



eviivo's Inherent Dynamic AI Pricing:
Architecture and Operating Principles



1. Key terms used in this document

- **Agentic harness:** An AI system that orchestrates an entire process autonomously, calling on specialist tools as needed
- **Competitive set ('Comp set'):** A custom-defined group of nearby or comparable properties that are tracked for competitive pricing context
- **Frontier model:** The most capable AI model available at any given time. This is regularly evaluated and swapped
- **Gradient Boosted Model (GBM):** A machine learning algorithm suited to structured data like booking and price information
- **Guardrails:** Any pricing limits configured by a property that the AI cannot breach
- **Inference:** The moment the AI generates a live pricing decision (as opposed to training)
- **Large / Small Language Model (LLM / SLM):** AI models that understand and generate natural language
- **Machine learning (ML):** Algorithms (mathematical rules and statistical methods) through which AI learns to identify patterns in order to make accurate predictions on new data
- **SHAP (SHapely Additive exPlanations):** A graphical method for showing how much each individual factor in a machine learning model is likely to push a predicted price up or down

2. Summary of eviivo's Inherent Dynamic AI Pricing

eviivo's **Inherent Dynamic AI Pricing** is powered by an agentic AI system. It combines traditional machine learning (for demand forecasting and price optimisation) with frontier large language models (LLMs) and other generative AI techniques to deliver contextual reasoning, strategy interpretation, and explainability.

The system operates as an **agentic harness** rather than a linear pipeline. This means that AI orchestrates the entire process, from deciding what data to retrieve (including live internet searches mid-flow) through to generating both the final rate and the natural language explanation of why it arrived at that rate. Within this harness, specialised components handle demand forecasting, goal interpretation, constraint validation, and more, **but the AI is present and reasoning at every stage** – not just waiting for a handoff.

This architecture is designed to evolve. Individual components can be updated or replaced independently as the AI landscape changes, and **the system continuously learns** from manual overrides, user feedback, and new market signals.

3. Architecture

eviivo's Inherent Dynamic AI Pricing can best be described as an agentic harness in which AI is present throughout, orchestrating, reasoning and intervening at every stage. The components below represent key stages of the process:

Data retrieval

The system ingests a rich and varied data payload spanning:

- **A property's historic performance**, based on unrestricted internal data held in the eviivo PMS
- General **external market feeds** and online **demand signals**
- **Competitor intelligence**, both generic (i.e., based on similar property types in the location) as well as a custom-defined comp set
- **Full amenities details**, as well as the property's review scores and the **reviews** collected in eviivo Suite (from both direct and OTA sources)
- **Weather forecasts**
- **Event calendars**, including eviivo's pricing calendar and additional local events provided by users

Critically, the agent can also search the internet live at inference time, meaning **the data sources available to any given pricing decision are not limited to a fixed list**. The system decides in the moment what context it needs – and takes action to get it.

Demand forecasting

Trained ML models and demand modelling techniques forecast demand and calculate candidate prices. Each night within a rolling 90-day forward window is priced individually for each rate plan per stay scenario, with contextual awareness of the surrounding 7-day and 60-day windows. Analysis is based on bespoke gradient boosted models (GBMs).

Contextual reasoning (generative AI)

Frontier large language models and other generative AI techniques interpret:

1. The commercial goals that a property or group has configured in eviivo Suite
2. The conditions attached to the above
3. Any injected user feedback

This layer reasons about the ML output in the context of the property's strategy and current conditions, translating high-level goals like "Maximise RevPAR once occupancy exceeds 70%" into concrete pricing behaviour.

Post-processing constraints

A constraint layer **validates that the proposed rates do not breach configured guardrails** except where conditional rules explicitly permit it. No rate is published that violates the property's hard limits.

Output and explainability

The final rate is returned to eviivo Suite ready for instant publication across all connected channels, along with **two layers of explanation** displayed on-screen:

1. SHAP values (from the ML side) provide **quantitative feature contributions**, breaking down exactly how much each factor influenced the price.
2. The agentic pipeline generates **a natural language explanation** that summarises the SHAP values.

Both layers are produced as part of the decision process itself, not retrospectively. This ensures that what the AI decided and what it tells you are always consistent.

If goal interpretation encounters an edge case, the system defaults to optimising for revenue growth, ensuring commercially sensible behaviour at all times.

4. Data Sources

The system draws on a wide range of data, combining eviivo's own PMS data with competitor intelligence, market signals, environmental context and live web research. Together, these sources give the AI a comprehensive, real-time view of the factors that should influence any given pricing decision.

Overall, eviivo's Dynamic AI Pricing engine has access to and can consume over 100 billion data points.



Source	What it includes
eviivo Suite (full PMS history)	Property details: Type, location, facilities, amenities, accessibility, policies Accommodation category: Capacity, size, bathrooms, amenities, rack rate, adult/child supplements, minimum/maximum price guardrails Pricing strategy: Rate plan model (flat, occupancy pricing, adult/child pricing), nightly prices by day of week or season, taxes, board basis, min/max stay restrictions, promos, pricing rules Policies: Cancellation and deposit policies, check-in/out day rules, open OTA channels Real-time booking data: Room night price, taxes, currency and any additional guest fees and charges. Restriction changes read live as they happen Inventory blocks: Owner blocks, cleaning, allocations Google Analytics demand signals: Search result sessions and page sessions per property per day Guest reviews: Sourced from own website and multiple OTAs. Includes overall score and sub-category scores (e.g., for cleanliness, value, location, host interactions, expectations, amenities and comfort) as well as the full textual detail of reviews OTA listing identifiers and commission percentages
General market data	Market data from reputable, publicly available sources cover: 1. Properties of a similar type in the area 2. A custom-defined competitor set For both, eviivo's AI pricing engine captures property type, location, facilities, room configuration, capacity, amenities, review scores (overall and sub-scores), host status, hosting duration, and location. The system also captures rates and prices by check-in date, including pre- and post-discount rates, tax amounts, cleaning fees, other fees, cancellation policies and fees, per-night-per-guest pricing and availability. Market-level booking feeds with ADR, room nights and revenue are aggregated by day, month and booking window, and segmented by channel, property type and geolocation.
Weather	City-level daily forecasts by geolocation, covering min/max temperature, min/max wind speed, rainfall, snowfall, total precipitation, cloud cover and relative humidity.
Events	Local events and seasonal periods sourced independently, in addition to events provided in eviivo Suite's pricing calendar (each with its own start and end dates, period names and event types) and a custom-defined list of local events that identifies each event's expected booking impact.
Dynamic web data (live at inference)	The LLM within the agentic harness can search the internet during processing. Data sources are therefore not limited to the above: the AI decides in real time whether additional context would improve a pricing decision and, if so, retrieves it.

5. How the Pricing Works


Optimisation window

This is based on a rolling 90-day forward basis, not per calendar month. Each night is priced individually per rate plan, with contextual awareness of 7-day and 60-day surrounding windows.

Goal interpretation

eviivo's AI settings support custom-defined pricing strategies with constraints and goals set at multiple levels. Goals include maximisation of RevPAR, Occupancy or Price, and these cascade down from the group portfolio level to the property and accommodation category levels (unless explicitly overridden).

This hierarchy is fully granular. For example, a boutique hotel group might decide to maximise occupancy at the portfolio level, but override this to maximise RevPAR for a flagship property, then override again at the category level so premium suites chase rate while standard rooms fill the house, with all goal configurations running simultaneously. eviivo's AI interprets these layered goals alongside conditional logic (such as occupancy thresholds) and translates them into pricing behaviour.



The screenshot shows the 'AI Settings' interface. At the top, there's a toggle switch labeled 'Active' which is turned on. Below this, there are three tabs: 'Group', 'Properties' (which is selected), and 'Categories'. Under the 'Properties' tab, there's a dropdown menu 'Select your properties' with '(6 selected)' next to it. Below that, there are three goal options: 'Maximize RevPAR' (selected), 'Maximize occupancy', and 'Maximize price'. Under the 'Pricing flexibility' section, there's a question 'Enable discounts below minimum price?' with 'Yes' and 'No' buttons. Below this, there's a section 'Allow the minimum to drop to' with a value of '\$100'. There are two checkboxes: 'Only if my occupancy level is still less than 50 %' (checked) and 'and there are more than 30 days left before check-in.' (checked). At the bottom, there's a section 'Minimum stays' with a question 'Allow minimum stay overrides?' and 'Yes' and 'No' buttons. There are 'Back' and 'Save' buttons at the bottom of the interface.

Learning from overrides

Manual overrides are treated as features for the next pricing cycle. eviivo's AI learns from both the fact that a price was changed and what it was changed to, incorporating this signal into subsequent recommendations.

User feedback

Natural language feedback can be injected into the AI's reasoning. Feedback is ingested daily, provided it is current and less than three days old. During initial trials, feedback was injected after an unexpected technical incident made some rooms unavailable, and AI Pricing automatically adjusted rates to a more aggressive strategy to make up any deficit before month end.

Algorithms and models

- **Machine learning:** Trained ML models and demand modelling techniques. Gradient boosted models (GBMs) make up the primary approach.
- **Generative AI (LLM):** Frontier LLMs are used for contextual reasoning, goal interpretation, and explainability. Models are regularly evaluated and switched as using a single model would lower accuracy.
- **Generative AI (non-LLM):** Other generative AI techniques in use alongside LLMs.
- **Explainability:** A SHAP graph is used to provide an easy visualisation of the quantitative contribution of each group of pricing factors.
- **Experimental:** Small language models (SLMs) are used where relevant and specific as a way to improve performance.

6. Explainability

Every pricing decision is accompanied by a full explanation, generated as part of the decision process itself. This operates at two levels:

1. **At the quantitative level,** SHAP values from the ML models show exactly how much each feature contributed to the recommended price. These drive the factor visualisations in the product, grouped into categories including:
 - Accommodation type and capacity
 - Booking history and occupancy
 - Booking window and pickup rate
 - Calendar and availability
 - Competitor data
 - Events
 - Guest reviews
 - Market and location
 - Policies and restrictions
 - Pricing history, seasonality and weekends
 - Weather
2. **At the reasoning level,** the agentic pipeline generates a natural language explanation of the decision, drawing on the same context it used to arrive at the price. The specific LLM producing the explanation may change over time as models are evaluated and swapped, but the explanation always comes from within the same decision-making process, not from a separate system looking at the output after the fact. This feature can be expected to change the most over time as LLMs evolve and become increasingly powerful.

7. Metrics and KPIs

The following core KPIs are aggregated by accommodation category and used by eviivo's Inherent Dynamic AI Pricing:

- **Available inventory:** Total inventory minus blocked inventory
- **Net occupancy %:** Inventory sold / available inventory × 100
- **Average Daily Rate (ADR):** Net room revenues / sold inventory
- **RevPAR:** Net room night revenues / available inventory (equivalently: occupancy % × ADR)
- **Pace % (inventory):** Inventory sold at point in time / available inventory at period end × 100
- **Pace % (revenue):** Room revenue sold at point in time / room revenue at period end × 100
- **Pickup:** New inventory or revenue sold at a point in time, minus cancellations at same point in time
- **Pickup rate:** Pickup / days in period

Forecasting uses both a simple method (stays-to-date scaled by last year's ratio) and a more advanced method that uses forward bookings already on the books (also scaled by last year's ratio), while accounting for the potential impact of future cancellations.

8. A Case Study of eviivo's Inherent Dynamic AI Pricing

A 3-star London hotel introduced **eviivo's Inherent Dynamic AI pricing** solution for the first time at the end of June 2026. A year earlier, the hotel had achieved **83% occupancy** in July, with a RevPAR of £128 and revenues of about £91,000.

The main goal set by the hotelier was unequivocal: **maximise RevPAR**.


Beating an already very reasonable occupancy level of 83% was a challenge, especially since:

- All major KPIs were behind at the beginning of July
- The hotel lost two substantial group reservations on July 8, just days before the cancellation fees would have kicked in

The following chart shows how AI reacted instantly to this hotel's specific context. It plots the hotel's monthly position in July, day by day, and compares it to the hotel's performance at the same point in time a year earlier. All numbers are expressed as a **% year-on-year improvement**.

KPI Trendline – YoY Growth

Tracking RevPAR, Revenue, Occupancy, and ADR growth daily
July 2026 compared to July 2025



From July 1, the hotel was moved to AI Pricing on full AI auto-pilot, starting from a low point of occupancy (in comparison to the previous year). Therefore, AI initially dropped the price until occupancy levels and revenues started to respond positively.

On July 8, the hotel received a large, unexpected cancellation and lost substantial revenues as shown on the chart. eviivo's AI Pricing adopted a more aggressive pricing approach in order to make up any lost ground.

An opaque web-exclusive promo and a close-out rule to temporarily shut down high-commission channels were activated towards month-end to encourage direct business and reduce costs.

The results:

- **RevPAR grew by 8.4%** from £128 to £139
- **Revenues grew by nearly 7%** to £97,300
- **Occupancy improved by 8.3%** as it increased from 83% to 90%
- **ADRs improved by 0.2%** to £154.50.

Summer 2026 saw London occupancies range between 81% and 88% in July and within 75% and 81% in August. This 3-star hotel – on full AI autopilot – outperformed the market at 90% for both months, delivering 8.3% growth in July and 20% in August.

Anecdotally, at some point during the trial, AI did what no human expected: for a few days, it made the breakfast-inclusive price cheaper than the non-refundable room-only rate. AI had simply worked out that flexibility seemed to command a much higher premium than breakfast.

9. Operating Principles

Agentic, not monolithic. eviivo's Inherent Dynamic AI Pricing operates as an orchestrating harness across the entire process, not as one step in a sequential handoff. The AI model reasons about data retrieval, forecasting, goal interpretation, and constraint validation as a unified decision, not a chain of isolated tasks.

LLM-flexible. Frontier models are used but regularly swapped as the market develops. SLMs are also used to optimise cost and performance. This versatility ensures that eviivo's Inherent Dynamic AI Pricing can always benefit from the latest LLM developments.

Goal-directed with a revenue safety net. The system pursues user-defined goals and strategies but defaults to revenue growth optimisation if anything unexpected occurs.

Continuously learning. Manual overrides, user feedback and new data all feed back into the system's reasoning on subsequent pricing cycles. Every human intervention makes it smarter.

Explainable by design. The same agentic pipeline that makes pricing decisions generates the explanations (both the quantitative SHAP-based factor contributions and the natural language reasoning), ensuring that what the AI decided and what it tells you are always aligned.

Contextually aware in real time. The system does not rely solely on a fixed dataset. At inference time, AI can search the web, pull in new context, and factor information it has not previously seen into a pricing decision. The data sources that influence any given rate are bound only by what the AI model judges to be relevant.

10. Ethics and AI Safety

- **Personal information:** The model is built to ensure that no personal information is ever consumed or disclosed by eviivo's AI Pricing model and the GDPR guidelines are observed.
- **Credit card and payment information:** No credit card information is held within eviivo Suite as all cards are encrypted and tokenised under PCI DSS.
- **Market data:** External market and competitor data is obtained ethically from reputable, publicly available data sources, and the AI model does not use confidential guest data or payment information. The platform does not operate as a data-sharing hub, thereby preventing any risk of collusion.
- **Security:**
 - eviivo undergoes **rigorous annual independent audits for both NF525 and PCI compliance**. This involves on-site audits, taking place over weeks, and a series of penetration tests.
 - eviivo uses a **multi-layered security infrastructure**, running its own firewall, its own intrusion detection, and virus/trojan and ransomware protection, as well as rate limits and a takedown service to address phishing.
 - eviivo operates a **disaster recovery service** using a backup data centre, which is further protected by a full moat service to prevent DDOS attacks.
 - eviivo's own internal systems and development environments are entirely segregated from the production platform, and eviivo has implemented **multi-factor authentication** across the board.
- **AI security:** Generative AI LLMs and ML techniques used by eviivo's Inherent Dynamic AI Pricing read and transfer pricing data or amenities data from and into eviivo Suite via a secure API subject to full authentication, plus the security precautions discussed above. This ensures that the LLM and ML models themselves have no direct access to the production platform. Any AI processing is done in a segregated environment, and data consumed via webhooks is passed back via a secure API post-processing. The API does not support access to personal data, card information, or password or user profile information.
- **No collusion.** While eviivo's Inherent Dynamic AI Pricing model accesses multiple data sources to determine the optimal price, it does not disclose or share specific competitor prices, and they are only one of a multitude of factors taken into account. **eviivo does not provide or act as a pricing hub**. Each customer's pricing is optimised independently, and no customer's rates are visible to, or used to set, another customer's rates.

11. How AI Pricing Tools Differ

There are two broad dimensions to think about when evaluating how a pricing tool works and how well it will fit with your business:

1. **Operating model:** The way you engage with the tool
2. **AI type:** The type of intelligence that steers the tool

Three operating models

Deep configuration

You design the pricing strategy through detailed configuration settings, including but not limited to: base prices, minimum and maximum rates, early bird discounts, last-minute price bumps, gap rules, seasonal profiles, and day-of-week adjustments.

Pricing tools that use deep configuration and ML algorithms typically forecast demand and calculate rates. However, results depend heavily upon how well the tool is configured and set up, the extent to which it can be adapted to your specific needs, and how regularly it is maintained.

These tools reward expertise and time investment. They deliver their best results when someone is actively managing them.

Guided automation

Providing similar forecasting capabilities, but with a more hands-off approach, guided automation tools are designed to run with minimal configuration. You set a base price, set your limits, select your goal, and let the system handle the rest.

Several tools in this category offer ways to tweak the algorithm's direction, such as aggressiveness dials (enabling you to control how bold the algorithm is when making pricing decisions), strategy presets, and combinations of the two, allowing for a deeper level of control and customisation.

PMS-native pricing

While PMS pricing tools typically include elements of the above, they warrant their own category as they each provide unique offerings and integration advantages.

Two AI approaches

Hospitality pricing tools use ML techniques to forecast and predict things like future demand and price sensitivity. Where they differ is in how each tool turns those insights into rates, with two main options available.

Predictive ML

A predictive ML algorithm takes historical data, market signals, and parameters you have set, then calculates the price that best fits those inputs. This is the way the vast majority of pricing tools work, and the best implementations can yield exceptional results. When you want to change the tool's behaviour, you change a setting, often by adjusting a dial or rule. However, these tools can be time-consuming to configure for the best results and to manage day-to-day.

Inherent Dynamic AI Pricing

With the rise of AI performance, an emerging approach consists of building an entire ML model directly within a generative AI structure. Generative AI is a type of ML model that generates new content (such as text or images) to a level where it behaves comparably to a human. Popular generative AI tools include Claude, ChatGPT and Gemini.

This approach is able to **create bespoke pricing strategies based on unique simple goal settings**. It allows AI to consume over 100 billion data points and act more dynamically (i.e., with more "common sense") when it comes to strategy changes based on performance or the most current context, such as an unexpected large cancellation event mid-month.

This **pairs effectively with natural language AI explainability**, which, while used widely across various vendors, works significantly better when the underlying decisions are made by the same AI. For many providers, the underlying tool is usually built on their own ML algorithms.

12. Five questions to ask if you're considering price automation

1. Who manages pricing day to day?

If you have a revenue manager who enjoys building comp sets, tweaking prices, updating extranets and testing price strategies, a deep configuration tool will reward that effort.

If your main focus is to increase your reputation and direct business, consider automation rather than manual revenue management. Letting AI do the groundwork enables you to focus on price analytics, monitoring AI, setting group strategies and boundaries for negotiated corporate rates, or working on the digital positioning of your products. Price matters, but it is only one component of the mix.

If pricing is simply one of twenty things you handle, you should look for a tool that executes rather than advises. Good tools let you keep your hand on the brakes while eliminating the most tedious pricing chores.

2. Can the tool enact different strategies for different parts of your business?

A single pricing objective applied uniformly may be fine for a five-room guesthouse. However, you may require more flexibility if you run multiple properties or a hybrid portfolio and have disparate accommodation categories across different locations and with different commercial roles.

In that instance, check whether:

- The tool lets you vary your goals and strategies by segment (i.e., at group, property, and category levels), under a single login and from a single calendar
- Those variations inherit sensible defaults so you are not configuring everything from scratch

3. How rigid are the guardrails?

Fixed minimum prices protect your rate but cannot respond to changing conditions. Some tools let the AI flex its own guardrails when specific conditions you define are met, like low occupancy, a long booking window, or calendar gaps. Having that conditional flexibility can recover revenue that a hard floor would leave on the table.

4. Where do the rates come from?

Rates either originate in a third-party tool and get pushed into your PMS, or they are generated inside the PMS. A pricing engine inside your PMS removes a sync step, a second login, and an additional support relationship. A standalone tool may offer more configuration depth.

Both options are valid; the question is which trade-off suits you.

5. How does the cost scale?

Percentage-of-revenue fees offer the strongest risk/reward approach: they rise every time performance improves and lower when business drops. For example, if pricing costs you 1% of revenue but delivers a 5% improvement, you win.

Per-listing or per-room fees stay flat and also support volume discounts to help you scale and improve your unit metrics as you grow.

For a growing portfolio, the cost difference can compound significantly over a few years.

Choosing what fits

The best AI pricing tool is the one that matches how you actually want to work.

If you want maximum configurability and enjoy pricing work while having the expertise to use it, then a deep configuration tool will serve you well.


If you want pricing that runs itself with sensible direction-setting, several excellent guided automation options exist.

eviivo's Inherent Dynamic AI Pricing was built for you if you want to:

- Define differentiated strategies for different parts of your business
- Set conditional rules for when AI may flex its own guardrails
- Have it all executed within the same platform you use to manage everything else – secure in the knowledge that you can easily monitor results and put your hand on the brakes at any time

Ready to try eviivo's Inherent Dynamic AI Pricing?
Book your free demo with one of our experienced consultants below.

BOOK FREE DEMO



eviivo

FOR MORE INFORMATION, CONTACT

SALES@EVIIVO.COM | +44 (0)20 3103 3500