PERTLINK VIEWPOINT · THE METER SERIES · PLAYBOOK

Tokenomics *in Hospitality*

A playbook for **budgeting**, **reading** and **governing** the AI meter, for owners, general managers, the C-suite, Chief AI Officers and the vendors who serve them.



TERENCE RONSON

Founder and Managing Director, Pertlink Limited
Chair, HFTP AI Collective's AI Education & Training Subcommittee
Founding member, PAIBA, the Philippine AI Business Association

FIRST EDITION · PUBLISHED 5 NOVEMBER 2026

CONTENTS

Inside this *playbook*

How to use this playbook 3

Introduction: The meter is running 4

01 Token literacy: every word runs the meter 5

02 The meter starts at the look 9

03 Budget the meter, not the machine 13

04 Read the meter 16

05 Govern the meter 19

06 Role playbooks 22

- Owners and investors 22
- General managers and operators 23
- The C-suite and the CFO 24
- Chief AI Officers 25
- Vendors and technology partners 26

Formula card and TCPG worksheet 27

The 2027 budget checklist 28

Glossary and sources 29

WHY THIS PLAYBOOK

On 18 October, Pertlink sent an AI agent to book a one-night stay on the websites of the world's ten largest hotel groups. Every page it read ran a meter, and someone will be invoiced for it. The meter runs inside the hotel too: in the concierge chat that re-reads the whole stay, in the review reply a manager rewrites, and in the AI hidden in a vendor fee. Few hotels can yet say what any of it costs per guest, and 2027 budgets are being set now.

I compiled this playbook so that owners, GMs, finance teams, Chief AI Officers, and vendors can read the same meter, in hotel language.

It is not a commercial publication. Nobody sponsored it, and it is vendor- and brand-agnostic. It recommends no product, platform, or model, and it names companies only as sources or public examples. The decisions stay with you. It's here for awareness and education.

Read it with one aim: to walk into your 2027 budget meeting able to say what AI costs per guest. If nobody on your team can explain a token yet, that's fine, because chapter 1 starts there. It's free, so please share it.

Terence Ronson

WHAT YOU'LL TAKE AWAY

- A number to manage: TCPG and TCPB, worked through step by step
- A shared language for everyone at the table, from owner to vendor
- A budget method built for usage pricing: rolling forecasts and written triggers
- The questions to ask vendors and AI agent platforms before you sign
- A checklist to take into the 2027 budget meeting

HOW TO USE THIS PLAYBOOK

Start where *your job* is

Hotels make AI decisions at five different tables: the owner's, the GM's, the finance director's, the Chief AI Officer's and the vendor's. Each needs the same arithmetic and a different first page.

Chapter	Page	Owners and investors	GMs and operators	C-suite and CFO	CAIO	Vendors
Introduction	4	●	●	●	●	●
01 Token literacy	5	○	●	○	●	●
02 The meter starts at the look	9	●	○	●	●	●
03 Budget the meter	13	●	●	●	○	○
04 Read the meter	16	○	●	●	●	●
05 Govern the meter	19	●	●	○	●	○
06 Your role playbook	22	●	●	●	●	●
Formula card and worksheet	27	○	○	●	●	●
2027 budget checklist	28	●	●	●	○	○

● Essential ○ Useful background

If you only have ten minutes

OWNERS
Read the introduction, then your playbook on page 22 and the checklist on page 28.

GMs
Read the token literacy table on page 6, the five-night example on page 7, then page 23.

C-SUITE
Read the rolling forecast on page 13 and why the bill jumps on page 17, then page 24.

CAIO
Read right-sizing the model on page 8 and the Decision Envelope on page 19, then page 25.

VENDORS
Read what's on the bill on page 16, then page 26: it is what your clients will ask for.

INTRODUCTION

The meter *is running*

Heading into 2027, the hard part is no longer adopting AI. It is measuring it, governing it and paying for it responsibly.

For decades, hotel technology behaved like furniture. You bought it, renewed it once a year, and the forecast held. AI doesn't behave like that. Every guest message it answers, every review it drafts and every rate an agent checks runs a meter. The more it is used, the more it costs.

That is tokenomics: the economics of the token, the unit AI providers bill in. This playbook is about reading that meter at hotel level, per guest and per booking, and deciding who pays for it.

Why now

1. **Pricing moved to the meter.** Early in 2026, many frontier AI companies moved to usage pricing where, in McKinsey's Wasim Lala's words, "the more you eat, the more you have to pay." The same amount of AI could cost two to three times as much, instantly. Expect hotel software to follow.
2. **The meter moved forward.** Guests increasingly begin their trip inside AI assistants that shortlist hotels before anyone clicks. Those agents run a meter on every look, and someone will be invoiced for it.
3. **The rules are being written.** In June, the Linux Foundation announced its intent to launch a Tokenomics Foundation, working with the FinOps Foundation on open standards for measuring AI token spend. Booking.com is among its first supporters. No hotel group is.

> *"AI doesn't come with a subscription fee. It comes with a meter. If you can't tell your owner what AI costs per guest, you don't have an AI budget. You have a hope."*
>
> TERENCE RONSON, SHIJI INSIGHTS, OCTOBER 2026

The good news is that AI cost can be controlled. "With the right visibility, it's solvable," says McKinsey's Kaavini Takkar. This playbook is that visibility, in five steps and five role pages.

01	02	03	04	05
Learn it	**Find where it starts**	**Budget it**	**Read it**	**Govern it**
What a token is and why some cost more	At the look, before the guest says hello	Rolling forecasts, TCPG, three lines	Where spend hides, why bills jump	Control plane, envelopes, override rates

01

TOKEN LITERACY

Every word *runs the meter*

What a token is, why some cost more than others, and why the words nobody reads are still on the bill.

ESSENTIAL FOR OWNERS GMS C-SUITE CAIO VENDORS

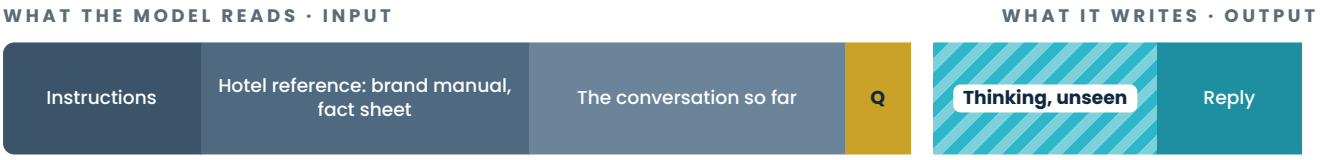
What you are paying for

Hotels taught a generation of managers to read a utility bill: kilowatt-hours, peak tariffs, the standing charge. AI needs the same literacy. Call it token literacy.

"A token is the smallest unit of information a model sees," says McKinsey's Wasim Lala. His colleague Kaavini Takkar puts it in billing terms: a token is "a unit of measure that AI providers use today to bill their clients." A short greeting is about six tokens. Providers price them per million.

Three facts that matter more

- **Not all tokens cost the same.** What the AI writes, output, is "way more expensive" than what it reads, input.
- **Some tokens are invisible.** Reasoning models generate thinking tokens "that we, as users, don't get to see. But we're essentially paying for it."
- **The guest's question is the smallest part.** Everything sent with it is counted every time.



Q = the guest's question: "What time is breakfast?"

Output is priced higher per token

FIGURE 1.1 · Anatomy of one AI reply. Pertlink illustration: proportions are indicative, not measured. The guest's question is the smallest thing the model reads.

Read it like a utility bill

	THE ELECTRICITY BILL	THE AI BILL
The unit	Kilowatt-hour	Token, priced per million
The tariff	Peak and off-peak rates	Output dearer than input; bigger models dearer than small ones
Hidden use	Equipment left on overnight	Thinking tokens, re-read conversations, pilots nobody switched off
The useful number	Cost per occupied room	Cost per guest served: TCPG

Price lists speak in millions of tokens. Budgets should speak in guests. That is the idea behind Token Cost Per Guest, TCPG, which runs through every chapter of this playbook.

Quotations: Hämäläinen, Takkar and Lala, What does AI really cost?, McKinsey, 14 October 2026.

01 · TOKEN LITERACY

Token literacy in *eight ideas*

The clearest primer we have seen is McKinsey's explainer *What does AI really cost?* Here it is in its experts' own words, with what each idea means at a hotel.

THE IDEA	WHAT MCKINSEY'S EXPERTS SAY	WHAT IT MEANS AT A HOTEL
What a token is	*"A token is the smallest unit of information a model sees." (Lala) "At its core, it's a unit of measure that AI providers use today to bill their clients." (Takkar)*	Every word a guest types, and every word the hotel's AI reads or writes in reply, is counted.
How big a token is	*"Hi ChatGPT, how are you?" could be "about six tokens." (Takkar)*	"What time is breakfast?" costs almost nothing on its own. What matters is everything sent along with it.
How tokens are priced	*"Typically, billing happens per million tokens." (Takkar)*	Price lists speak in millions. Budgets should speak in guests: that is TCPG.
Input and output	*Input is "the prompt you're giving the system"; output is "the output the system is generating." "Output tokens are way more expensive than input tokens." (Takkar)*	A long, effusive reply to every review costs more than a short, warm one. Length is a cost decision.
Thinking tokens	*Reasoning models generate "intermediate thinking tokens… that we, as users, don't get to see. But we're essentially paying for it." (Takkar)*	A reasoning model pondering breakfast hours is paying for thought nobody needed.
Context windows	*"People are providing as much reference as possible to get the best outcomes." (Takkar)*	Sending the whole brand standards manual with every guest message is paid for every time.
Memory	*The models are "stateless": to remember something, "you need to feed everything back into it." (Hämäläinen)*	A guest conversation that runs all week is re-read, in full, every time the AI replies.
Caching	*Takkar warns of teams "forgetting how to cache and forgetting how to reduce my context window."*	The hotel fact sheet doesn't change between guests. Cached, it needn't be paid for at full price each time.

Quotations: Lari Hämäläinen, Kaavini Takkar and Wasim Lala, What does AI really cost?, McKinsey Explainers, 14 October 2026. Hotel column: Pertlink.

TEST YOUR TEAM

Can someone in each department explain every row of this table in their own words? If nobody can, that is where your AI budget should start. Token literacy belongs in every AI training programme, HFTP's *AI in Hospitality: Essentials* included.

01 · TOKEN LITERACY

Memory: the *hidden multiplier*

AI models don't remember. To keep a conversation going, the whole conversation is sent again with every reply, and paid for again.

Take a guest who chats with the hotel's AI concierge across a five-night stay, adding roughly 2,000 tokens of conversation a day. Because the model is stateless, each day's reply means re-reading the whole thread: 2,000 tokens on day one, 4,000 on day two, rising to 10,000 on day five. By checkout, the model has read 30,000 tokens to hold a 10,000-token conversation.

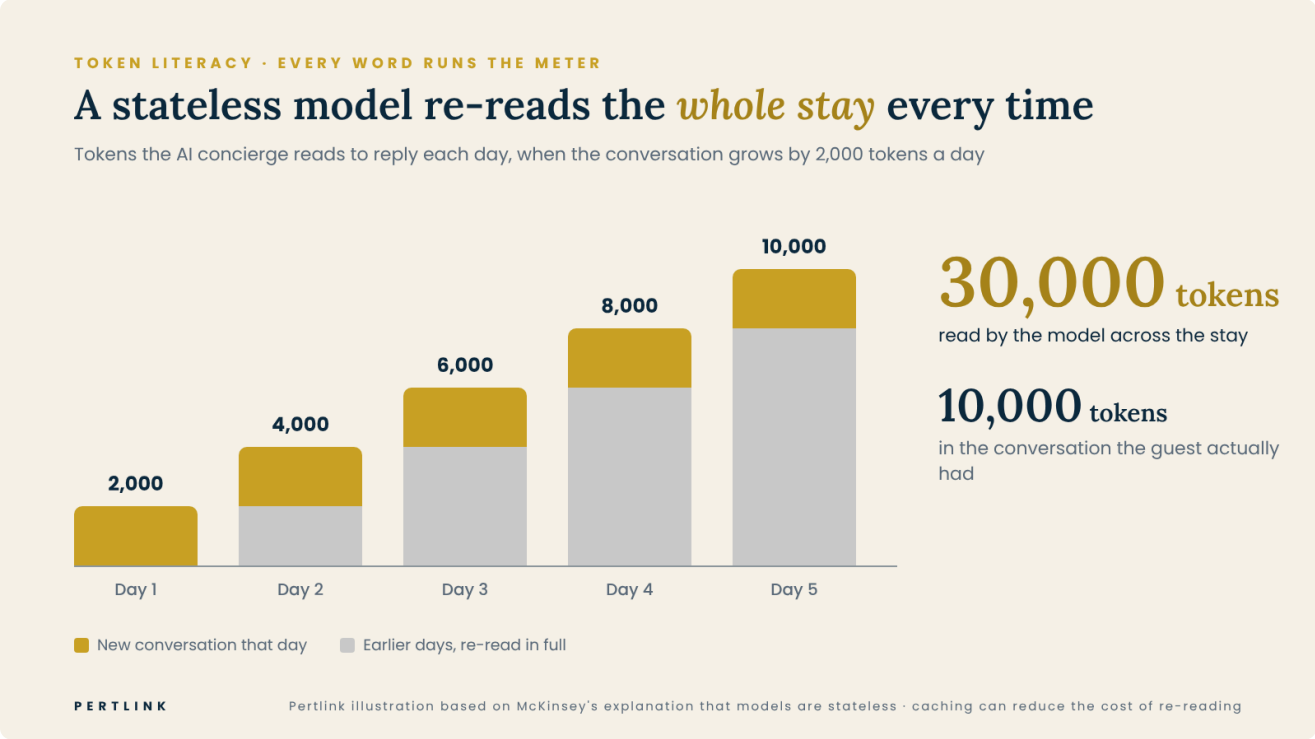


FIGURE 1.2 · Pertlink illustration, based on McKinsey's explanation that models are stateless. Caching can soften the cost of re-reading.

The longer the stay, the worse it gets

The arithmetic is simple. If a conversation grows by the same amount each day, the model reads that amount multiplied by 1 + 2 + 3 and so on, up to the number of days. Against the conversation itself, that is a ratio of (n + 1) ÷ 2.

Three times the words at five nights. Five and a half times at ten. Long stays, serviced apartments and resort guests who chat all week are where this shows first.

What to do about it

- **Summarise, don't re-read.** Carry a short summary of the stay forward instead of the full transcript.
- **Cache what doesn't change.** The fact sheet, policies and menus are the same for every guest.
- **Trim the context.** Send the relevant page of the brand manual, not the manual.
- **Ask your vendor** how its guest messaging handles all three. If it can't say, it probably doesn't.

The guest sees a friendly chat. The meter sees three times the words.

01 · TOKEN LITERACY

Right-size *the model*

Breakfast hours don't need the most powerful model in the building. Some questions don't need a model at all.

1	**No model** Rules, templates, look-ups	Breakfast hours from the fact sheet. Booking confirmations. Room status. Routing a towel request to housekeeping.	○○○
2	**Small or open-weight model** Fast, cheap, can run where you control it	Common guest questions. Sorting feedback by theme. Translating a short message.	●○○
3	**Frontier model** The capable general-purpose tier	A personalised offer. A delicate complaint. A review reply in the hotel's voice.	●●○
4	**Reasoning model** Thinks before it answers, and bills for it	Multi-step service recovery. Analysing a month of revenue data. Planning a complex group itinerary.	●●●

FIGURE 1.3 · The routing ladder. Pertlink illustration; tasks and relative costs are indicative.

Why it matters

McKinsey's Wasim Lala describes one client whose usage and spend doubled in two months as users moved from Claude Sonnet to the more advanced Opus. Nobody changed the price. People changed the default.

Takkar's question to companies is blunt: "Why didn't you route this particular request to an open-weight model?" Hämäläinen favours open-source models "where you can control what the model does and where it runs."

Five habits that shrink the bill

1. Send each task to the cheapest model that does it well, or to no model.
2. Cache what doesn't change between guests.
3. Summarise long conversations instead of re-reading them.
4. Keep replies as long as they need to be. Length is a cost decision.
5. Teach prompt discipline. Lala's test: "Am I teaching my employees how to write the prompts so I'm not overloading the context window?"

02

THE METER STARTS AT THE LOOK

The meter starts *before hello*

AI agents shop for guests before anyone says hello. Every look has a cost of sale, and it tends to land on the hotel.

ESSENTIAL FOR OWNERS GMS C-SUITE CAIO VENDORS

Guests increasingly begin their trip inside AI assistants that shortlist a handful of hotels before anyone clicks. “The booking used to be the battle,” Terence Ronson told Shiji Insights. “Now the battle is the context before the booking.”

That context now carries a meter. In Return on tokens: Who pays for the free travel agent?, Eugenia Dunaeva, an independent director at TravelAI who spent 11 years at Expedia, observes that most travel executives can recite their look-to-book ratio, yet hardly any can say how many tokens a booking takes.

Her worked example: an agent fields 1,000 requests and 100 become bookings that stick, earning $30 each. That is $3,000 in revenue. At 60 cents of model and tool cost per request, the token bill is $600, because it is paid on all 1,000 requests, not just the 100 that converted.

RETURN ON TOKENS · ROT

(value of output − cost of tokens) ÷ cost of tokens × 100

Packy McCormick with Markie Wagner, Not Boring, June 2026. Dunaeva’s example: ($3,000 − $600) ÷ $600 × 100 = 400%.

TOKEN COST PER BOOKING · TCPB

look and book tokens ÷ bookings that become stays

Pertlink. Dunaeva’s example: $600 ÷ 100 = **$6 a booking**, a fifth of what the agent earns, before payment, distribution and support take their share.

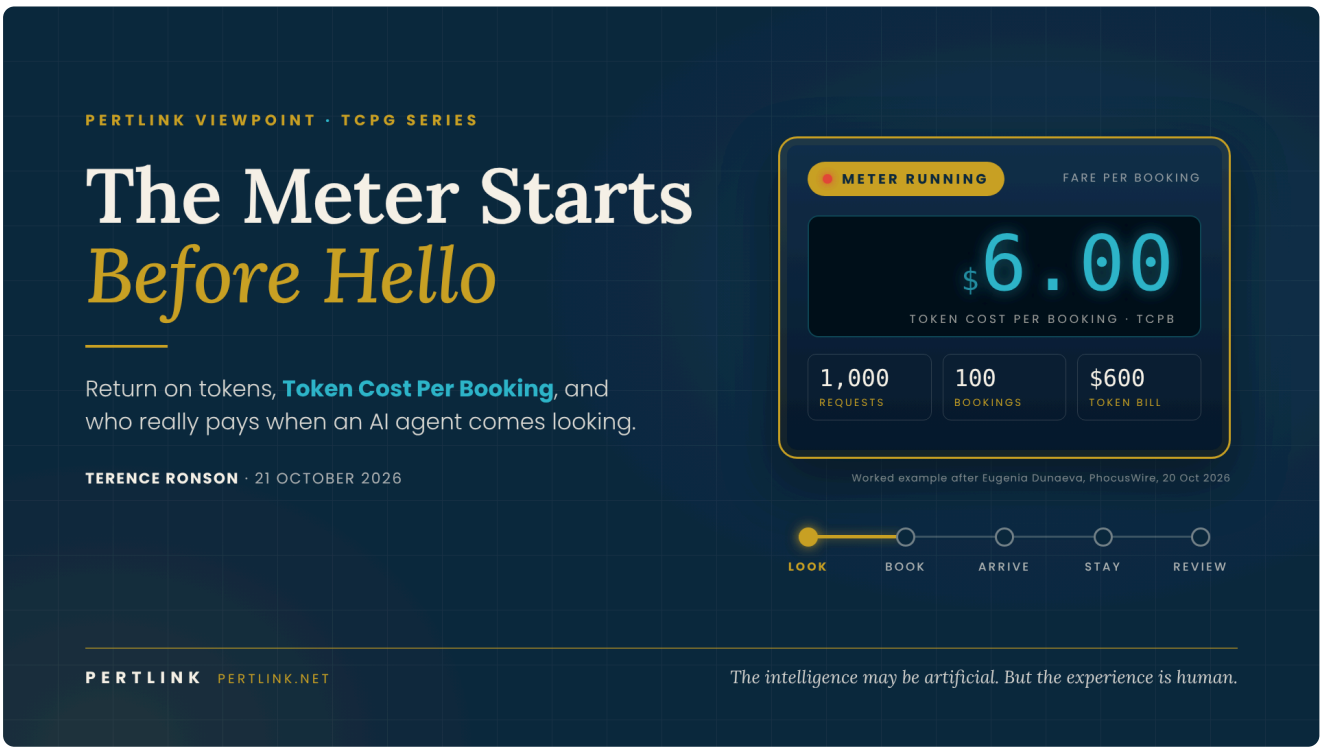


FIGURE 2.1 · Hero graphic from The Meter Starts Before Hello, Hospitality Net. Worked example after Eugenia Dunaeva, PhocusWire, 20 October 2026; TCPB is Pertlink’s.

02 · THE METER STARTS AT THE LOOK

Nobody made a mistake. *The meter ran.*

The agent's cost of sale isn't fixed. Change only how it behaves, and the cost of selling each room doubles.

SCENARIO	TOKEN BILL	BOOKINGS THAT STICK	REVENUE	TCPB	ROT
Dunaeva's example	$600	100	$3,000	$6	400%
The agent checks twice as often	$1,200	100	$3,000	$12	150%
Half its bookings are spare rooms, later released	$600	50	$1,500	$12	150%

1,000 requests at 60 cents each, $30 earned per booking that sticks. First row: Dunaeva, PhocusWire. Other rows: Pertlink illustration.

The spare rooms are not hypothetical. With free cancellation, Dunaeva notes, an agent can book the best room, then the second best, then the third, and release the extras at the last minute. Your revenue system may read that as demand.

Who pays: we have seen this film before

Right now, only the agents' investors pay for the looking. Instinct is free to travellers, and so is Meta's Muse, up to a usage limit. Booking Holdings' new agent, Vayo, is free in beta. There are only three places the cost can go.

INVESTORS

Won't fund it forever

Subsidised looking is a launch strategy, not a business model.

TRAVELLERS

Won't pay

Free is the hook. Charging the guest breaks the product.

THE SUPPLIER

That leaves the hotel

Instinct's founder, Noah Shinn, has already floated boutique hotels paying up to 30% per booking.

Hotels have seen this film. Online travel agencies and metasearch were free for travellers too. Hotels funded them through commission and paid per click. And the route may not be direct: with Expedia selling hotels inside Muse as merchant of record, an agent's cost of sale can reach the hotel inside a commission it already pays.

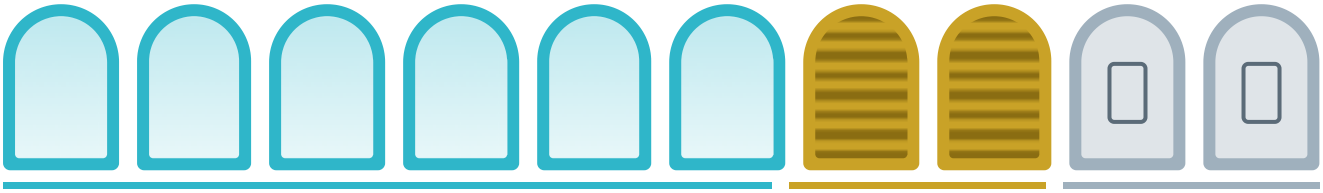
A platform's cost of sale sets the floor under what it will charge. It doesn't set the price. Platforms price at value, which is why the figure floated is 30% and not $6.

02 · THE METER STARTS AT THE LOOK

Make yourself *cheaper to sell*

Part of an agent's cost of sale depends on you. Every extra page it has to read, every form it fills in twice, every block it has to retry, adds to its bill.

A hotel whose rates, rooms and policies are clear and readable by machines is cheaper to sell than one the agent has to fight. Cheaper to sell is a negotiating position. Most hotels haven't taken it, because they haven't decided anything.



6 reached checkout
The agent could complete a booking.

2 blocked
By bot-management software, not a published AI policy.

2 behind apps
Booking kept behind apps and SMS logins.

FIGURE 2.2 · On 18 October 2026, Pertlink sent an AI agent to book a one-night stay on the websites of the world's ten largest hotel groups. A repeat test showed the outcome could change from one day to the next, and from one agent to another. Source: Pertlink, *Are Hotels Closing Their Doors to AI Agents?*, as reported in The Meter Starts Before Hello.

Hotels haven't closed their doors to AI agents. Most haven't decided whether to open them. A firewall setting isn't a policy, and it can't be priced.

Agents will identify themselves

The tools to price the look are arriving. On 9 October, Sierra published a draft Personal Agent Protocol, known as Poppy, which lets a traveller's agent identify itself, get the customer's permission and work with a business. Launch partners include Meta, Instinct, Shopify, Stripe and Walmart. Its 35 design partners include United Airlines and Hertz.

Airlines and car rental are at the table. No hotel group is on either list.

Your direct channel has a TCPB too

The booking assistant on your website, the pre-arrival messaging and the rate tools all run a meter. Know that number, and you can say what a direct booking costs you in tokens, beside what an agent platform will ask. Without it, there is nothing to negotiate with.

The hotels that are cheap to sell and know who is looking will negotiate. The ones that aren't will be invoiced.

02 · THE METER STARTS AT THE LOOK

Bumps *in the road*

No argument survives without its objections. Here are the ones a sceptical board member will raise, and what the evidence says.

THE OBJECTION	THE EVIDENCE	WHAT IT MEANS FOR HOTELS
"Token prices will fall, and the problem will go with them."	Agents are getting leaner. But the Linux Foundation reports that per-token costs fell heavily in 2023–2025, have since levelled off, and that new model prices are rising. McKinsey finds the effective cost of AI is already climbing for many enterprises, because usage grows faster than unit prices fall.	Don't plan on the cost disappearing. Cheaper looks invite more looking.
"Platforms will make their money elsewhere."	Advertising, paid placement, subscriptions and data are all options. But, as Dunaeva points out, an agent paid by some hotels and not others is no longer a neutral assistant.	If the bill arrives as a ranking fee rather than a commission line, it is harder to see, and the hotel still pays it.
"The cost will be absorbed by intermediaries."	With Expedia set to sell hotels inside Meta's Muse, the agent's cost of sale may sit between two platforms before it reaches the hotel.	That doesn't remove it. It buries it inside a commission you already pay, where it is harder to challenge.
"Regulators won't allow it."	Some will push back. In July, China's market regulator imposed a fine and confiscated gains totalling about US$763 million on Trip.com Group over exclusive deals with hotels (PhocusWire).	Expect scrutiny of pay-to-rank agents, but don't wait for a regulator to negotiate your commission.
"We'll just block them."	Blocking an agent means turning away the guest behind it. In Pertlink's test, the blocking came from security software, not any stated AI policy.	The choice isn't open or closed. It is on what terms. Some hotels may not know they are turning agents away.
"Free cancellation is a guest benefit."	It is, for guests. For an agent holding three rooms to release two, it is free inventory, and your revenue system can read it as demand.	Cancellation terms for agent bookings will become a pricing tool. Decide them before the platforms do.

THE STRESS TEST

McKinsey advises stress-testing every AI business case both ways. In its AI value curve research, a 30 percent rise in the effective cost of AI cuts the global value pool from $2.52 trillion to $1.67 trillion, a 34 percent drop. A few dollars of TCPB passed down the chain matters more to a hotel than to a bank.

03

BUDGET THE METER

Budget the meter, *not the machine*

Put AI on a rolling forecast, count it per guest, and fund the people as seriously as the platform.

ESSENTIAL FOR OWNERS GMS C-SUITE CAIO VENDORS

From licence to meter

Most hotel technology budgets assume software behaves like a fixed asset: licences, annual renewals, predictable costs. AI breaks that model. Every guest interaction an AI system handles incurs a cost, and that cost grows as adoption grows. Usage can rise quickly with guest volumes, adoption rates and changing vendor pricing, creating swings an annual plan will miss.

> *"Keep the annual budget for payroll and capex. Put AI on a rolling monthly or quarterly forecast, with a clear trigger for when usage, and cost, outruns the plan. Flexibility without a trigger is just drift."*
>
> TERENCE RONSON, SHIJI INSIGHTS

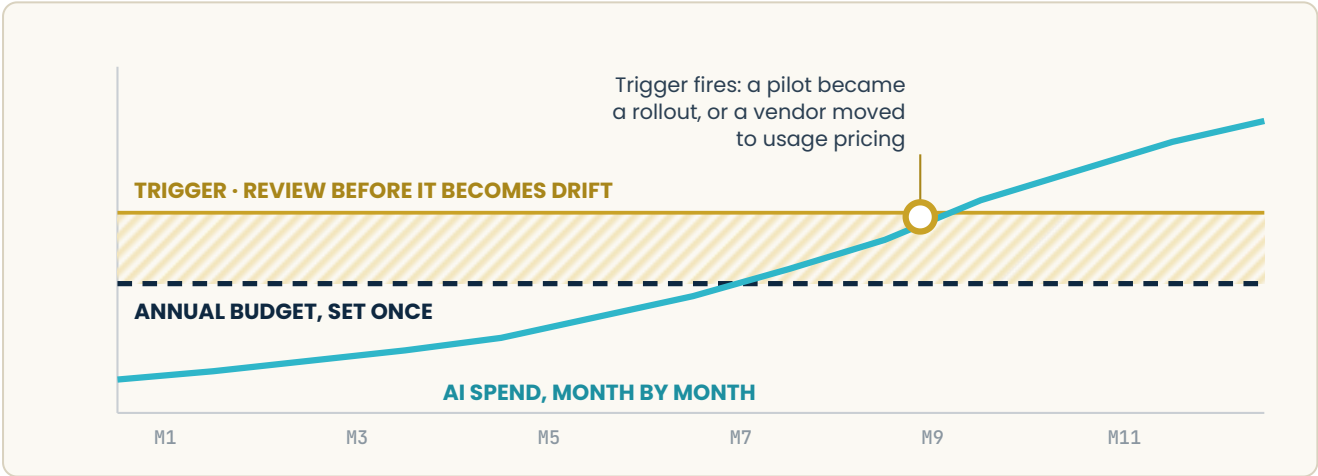


FIGURE 3.1 · An annual budget set once, against AI spend that moves with use. Pertlink illustration; no real figures.

What the trigger should watch

- **Spend against forecast,** month by month, by property.
- **TCPG rising without outcomes.** A rising TCPG beside rising revenue is a good month. Beside flat outcomes, it is waste.
- **A vendor pricing change,** above all a move from per-room fees to usage.
- **A pilot becoming a rollout.** McKinsey's research finds AI spend "nearly quadruples" from pilot to production.

Source: Waisun Chan, Budget the Meter, Not the Machine, Shiji Insights, 19 October 2026.

03 · BUDGET THE METER

Token Cost *Per Guest*

The number every hotel needs and almost none keeps: the AI cost of serving one guest across the whole journey, from the first pre-arrival question to the reply to their post-stay review.

TOKEN COST PER GUEST · TCPG

fully loaded AI cost across the guest journey ÷ guests served

Pertlink. Counted on successful outcomes only: guests served, not conversations held.

What "fully loaded" means

- Provider accounts and staff subscriptions
- Cloud compute behind every model call
- The AI share of software you already buy
- Orchestration, retries and failed attempts
- The person who checks, corrects or rewrites the output

McKinsey now counts AI cost the same way. Its AI value curve prices work done with AI as "model and token costs as well as any human review, validation, and rework required."

Never read it alone

TCPG is a cost, not a verdict. A rising TCPG next to rising revenue or better guest scores is a good month. Next to flat outcomes, it is waste.

"A busy chatbot is not a profitable chatbot," Ronson told Shiji Insights. "The most common ROI mistake I see is counting the conversations and never counting the cost of each one, or what happened to the guest afterwards."

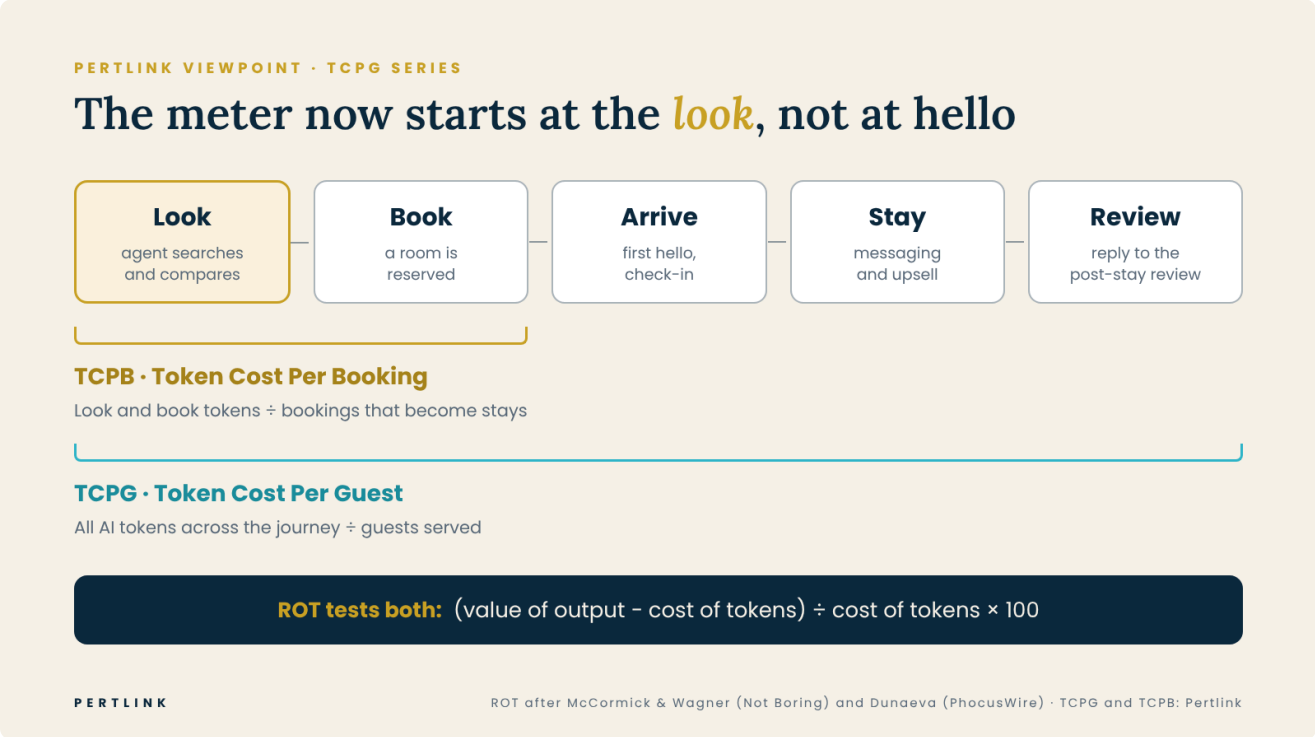


FIGURE 3.2 · TCPB is the front end of TCPG. ROT is the test both have to pass. Pertlink; ROT after McCormick & Wagner and Dunaeva.

03 · BUDGET THE METER

Three lines every *2027 budget* needs

"None of these are glamorous," Ronson says. "All of them are the difference between an AI strategy and an AI expense."

1

AI literacy for every department

Funded as a ring-fenced line, not from whatever is left over.

2

Data foundations and governance

With the people and processes to sustain them.

3

Metered AI spend with TCPG tracking

So usage, cost and guest outcome appear in the same report.

After Waisun Chan, Budget the Meter, Not the Machine, Shiji Insights, 19 October 2026.

> *"Buying AI without training your people is like installing a grand piano in a hotel with no pianist. It looks impressive in the lobby. Nobody gets any music."*
>
> TERENCE RONSON, SHIJI INSIGHTS

Three kinds of AI, three different sums

"Today, we're using AI very much as a blanket term," says McKinsey's Lari Hämäläinen. "If you simplify it, there are three different types of AI." Put all three in one budget line and a staff subscription, a guest-facing agent and a transformation programme end up judged by the same yardstick, and all three look wrong.

KIND OF AI	AT A HOTEL	HOW TO JUDGE IT
Personal tooling	The sales manager drafting proposals; the GM summarising the month's reports.	A cost of doing business. Cap it, govern it, and don't expect a return on every prompt.
Automation and efficiency	Review replies, invoice matching, the guest questions asked a hundred times a week.	Cost per successful output, TCPG and TCPB. The work is countable, so count it.
Rebuilding for AI	Redesigning the front office so AI carries the routine and people own the moments that matter.	Over a longer horizon, on guest outcomes as well as cost. This is the thinking behind HXO.

AND KEEP THREE OUTCOMES APART

Efficiency is doing the same work faster. **Cost saving** is removing spend. **Value creation** is revenue or loyalty that didn't exist before. Business cases often blur them. Each needs its own measure.

04

READ THE METER

Read *the meter*

Where AI spend hides, why the bill jumps, and how to count only what worked.

ESSENTIAL FOR OWNERS GMS C-SUITE CAIO VENDORS

Before anyone can count cost per guest, they need to know what they are counting. "AI cost is perhaps not hard to decipher," says McKinsey's Kaavini Takkar, "but it is very complicated." She names four places AI spend hides. Each has a hotel address.

1

Direct spend with AI providers

The hotel's own accounts with the model makers, and the staff subscriptions somebody put on a corporate card.

2

Cloud compute

Behind every model call. Often billed to IT, or to a technology partner, rather than to the department using the AI.

3 OFTEN THE LARGEST

AI inside software you already buy

The guest messaging platform, the review tool, the revenue system. For many hotels this is the largest AI line, and the least visible, because it sits inside a vendor fee.

4

Experimentation

The pilot at one property, the trial a department head signed up for, the tool nobody switched off.

An agent adds its own lines

Takkar's starting point is simple: "Every single time an agent receives a prompt, it is going to make model calls. You multiply your model calls by the per-token price." Then come the costs on top: human validation, monitoring, storage of the hotel's knowledge for the AI to search, the plumbing that connects the agent to other systems, and the cost of failure.

Failure runs the meter twice

Takkar's example is an airline rebooking agent that offers a flight, waits, then has to say: "This flight is actually already gone. Let's go through this process again."

The hotel version: the upsell agent that sells a room type already sold out, or the concierge that confirms a spa slot someone else has taken. The guest gets an apology. The hotel gets the bill for both conversations.

Quotations: What does AI really cost?, McKinsey, 14 October 2026. Hotel examples: Pertlink.

04 · READ THE METER

Why *the bill jumps*

"If the price per token keeps falling, why are the bills going up?" asks one of McKinsey's section titles. Because price per token was never the number that mattered. Volume, design and discipline are.

WHAT CHANGED	WHAT HAPPENED	SOURCE	THE HOTEL EQUIVALENT
Pilot to production	AI spend "nearly quadruples" on average, according to McKinsey research.	Takkar	A pilot at one property, rolled out across a portfolio
Per-seat to usage pricing	Early in 2026, many frontier companies moved to pricing where "the more you eat, the more you have to pay." The same AI could cost two to three times as much, instantly.	Lala	Vendor contracts moving from a fee per room to a charge per use
A poorly designed agent	"The same agent could consume up to 30 times as many tokens" as a well-designed one.	Lala	Guest messaging that re-sends the manual and re-reads the whole stay with every reply
A bigger model by default	One client's usage and spend doubled in two months as users moved from Claude Sonnet to the more advanced Opus.	Lala	Every question, breakfast hours included, sent to the most powerful model
No guardrails	A bank spending $2 million to $3 million a year saw projections go to more than $25 million once everyone had every tool and model.	Takkar	Every department given every tool before anyone set the rules

Source: What does AI really cost?, McKinsey, 14 October 2026. Hotel column: Pertlink.

Only one of these is a change in price. The other four are management decisions that nobody knew they were making.

Prices aren't falling the way you think

The Linux Foundation reports that per-token costs fell heavily in 2023–2025, have since levelled off, and that new model prices are rising. In June it announced its intent to launch a Tokenomics Foundation, with the FinOps Foundation, to set open standards for measuring AI token spend.

Booking.com is among its first supporters. Its IT finance lead, Chris Reed, gave the reason: at its scale, "small differences in efficiency compound into very large numbers." One of the world's largest booking platforms is helping write the rules for how tokens are counted. Hotels aren't in the room.

WHAT STANDARDS WON'T DO

Standards bodies will define the unit: the token, its price, the benchmark. They won't define a hotel's denominator. If hospitality doesn't decide what tokens should cost per guest and per booking, the platforms between the guest and the hotel will decide, and invoice accordingly.

04 · READ THE METER

Count only *what worked*

How do you know if an agent is worth it? "One good metric to look at would be the cost of the agent per successful output," says McKinsey's Takkar. Then compare it with "the fully loaded cost of a human or of another tool that had been doing the job."

Failures count

The upsell that sold a room already gone isn't a successful output. Neither is the review reply a manager had to rewrite. Their tokens go into the numerator, but they add nothing to the denominator. That is why TCPG is divided by guests served, and TCPB by bookings that become stays, never by conversations held.

Frequency decides

Takkar contrasts the thousands of rebooking requests an airline gets every day with an agent that audits tax filings "perhaps once a year." The first earns its build cost back quickly. The second may never. Guest messaging runs all day, every day. The annual budget pack runs once. An agent may suit the first long before it suits the second, however impressive the demo.

Where the meter runs, stage by stage

STAGE	WHAT AI DOES	COST	WHAT DRIVES IT	KEEP IT DOWN
Look The guest's agent shops	Agents ask about rates and rooms; your booking AI replies; many hotels compared at once	●●○	**Volume.** Paid on every look, and most looks never book.	Publish clean, machine-readable content so agents need fewer calls.
Book Booking confirmed	Confirmations and payments; pre-arrival questions; personalised offers	●○○	**Short, one-off messages.** Personalised offers cost more.	Template the routine. Spend tokens where the words are personal.
Check-in and stay Arrival and the stay	Digital check-in and ID; concierge chat all stay; requests routed to teams	●●●	**Context.** Every reply re-reads the stay: a 10,000-token chat can mean 30,000 read.	Summarise the stay. Send simple questions to a smaller model.
Check-out Departure and review	Folio and invoice queries; feedback analysis; review replies a manager edits	●○○	**Rework.** A reply a manager rewrites is paid for twice: tokens and time.	Track how often managers rewrite: your AI Override Rate.

TABLE 4.1 · Relative token cost per stage, Pertlink illustration: ●○○ low, ●●○ medium, ●●● high. Look and Book roll up into TCPB; every stage rolls up into TCPG. A landscape version of this map is available as the graphic *Where the Meter Runs*.

> *"AI will be worth it as long as you can measure it. Without measurement, there can be no success."*
>
> WASIM LALA, MCKINSEY

05

GOVERN THE METER

Govern *the meter*

One view of the spend, clear rules for what AI may decide, and a measure of how often people step in.

ESSENTIAL FOR OWNERS GMS C-SUITE CAIO VENDORS

One place to see the spend

McKinsey's Takkar starts with a single AI control plane: "one place where I can bring all these different sources of my AI spend together." For a hotel group, that means provider accounts, cloud, pilots and the AI inside every vendor fee, in one view.

How do you measure token burn when vendors don't show it? You bring every bill into that view, and you make usage reporting a condition of the contract. Pertlink makes the wider case for oversight in Who Controls the Controls? and Part Two.

What AI may decide

A **Decision Envelope** documents what an AI system may decide alone, what it must escalate and what it may never do. It belongs in the budget, not just the policy manual, because every boundary has a cost: escalations take staff time, and decisions outside the envelope carry liability.

That liability is real. In Moffatt v Air Canada, a Canadian tribunal held the airline liable for incorrect refund guidance from its website chatbot.



FIGURE 5.1 · A Decision Envelope for a guest messaging agent. Pertlink illustration; every hotel writes its own.

AI OVERRIDE RATE

AI decisions corrected or reversed by staff ÷ AI decisions taken

Pertlink. "If nobody ever overrides the AI, either it's perfect or nobody's watching." Terence Ronson, Shiji Insights.

Read the rate by agent and by department. Zero is a warning, not a triumph. A spike says the envelope or the AI needs work. And every override has a token cost: the AI's attempt and the person's correction both sit in TCPG.

05 · GOVERN THE METER

Worked example: *the upgrade decision*

Rules, AI and people in one decision. The yes-or-no gates run on rules at near-zero token cost. Tokens are spent only where words are needed. Every path is logged and counted per guest.

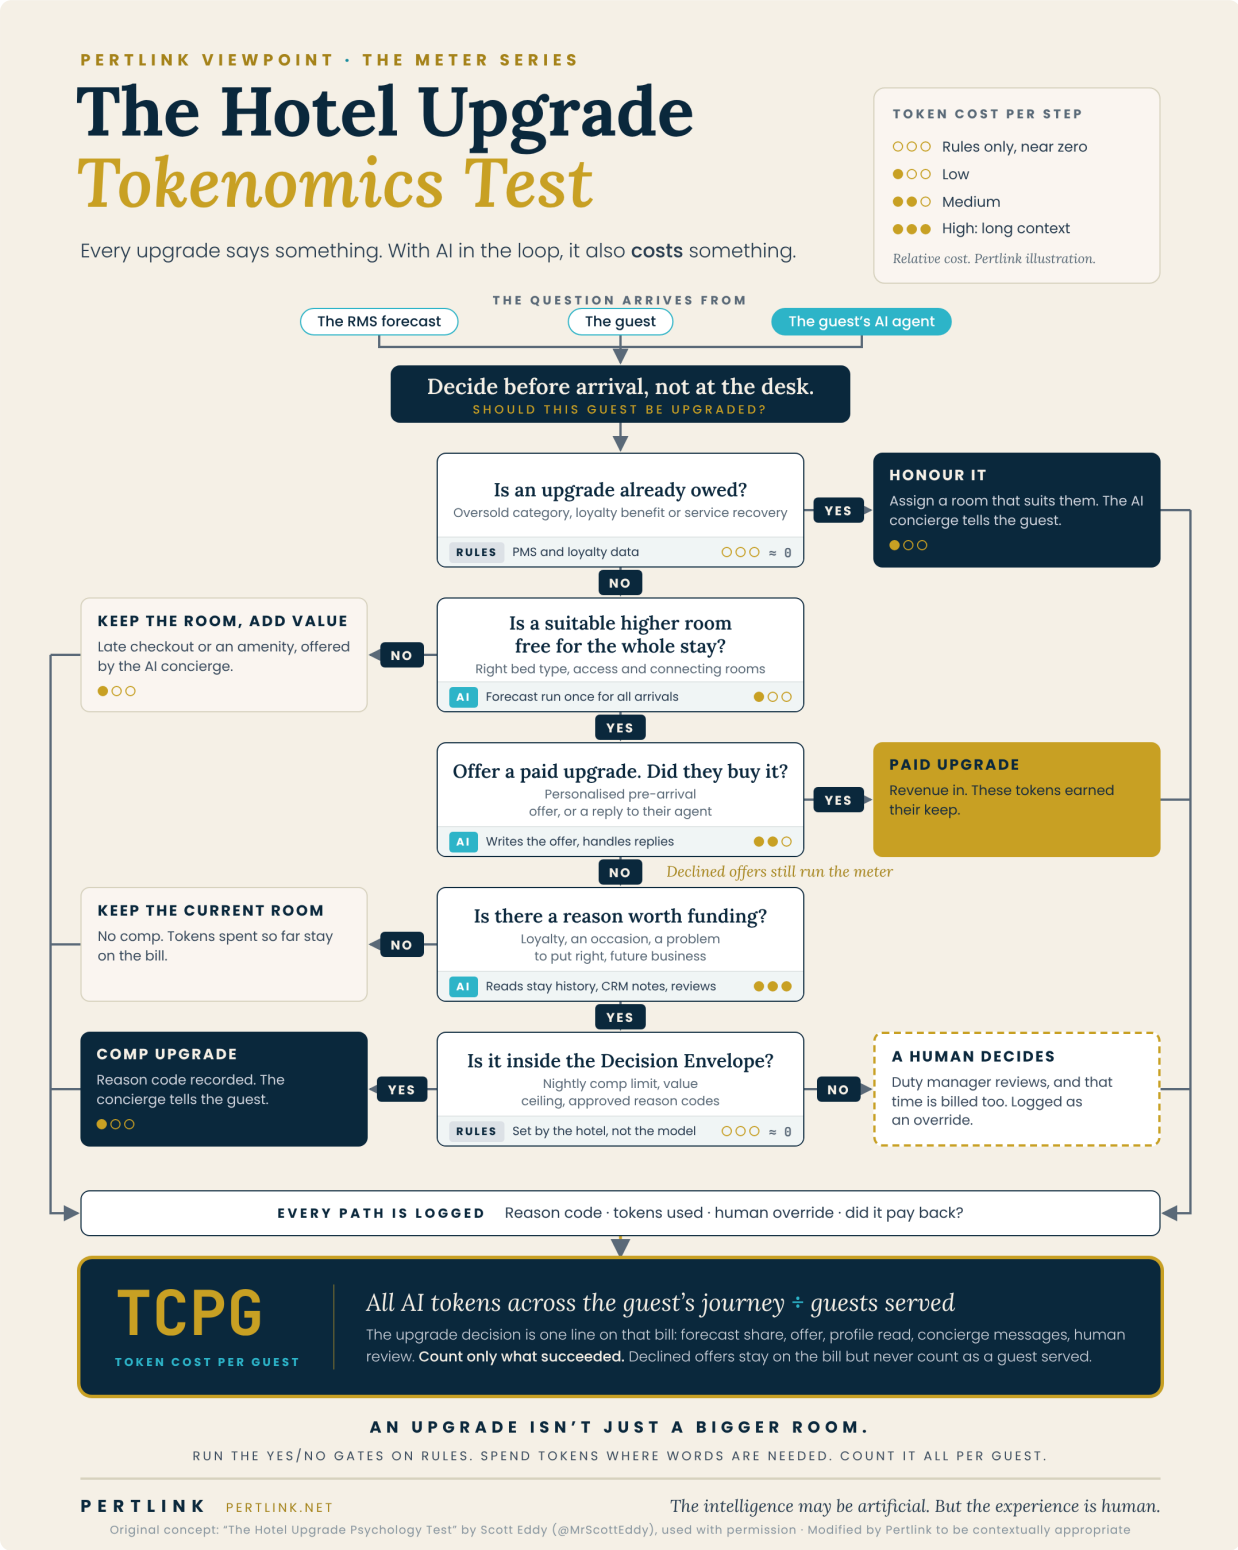


FIGURE 5.2 · The Hotel Upgrade Tokenomics Test. Original concept: "The Hotel Upgrade Psychology Test" by Scott Eddy (@MrScottEddy), used with permission. Modified by Pertlink to be contextually appropriate. Token costs are relative.

05 · GOVERN THE METER

Govern for *people*, not just platforms

The point of reading the meter isn't to spend less on AI. It is to spend it where it gives people more time with guests.

The Human Experience Orchestrator

Pertlink's HXO model judges every AI investment with two questions. The second matters more.

QUESTION ONE

Does it improve efficiency or reduce the cost of serving guests?

QUESTION TWO

Does it create space for meaningful human interaction?

Boring AI pays for exciting AI

Start behind the scenes, where data is cleaner, the risk to guests is lower and returns are easier to prove: invoice matching, linen, forecasting, the morning briefing. "Get the back of house working first," Ronson says, "then use what it saves to fund what guests will actually see."

Budget for the pianist

Revenue managers, housekeeping supervisors, engineers, HR teams and GMs will all work with AI. Ring-fence their training. Make token literacy part of it, so every department can tell a cheap question from an expensive one.

AI economics are local

AI costs much the same in Manila as in Manhattan. Wages don't. Pertlink's reading: for hotels across much of Asia-Pacific, the margin AI has to work with is thinner, so routing, caching and right-sizing matter more, not less. Fragmenting data-localisation rules across the region also shape where guest data can sit and which AI vendors a group can use. Those are budget decisions as much as legal ones.

The real advantage is judgment

By 2030, access to AI will be increasingly standardised. Most competitors will have similar tools. What won't be commoditised is organisational judgment, operational discipline, and the ability to create memorable guest experiences.

The intelligence may be artificial. But the experience is human.

06 · ROLE PLAYBOOK

Owners and investors

YOUR QUESTION

What does AI cost per guest, and is it buying anything?

YOUR NUMBERS

TCPG by property, every quarter
Beside RevPAR, GOP and guest scores. Rising TCPG with rising outcomes is a good quarter. With flat outcomes, it is waste.

TCPB on direct bookings
Your benchmark when an agent platform names its price. Without it, there is nothing to negotiate with.

AI spend against forecast
And whether a trigger fired, and what management did about it.

AI Override Rate
Proof that someone is watching. Zero is a warning.

DO THIS QUARTER

- [ ] Ask management for TCPG and TCPB, counted on successful outcomes only.
- [ ] Require AI spend on a rolling forecast with a written trigger.
- [ ] Approve a ring-fenced AI literacy line, not leftovers.
- [ ] Stress-test AI business cases for a 30% rise in AI costs, both ways, as McKinsey advises.
- [ ] Decide your policy on AI agents: what they may do and on what terms. Don't leave it to the firewall.
- [ ] Make usage reporting a condition of management and vendor contracts.

RED FLAGS

- Business cases that count conversations, not cost or what happened to the guest.
- "Tokens consumed" or "users onboarded" reported as success. That is tokenmaxxing.
- A pilot approved without a portfolio cost estimate. Spend "nearly quadruples" from pilot to production.

ASK MANAGEMENT

- *What does AI cost us per guest, and what did it buy?*
- *Where are all the places we pay for AI?*
- *What happens to our TCPB if an agent platform charges us?*
- *Who decided how our booking engine treats AI agents: us, or a security vendor?*

Pertlink recommendations, drawing on chapters 2, 3 and 5. McKinsey figures: What does AI really cost? and The AI value curve.

06 · ROLE PLAYBOOK

General managers and operators

YOUR QUESTION

Where is the meter running in my hotel, and who is reading it?

YOUR NUMBERS

TCPG for the hotel, monthly
Read beside the month's guest scores and revenue, never alone.

Cost per successful output
For your busiest AI jobs: guest messaging, review replies, upsells.

AI Override Rate by department
Who is correcting the AI, how often, and why.

Failure loops
Offers for rooms already sold, replies rewritten, requests sent to the wrong team.

RED FLAGS

- Week-long guest chats with no summarising.
- Long, effusive replies to every review. Length is a cost decision.
- Pilots nobody switched off.

DO THIS QUARTER

- [ ] Name a meter-reader in every department that uses AI.
- [ ] List every AI tool in the building, including the ones on corporate cards.
- [ ] Send simple guest questions to templates or the smallest model that answers well.
- [ ] Stop sending the full brand manual with every message. Cache the fact sheet.
- [ ] Summarise long stays instead of re-reading the whole chat.
- [ ] Start with back-of-house AI, where data is cleaner and returns are easier to prove.

ASK YOUR TEAM

- *Who here can explain what a token is?*
- *Which model answers our simplest guest question, and does it need to?*
- *When did someone last override the AI, and why?*
- *What did our guest messaging cost per guest last month?*

Pertlink recommendations, drawing on chapters 1, 4 and 5.

06 · ROLE PLAYBOOK

The C-suite and the CFO

YOUR QUESTION

How do we budget, account for and forecast a cost that moves with use?

YOUR NUMBERS

Total AI spend in one view
Providers, cloud, AI inside vendor fees and experiments, by property and brand.

TCPG and TCPB by brand, region and channel
Counted on successful outcomes: guests served, bookings that become stays.

Forecast variance and trigger events
Month by month, with the reason for each trigger.

Cost per successful output
Against the fully loaded cost of the person or tool that did the job before.

RED FLAGS

- A vendor moving from per-room fees to usage pricing. The same AI can cost two to three times as much, instantly.
- Bigger models by default. One client's spend doubled in two months from a model switch.
- Access before guardrails. One bank's projections went from $2–3 million a year to more than $25 million.

DO THIS QUARTER

- [ ] Move AI to a rolling monthly or quarterly forecast. Keep the annual budget for payroll and capex.
- [ ] Set triggers: spend over forecast, TCPG rising without outcomes, a vendor pricing change.
- [ ] Budget the three kinds of AI separately: personal tooling, automation, rebuilding.
- [ ] Measure efficiency, cost saving and value creation separately.
- [ ] Build an AI FinOps capability. McKinsey calls it "a requirement, not an option."
- [ ] Model a 30% rise in AI costs, and a pilot-to-portfolio jump.

PUT IN VENDOR CONTRACTS

- *Usage and cost reporting per property and per task.*
- *Notice periods and caps on pricing changes.*
- *Which model handles which task, and our right to change it.*
- *Where guest data is processed and stored.*

Pertlink recommendations, drawing on chapters 3 and 4. McKinsey figures: What does AI really cost? and The AI value curve.

06 · ROLE PLAYBOOK

Chief AI Officers

YOUR QUESTION

Is every token going to the right model, for a task worth doing, inside the rules?

YOUR NUMBERS

Tokens per task, by model tier
And the share of tasks routed to no model at all.

Context size, cache hit rate, chat length
The three levers behind most runaway bills.

Retries and failure loops per agent
Every loop runs the meter again.

Override Rate per agent and boundary
Mapped to each Decision Envelope, with TCPG and TCPB from telemetry, not estimates.

RED FLAGS

- A poorly designed agent can use up to 30 times the tokens of a good one.
- Shadow AI: tools bought on cards, outside the control plane.
- An override rate of zero.

DO THIS QUARTER

- [ ] Build a single AI control plane: every source of AI spend in one view.
- [ ] Publish a routing policy: rules first, small model next, frontier and reasoning only when needed.
- [ ] Set context standards: cache static facts, summarise history, trim prompts.
- [ ] Write a Decision Envelope for every agent: alone, escalate, never.
- [ ] Set a policy for outside agents: how they identify, what they may book, on what cancellation terms. Track Sierra's Personal Agent Protocol.
- [ ] Run token literacy for every department, starting with HFTP's AI in Hospitality: Essentials.

ASK YOUR VENDORS

- *Show me tokens per task and the model used.*
- *Can we route tasks to a cheaper or open-weight model?*
- *How do you cache and summarise?*
- *Where are overrides and escalations logged?*

Pertlink recommendations, drawing on chapters 1, 2, 4 and 5. Agent design figure: Wasim Lala, What does AI really cost?

06 · ROLE PLAYBOOK

Vendors and technology partners

YOUR QUESTION

Can my hotel clients read my meter, and would they like what they see?

YOUR NUMBERS

Usage and cost per property, per task
The first thing a CFO following this playbook will ask for.

The model behind each task
And whether the hotel can route it elsewhere.

TCPG-ready data
Exports a hotel can drop straight into its per-guest and per-booking numbers.

Override and escalation logs
Mapped to the hotel's Decision Envelope.

RED FLAGS

- ! AI cost buried in a seat fee. For many hotels, AI inside software they already buy is the largest AI line and the least visible.
- ! Agents that re-send the manual and re-read the whole stay with every reply.
- ! No way for the hotel to see, or reverse, what the AI decided.

DO THIS QUARTER

- ☐ Add a TCPG-ready usage report to your product.
- ☐ Offer usage alerts and caps that hotels can set themselves.
- ☐ Route simple tasks to cheaper models, and say so.
- ☐ Price on successful outcomes where you can, not on conversations held.
- ☐ If you sell distribution or booking tech, make rates, rooms and policies machine-readable. Cheaper to sell is your clients' negotiating position.
- ☐ Give notice before any move from per-room to usage pricing.

WHAT HOTELS WILL ASK YOU

- ? *What does your AI cost us per guest?*
- ? *Which model does what, and can we change it?*
- ? *How will your pricing change next year?*
- ? *Can we see every decision the AI made, and override it?*

Pertlink recommendations, drawing on chapters 2 and 4. Transparency is a feature: the vendor that helps a hotel read the meter is the last one it switches off.

FORMULA CARD

Six sums *worth knowing*

RETURN ON TOKENS · ROT

(value of output − cost of tokens) ÷ cost of tokens × 100

`($3,000 - $600) ÷ $600 × 100 = 400%`

McCormick & Wagner, Not Boring

TOKEN COST PER BOOKING · TCPB

look and book AI cost ÷ bookings that become stays

`$600 ÷ 100 bookings = $6`

Pertlink, on Dunaeva's example

TOKEN COST PER GUEST · TCPG

fully loaded AI cost across the journey ÷ guests served

`Use the worksheet below`

Pertlink

COST PER SUCCESSFUL OUTPUT

total cost of the agent ÷ successful outputs

`Compare with the fully loaded cost of the person or tool it replaces`

Kaavini Takkar, McKinsey

AI OVERRIDE RATE

AI decisions corrected or reversed ÷ AI decisions taken

`Zero is a warning. Spikes are signals.`

Pertlink

THE STATELESS RE-READ

tokens added per day × n(n+1) ÷ 2

`Five nights at 2,000 a day: 2,000 × 15 = 30,000 read, three times the chat`

Pertlink arithmetic

TCPG worksheet

Fill in one per property, per month. Count only successful outcomes in the denominator.

Line	This month	Last month	Notes
Provider accounts and staff subscriptions			
Cloud compute for AI			
AI share of vendor fees			
Orchestration, retries and failed attempts			
Human review, validation and rework (hours × rate)			
Experiments and pilots			
Total AI cost			
÷ Guests served			
= TCPG			
Beside it: revenue, guest score, AI Override Rate			

CHECKLIST

The *2027 budget* checklist

Eleven questions and three lines, before the budget is signed. Tick them in the meeting, not after it.

READING THE METER

- [ ] Can we list every place we pay for AI? Provider accounts, cloud, pilots, and the AI inside every vendor fee.
- [ ] Do we know our TCPG and TCPB, counted on successful outcomes only?
- [ ] Is each AI use judged by the right sum: personal tooling, automation, rebuilding?
- [ ] Which model answers our simplest guest question, and does it need to?
- [ ] What happens to the bill when a pilot becomes a portfolio rollout, or a vendor moves to usage pricing?
- [ ] Who on the team can explain what a token is? If nobody can, start there.

AI AGENTS, FOR THE BOARD

- [ ] What does an agent booking cost the platform to make, and what will it ask us to pay?
- [ ] What is our own TCPB on direct bookings, and what happens to it if AI costs rise 30%?
- [ ] How much of our booking-engine traffic is machines, and who decided how we treat it: us, or our bot-management vendor?
- [ ] Are we cheap or expensive for an agent to sell?
- [ ] If a personal agent identified itself tomorrow, what would we let it do, on what cancellation terms, and at what price?

And the three lines

- [ ] AI literacy for every department, ring-fenced
- [ ] Data foundations and governance
- [ ] Metered AI spend with TCPG tracking

Questions adapted from Pertlink's *Read the Meter* and *The Meter Starts Before Hello*. The three lines: Terence Ronson in Shiji Insights.

SIGNED OFF BY

OWNER	GENERAL MANAGER	CFO	CHIEF AI OFFICER

Date of review: _____ Next trigger review: _____

Every word runs the meter. Make sure someone in the building can read it.

GLOSSARY

Words on *the meter*

Agent AI that takes actions for someone, such as searching, booking or replying, making many model calls per task.

AI control plane One place where every source of AI spend and activity is brought together.

AI FinOps Managing AI cost like any other factor cost, down to the cost per useful outcome.

AI Override Rate How often staff correct or reverse AI decisions. Pertlink.

Caching Storing content that doesn't change so it isn't paid for at full price each time.

Context window Everything sent to the model with a request: instructions, reference, conversation.

Cost per successful output An agent's total cost divided by the tasks it completed successfully.

Decision Envelope What an AI may decide alone, must escalate, and may never do. Pertlink.

HXO Human Experience Orchestrator: AI carries the routine, people own the moments that matter. Pertlink.

Input tokens What the model reads.

Model routing Sending each task to the cheapest model that does it well, or to no model.

Open-weight model A model whose weights are published, so a business can run and control it.

Output tokens What the model writes. Priced higher than input.

Personal agent An AI acting for a traveller, for example shopping for hotels.

ROT Return on tokens: (value − token cost) ÷ token cost × 100. McCormick & Wagner.

Stateless Keeping no memory between calls, so history must be sent again.

TCPB Token Cost Per Booking: look and book AI cost ÷ bookings that become stays. Pertlink.

TCPG Token Cost Per Guest: fully loaded AI cost of serving one guest across the journey. Pertlink.

Thinking tokens Hidden working tokens that reasoning models generate, and bill.

Token The smallest unit of information a model sees, and the unit AI is billed in.

Tokenmaxxing Rewarding token consumption as though it were progress.

Tokenomics The economics of tokens: what AI costs per unit, per task and per guest, and who pays.

SOURCES

Where the *numbers* come from

1 Chan, W. (19 October 2026). Budget the Meter, Not the Machine: Planning for 2027 When AI Runs on Usage. Shiji Insights.

2 Dunaeva, E. (20 October 2026). Return on tokens: Who pays for the free travel agent? PhocusWire.

3 Hämäläinen, L., Takkar, K., and Lala, W. (14 October 2026). What does AI really cost? McKinsey Explainers.

4 Van der Veken, L., Durth, S., Catlin, T., Alatovic, T., and Isherwood, T., et al. (14 October 2026). The AI value curve: The economics of applying AI to existing work. QuantumBlack, AI by McKinsey.

5 McCormick, P., with Wagner, M. (10 June 2026). Return on Tokens (ROT). Not Boring.

6 The Linux Foundation (3 June 2026). Linux Foundation Announces the Intent to Launch the Tokenomics Foundation.

7 Taylor, B., and Bavor, C. (9 October 2026). Sharing a draft of Personal Agent Protocol. Sierra.

8 American Bar Association (February 2024). BC Tribunal Confirms Companies Remain Liable for Information Provided by AI Chatbot. Business Law Today. On Moffatt v Air Canada.

9 HFTP Academy. AI in Hospitality: Essentials. Course.

10 PhocusWire (24 September 2026). Report on Expedia hotel booking inside Meta's Muse. PhocusWire (27 July 2026). Report on the Trip.com Group antitrust penalty.

11 Ronson, T. (21 October 2026). The Meter Starts Before Hello. Pertlink Viewpoint; Hospitality Net. Includes Pertlink's 18 October agent booking test.

12 Ronson, T. (22 October 2026). Read the Meter. Pertlink Viewpoint; Hospitality Net.

13 Ronson, T. (24 July and 5 August 2026). Who Controls the Controls? and Part Two. Hospitality Net.

The worked example of 1,000 requests and 100 bookings is Dunaeva's. Quotations from McKinsey's experts are from *What does AI really cost?*; the value curve and its stress test are McKinsey's. TCPG, TCPB, HXO, the Decision Envelope, the AI Override Rate, the scenarios, the five-night example, the routing ladder and the role playbooks are Pertlink's.

RIGHTS AND PERMISSIONS

© 2026 Pertlink Limited. You may copy and share this playbook unchanged, with credit to Pertlink, for non-commercial purposes. All other rights reserved.

Quotations, figures, data and illustrations from third parties remain the property of their respective owners. They appear here with attribution and links for comment, review and education, and are not covered by the notice above. Where a graphic adapts another creator's work, the original creator is credited beside it.

Company, product and publication names are trademarks of their respective owners. Mentioning them does not imply endorsement of, or by, Pertlink.

This playbook is general information for education, not financial, legal, or procurement advice. Figures are worked examples, so check current prices and terms with your providers. Links were checked in October 2026.

Every effort has been made to keep this playbook accurate, but errors may remain, and Pertlink cannot accept responsibility for decisions made on the strength of it.

If you hold rights in material used here and believe it is used incorrectly, contact Pertlink at terence@pertlink.net and we will correct it promptly.

PERTLINK
PUTS THE SI INTO HOSPITALITY

About Pertlink

Pertlink Limited is a boutique hospitality technology and AI advisory consultancy, founded in Hong Kong in 2000 and working across the Philippines and Asia-Pacific.

About the author

Terence Ronson is Founder and Managing Director of Pertlink and Chair of the HFTP AI Collective's AI Education & Training Subcommittee. An inductee of the HFTP International Hospitality Technology Hall of Fame and the CHTA Hall of Fame, he has spent more than 25 years opening hotels across Asia. He is a founding member of PAIBA, the Philippine AI Business Association, and developed TCPG, TCPB, the Human Experience Orchestrator and the Hybrid Hospitality Framework.

THE METER SERIES

Budget the Meter, Not the Machine Shiji Insights, 19 October 2026

The Meter Starts Before Hello Hospitality Net, October 2026

Read the Meter Hospitality Net, October 2026

Tokenomics in Hospitality Video overview and this playbook

The intelligence may be artificial. But the experience is human.

 © 2026 Pertlink Limited