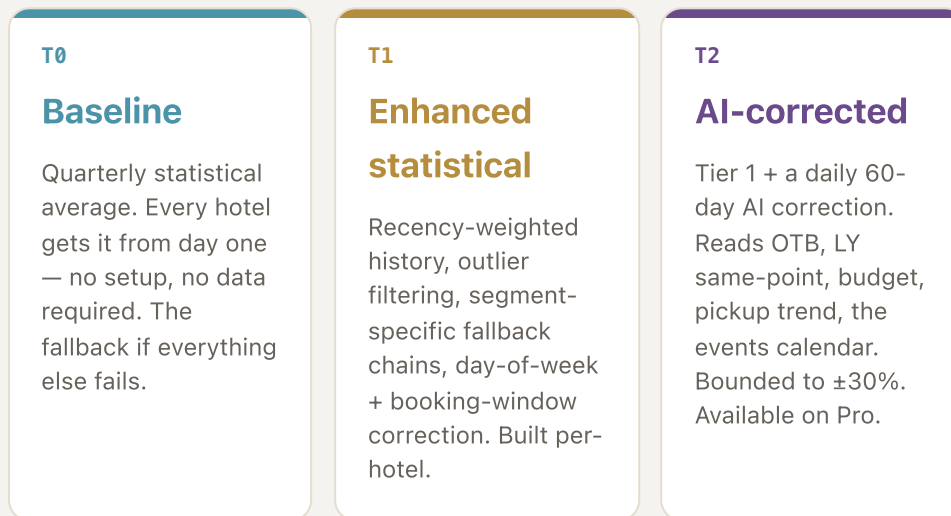
Human commits. AI assists. Every forecast measured.

Peaqlus forecasting is a three-tier ladder. The lowest tier is the statistical baseline — every hotel gets it, no setup required. The middle tier is the enhanced statistical layer, built per-hotel from recency-weighted history. The top tier adds an AI correction with a per-day reasoning trail. All three are measured monthly against actuals, on the hotel's own data, with the accuracy report published in a fixed format.

THREE TIERS

A ladder, not a magic box.

Every Peaqlus customer starts at the bottom tier and grows up the ladder as the data and the team mature. Each tier is its own forecast, visible alongside the others, comparable on the same accuracy axes.



WHAT'S THE SAME ACROSS ALL THREE

Each tier produces a daily 60-day pickup forecast. The Pricing Calendar reads the forecast natively, so the rate-setting decision sees the expected occupancy alongside the competitive rate. The Pricing Engine, the Revenue Meeting, and the Daily Briefing all read the same forecast — no version drift, no "which forecast are we looking at" moments.

WHAT CHANGES BETWEEN TIERS

Forecast accuracy. T1 typically lands at 10-18% MAPE on 14-day horizons after the first 90 days of data history. T2 typically improves on T1 by a further 3-6 percentage points. The accuracy is measured, not claimed.

THE DIFFERENTIATOR

We measure the forecast. Every month, on your data.

A monthly accuracy report lands in your inbox. Three error metrics (MAPE, SMAPE, MAE) across three horizons (7-day, 14-day, 30-day). Raw T1 vs AI-corrected T2 side by side. The actual numbers, on your actual data — not industry benchmarks, not a one-time validation study.

01 – MAPE / SMAPE / MAE

Three error metrics, three horizons

MAPE for the relative error (percent), SMAPE for sparse-data safety, MAE for absolute room nights. Each computed across 7-, 14- and 30-day forecast horizons. The full table on one page — auditable, comparable, defensible.

02 – RAW vs AI, SIDE BY SIDE

You see the AI's contribution, monthly

The same accuracy table shows T1 (raw statistical) and T2 (AI-corrected) for every horizon. If the AI isn't earning its keep on your data, the report shows it. Most months, AI adds 3-6 MAPE points — but only your data answers for your hotel.

03 – PER-DAY AI REASONING

The AI explains every move

For any specific day, the Forecast Advisor shows the AI's reasoning: which inputs moved the forecast, by how much, why. "Friday Aug 22: +28%, because national holiday + LY same-day pickup +18 + neighbouring source-market peak." Auditable in your language.

04 – COMMITTED VERSIONS, NEVER REWRITTEN

The past gets a permanent record

Every committed forecast is preserved as an immutable version with a timestamp, user, and reason. The June forecast you committed in May stays exactly as it was — even three years later. Time Machine replays it, the accuracy report compares it. The forecast is part of the audit trail.

Most AI tools promise accuracy.
Peaqlus measures it — and publishes it.

PRICING

Statistical tiers free. AI tier on Pro.

Per-property monthly pricing. Tier 0 (Baseline) and Tier 1 (Enhanced Statistical) are included in every bundle — even Starter. Tier 2 (AI-corrected) ships with Pulse AI, which is included in the Pro bundle and available à la carte alongside Starter.

Starter

€109

/ property / mo

T0 + T1 forecasting included.

- + Tier 0 (Baseline)
- + Tier 1 (Enhanced Statistical)
- + Monthly accuracy report
- + Committed-version audit trail
- + BI Core, Insight Engine, Decisions

Pro

RECOMMENDED

€334

/ property / mo

All three tiers, full reasoning.

- + Everything in Starter, plus:
- + Tier 2 (AI-corrected)
- + Forecast Advisor (per-day reasoning)
- + Pulse Chat & AI Report Narrative
- + Sales Pipeline + Revenue Meeting

Starter + Pulse AI

€109 + €49

/ property / mo

Add T2 to Starter standalone.

- + Starter, plus:
- + Tier 2 (AI-corrected)
- + Forecast Advisor
- + Pulse Chat & Daily Briefing
- + AI Report Narrative

What the AI tier doesn't do

- → **It doesn't set rates.** The forecast feeds the Pricing module; the Pricing module sets rates. Two distinct decisions, in distinct modules.
- → **It doesn't promise infinite accuracy.** AI corrections stay within sensible bounds around the statistical baseline. If the model tries to propose a wild adjustment, the platform stops it. The measurement is the truth-teller, not the model's confidence.
- → **It doesn't replace human judgment.** Every forecast is committed by a person, with an attached reason if it differs from the AI suggestion. The AI proposes; the human commits.

See the accuracy report on your own data.

A 45-minute walkthrough. We start on our live demo property, then load a CSV of your last 90 days if you bring it, and run the three-tier comparison live. By the end, you've seen the monthly accuracy report on real numbers — yours or ours.

sales@peaqplus.com