

LIAMOL HOSPITALITY INTELLIGENCE

Luxury Hospitality AI Adoption Survey

An executive assessment of AI maturity, guest experience and responsible investment in luxury hospitality.



**Advance capability.
Preserve the art of hospitality.**

SEPTEMBER 2026 | EXECUTIVE SURVEY FRAMEWORK

No primary survey was conducted for this publication. Includes a proposed assessment model and clearly labelled illustrative scoring; no industry adoption rates are claimed.

01 / LEADERSHIP BRIEF

Executive perspective

The value of AI in luxury hospitality depends on the quality of the service it enables.

Luxury hospitality combines personal recognition, discretion and dependable execution. AI can support those qualities by making knowledge easier to find, anticipating operational needs and helping colleagues respond with context. It can also weaken them when an automated interaction makes an incorrect promise, exposes a private preference or removes human judgement from an emotionally important moment.

This report offers a survey-style assessment of that management challenge. Its recommendations are analytical judgements, not findings from hotel respondents. Public references provide selected examples and control guidance; the maturity model, prioritisation tools and roadmap are proposed Liamol frameworks.

Executive priority	Practical implication
Start with a service outcome	Define the guest or colleague problem and measure the complete workflow, including correction and escalation.
Distinguish use from maturity	A purchased feature is not proof of sustained adoption, safe operation or economic value.
Build on trusted information	Fix content ownership, guest identity and interface reliability before expanding autonomous actions.
Match autonomy to consequence	Drafting an answer and changing a reservation require different permissions, testing and approval.
Fund evidence before scale	Release investment when a bounded pilot demonstrates quality, ownership, controls and credible economics.

THE BOARD QUESTION

Where will AI help colleagues deliver more attentive service, and what evidence would make us stop, redesign or expand its use? Establish that decision rule before selecting a vendor.

02 / PUBLICATION APPROACH

Scope, evidence and reading guide

For owners, operators, general managers, commercial leaders and technology teams. The unit of assessment is a defined hotel or operating business, with portfolio policy recorded separately. A resort, city hotel and branded residence may require different use cases even when they share the same technology.

Read for	Pages
Executive direction and evidence boundaries	2-3
Maturity and use-case priorities	4-5
Guest, operational, analytical and commercial applications	6-9
Generative / agentic AI and workforce	10-11
Data, architecture, governance and security	12-15
Vendors, barriers and investment	16-17
Maturity scoring and illustrative example	18-19
Roadmap, outlook and recommendations	20-21
Reusable executive survey and evidence record	22-23
Sources and publication notes	24

What the report can and cannot establish

Public context: numbered citations identify a company announcement or published technical guidance. A product announcement shows a stated application, not independent evidence of results or sector-wide adoption.

Proposed framework: all score anchors, decision gates and recommended practices are tools for an internal review, not a validated industry benchmark.

Illustrative example: the worked score on page 19 is fictional. It is not a surveyed property, percentile or prediction.

No respondent dataset, sampling frame, fieldwork or empirical adoption distribution is available for this publication. Accordingly, no market percentages, invented quotations or purported statistical findings are presented. The title identifies the assessment subject and survey instrument, not completed primary research.

Evidence reviewed on 8 September 2026. Use the questionnaire to collect actual evidence before describing an organisation or cohort as assessed. Any percentages in promotional artwork require their own source or an illustrative label.

03 / BEYOND THE PILOT

Understanding adoption maturity

Record adoption status and maturity separately. Adoption describes where a use case is deployed. Maturity describes whether the organisation can operate it reliably, demonstrate its value and control its consequences. A small, well-governed pilot can be more mature than a broadly available but unmanaged tool.

Adoption status	Evidence required
Not in use	Confirmed absence; record whether deliberately excluded or not yet considered.
Exploring	Problem definition, accountable sponsor and evaluation activity.
Pilot	Bounded users, property scope, dates, test plan and a decision gate.
Production	Service owner, support, monitoring and approved operating procedure.
Scaled	Verified deployment and usage across the declared scope; local exceptions remain visible.

A capability progression, not a technology ranking

Absent: the relevant practice is confirmed missing.

Ad hoc: individual activity exists with limited repeatability.

Defined: ownership, procedures and bounded delivery are documented.

Managed: controls, usage and outcomes are monitored across the stated scope.

Improving: sustained evidence informs changes and failure tests are repeated. These levels become scores 0-4 in the model on page 18.

Assess each hotel against its intended service model. A deliberate choice to retain human concierge service is compatible with high maturity if the decision is evidenced and the supporting data and workflow controls are strong. Autonomy is a design choice; it is not the highest maturity level by default.

ASSESS THE REAL OPERATING UNIT

Ask what happens on a busy evening shift, when a guest changes plans or when a source system is unavailable. Demonstrations of routine happy-path behaviour are insufficient.

04 / SELECT FOR VALUE AND FEASIBILITY

Priority use-case portfolio

Use this qualitative shortlist to open an executive discussion. The suggested sequencing is a framework judgement, not a ranking of market adoption or expected returns. Confirm local volumes, data sensitivity, integrations and service constraints before committing.

Use case	Potential value / evidence	Suggested starting posture
Staff knowledge assistant	Faster access to approved procedures; track correct answers and total resolution time.	Read-only; source-linked answers and a named content owner.
Guest message drafting	More responsive, consistent communication; review corrections and guest complaints.	Colleague approves every outgoing message during pilot.
Demand and service forecasting	Better resource preparation; compare forecast error with a simple baseline.	Advisory output with manager review.
Revenue decision support	Explain demand signals and offer choices; measure net contribution.	Revenue manager approves rate or inventory changes.
Personalised recommendations	More relevant stays and ancillary offers; test guest acceptance without pressure.	Use declared preferences; confirm availability.
Agentic service coordination	Reduce handoffs for eligible requests; track complete and correct fulfilment.	Bounded actions only after integration and recovery tests.

A practical selection test

Score the business problem qualitatively for service impact, frequency, data readiness, integration complexity and consequence of error. Prefer a high-frequency workflow with clear ownership and reversible outputs. Reject a use case if its value relies on unavailable data, intrusive profiling or an unstaffed escalation queue.

FIRST PORTFOLIO

Select one colleague-facing pilot and one analytical pilot. Keep each small enough to compare with current practice, but representative of real shifts, languages and exceptions. A portfolio of disconnected experiments creates support cost without reusable capability.

05 / RECOGNITION WITH DISCRETION

Personalised guest experiences

Begin with information a guest has intentionally shared and that colleagues are authorised to use. The service aim is relevance: an appropriate room preference, a useful itinerary or an accessible communication channel. Avoid surprising guests with inferred personal traits or treating an old preference as a permanent instruction.

Journey moment	Application and service safeguard
Before arrival	Draft an itinerary from stated interests and current property content. Verify hours, capacity, prices and accessibility before confirmation.
Arrival	Summarise relevant preferences for the authorised host. Distinguish guest-confirmed requests from unverified notes and model suggestions.
During the stay	Recommend dining, spa or local experiences with live availability checks. Offer an easy handover to a colleague.
Service recovery	Prepare a factual case summary. A manager owns empathy, compensation and any promise that creates a commitment.
After departure	Draft follow-up using permitted data and channel rules. Correct or remove inaccurate preferences through a controlled workflow.

Measure service quality alongside convenience

Track eligible request resolution time, factual errors, colleague correction time, handover completion and guest-reported usefulness. Segment by language and request type. An answer can sound polished while giving an incorrect opening time; language quality and factual accuracy require separate checks.

In March 2026 Hilton announced a generative AI Planner for hotel discovery and stay planning, described as a beta at launch [1]. This is a public example of travel discovery, not evidence of luxury-segment penetration or realised guest satisfaction. The implication for luxury operators is to evaluate discovery and fulfilment as distinct capabilities.

SERVICE DESIGN RULE

Guests should be able to reach a person without repeating the entire conversation. Transfer only necessary context, make the handover visible and verify that someone has accepted responsibility.

06 / CAPACITY THAT IMPROVES SERVICE

Operational efficiency

Target the work around service delivery: searching for procedures, summarising shift notes, classifying requests and identifying unusual demand. AI creates value only if the complete task becomes better or less burdensome. Counting generated replies or model calls does not measure operational efficiency.

Workflow	Useful application	Required boundary
Housekeeping	Forecast workload and suggest assignment priorities.	Supervisors validate room status, skills and achievable workloads.
Engineering	Summarise fault history and suggest diagnostic steps.	Qualified staff approve safety-related work; preserve equipment interlocks.
Service desk	Classify and route guest requests with context.	Urgent cases follow explicit rules and a staffed escalation route.
Food and beverage	Support demand forecasts and purchasing analysis.	Validate event demand, stock constraints and food-safety procedures.
Administration	Draft reports and extract routine document fields.	Reconcile amounts and source records before posting or payment.

Define a credible operating baseline

Measure time from request receipt to verified completion, including waiting, review and rework. Compare similar volumes, occupancy, shifts and task types. Keep the existing process available during the pilot and document simultaneous changes in staffing or operating procedures.

Net capacity released = baseline task hours minus assisted task hours, including checking, corrections and ongoing support. Released hours become a cash saving only when an actual cost is avoided; otherwise, state how the time is redeployed to guest service. Pair any productivity gain with defect rates, missed requests and staff workload.

SCALE CONDITION

An operations leader must accept the exception workload. If routing is fast but requests wait in an unattended queue, the workflow has not improved.

07 / FROM ANALYSIS TO ACCOUNTABLE ACTION

Data-driven decision making

AI can help managers explore trends, summarise variance and test hypotheses. Trust depends on stable definitions and traceable calculations. Keep financial and operational measures in governed analytical systems; use language models to explain verified outputs rather than inventing calculations from incomplete context.

Decision	Measure and validation
Demand planning	Compare forecast error with a seasonal or recent-history baseline. Use time-based holdouts and prohibit future information in training inputs.
Guest sentiment	Review a multilingual sample against human labels. Separate service topics from inferred emotion and avoid treating review authors as all guests.
Asset intervention	Record alert precision, missed failures and maintenance outcome. Distinguish a prediction from a technician-verified diagnosis.
Management reporting	Reconcile occupancy, average daily rate and revenue per available room to the authorised reporting source. Display scope and dates.

Make the decision trail reviewable

Each material recommendation should show the question, data window, source, exclusions, method, uncertainty and accountable decision-maker. If the system cannot retrieve an authoritative answer, it should state the gap. Managers should be able to inspect the underlying report or record rather than relying on a fluent explanation.

Correlation does not show that AI caused an improvement. Marketing changes, seasonality, refurbishment or channel mix may explain a result. Where practical, use comparable holdout groups or a staged rollout; otherwise, label before/after evidence as observational and record likely confounders.

EXECUTIVE REVIEW

Use one page per decision: what changed, why it matters, confidence, proposed action and the consequence of being wrong. Escalate uncertainty when it could change the action.

Survey prompt: Which recurring management decision has an AI input, and can its latest recommendation be reproduced from retained source data?

08 / PROTECT CONTRIBUTION AND TRUST

Revenue and commercial applications

Separate predictive pricing and forecasting from generative communication. A forecast may inform a revenue management system; a language model may explain the forecast or draft an offer. Neither should be assumed to hold authority to change rates, terms or inventory.

Application	Executive test
Revenue management support	Does the recommendation account for capacity, restrictions, group business and displacement? Retain approved floors and commercial controls.
Ancillary offers	Is the offer relevant, available and deliverable? Measure incremental contribution after fulfilment cost, not only acceptance.
Direct-channel content	Are room attributes, inclusions, cancellation terms and accessibility details accurate across languages and channels?
Sales proposals	Can staff trace rates and terms to approved sources? Review commitments before a proposal reaches a client.
Marketing analysis	Are attribution and audience definitions transparent? Use permitted data and avoid sensitive inferences about individual willingness to pay.

Measure incremental value conservatively

Compare eligible treatment and holdout groups where feasible, accounting for cancellations, channel costs, discounts and displaced revenue. Track conversion together with net contribution, guest complaints and offer errors. A higher average booking value can reflect a different guest mix rather than an AI effect.

For a pilot business case, separate incremental contribution, avoidable operating cost and non-cash capacity. Subtract implementation, integration, licences, model usage, evaluation, review and support. Do not count the same labour hours both as cost savings and additional commercial capacity.

COMMERCIAL CONTROL

Require explicit approval for changes to price, inventory, deposits, refunds and contractual terms until the relevant workflow has demonstrated bounded authority, auditability and tested recovery.

09 / DIFFERENT CAPABILITIES, DIFFERENT AUTHORITY

Generative and agentic AI

Generative AI produces content such as text or summaries. **Predictive AI** estimates an outcome such as demand. **Agentic AI** combines reasoning with tools to pursue a task through multiple steps. These categories can overlap. A conversational interface alone does not prove that a system can safely complete a hotel transaction.

Authority level	Hotel example	Minimum operating condition
Retrieve / draft	Find a procedure or draft a reply.	Approved sources, access filtering and review appropriate to the task.
Recommend	Suggest an itinerary or service action.	Check constraints and distinguish advice from confirmed availability.
Prepare for approval	Assemble a reservation amendment.	Show exact proposed changes and consequences to an authorised approver.
Execute within bounds	Create a permitted service ticket.	Narrow permissions, validated inputs, audit trail and duplicate prevention.
Coordinate across tools	Book and reconcile several services.	Step-level state, transaction checks, recovery and clear stop conditions.

Design the failure path before the happy path

A spa booking agent should verify guest identity and availability, prepare the booking, obtain required approval, execute once and confirm the system record. A timeout must trigger a status check before retrying; a partial booking needs reconciliation. The model must not infer success from a plausible-looking reply.

Treat retrieved documents, messages and web content as untrusted inputs. Enforce permissions and action validation outside the model. Prompt injection and excessive agency are documented application risks [4,5]. A prompt telling an agent to behave safely is not an access-control mechanism.

AUTONOMY GATE

Grant only the tools, data and transaction scope needed. Stop on identity ambiguity, conflicting records or unavailable confirmation. Preserve a staffed manual route and an immediate way to disable execution.

10 / SKILLS, CONFIDENCE AND SERVICE JUDGEMENT

Workforce augmentation

Position AI around clearly defined tasks, with colleagues involved in design and evaluation. Knowledge of a property, an individual guest and the emotional context of a request remains important even when a system can produce a polished answer. A rollout succeeds when staff know when to use the tool, when to question it and how to obtain help.

Role	Learning priority	Evidence of readiness
Frontline colleagues	Verify sources, protect guest data and transfer exceptions.	Observed scenario practice across shifts and languages.
Managers	Interpret uncertainty, review overrides and monitor service impact.	Documented decisions using pilot evidence.
Content owners	Maintain approved procedures, provenance and review dates.	Assigned update process and stale-content checks.
Technology / security	Evaluate models, control tools and investigate incidents.	Completed failure tests and exercised recovery.
Executives	Set risk appetite, funding gates and accountability.	Approved use-case charter and stop/scale criteria.

Measure adoption without rewarding blind use

Define eligible users and tasks before measuring activity. Track successful assisted tasks, override reasons, confidence in using the tool and time spent correcting output. Low usage may reflect poor workflow fit or a reasonable preference for direct service; investigate before treating it as resistance.

Preserve manual competence through practice. Avoid making unsupported AI judgements the sole basis of individual performance assessments or staffing decisions. Explain what usage information is collected and involve appropriate workforce representatives where relevant.

CHANGE PLAN

Appoint a respected operational champion on each pilot shift. Give colleagues a short feedback route, publish fixes and recognise sound judgement when someone rejects an inappropriate recommendation.

Survey prompt: Can a colleague demonstrate an incorrect answer, report it and complete the task safely without the tool?

11 / TRUSTED CONTEXT BEFORE MODEL CAPABILITY

Data readiness

The most useful AI context is often ordinary operational information: current opening hours, approved service standards, accurate reservation state and guest preferences with clear provenance. Prioritise data required by the chosen use case rather than a large, undirected data consolidation programme.

Readiness dimension	Evidence to inspect
Ownership and definition	Named owner for each source; agreed definitions for guest, stay, room, revenue and service completion.
Accuracy and identity	Measured duplicate profiles, reconciled bookings and documented identity matching. Ambiguous guests remain unresolved.
Freshness and provenance	Source, effective date and update frequency retained; superseded policies removed from retrieval.
Access and purpose	Role- and property-appropriate access; only data necessary for the task is available to the AI application.
Lifecycle and quality	Retention, correction and deletion processes cover source systems, indexes, caches and logs. Quality exceptions have owners.

Start with a controlled knowledge collection

For a staff assistant, publish a limited set of approved documents with owner, version and review date. Test whether answers cite the correct source, refuse out-of-scope questions and respect access restrictions. An index must inherit document permissions; copying restricted content into a search layer must not broaden access.

For guest-facing applications, distinguish stable policy from live transactional data. A static brochure can describe the spa; it cannot confirm today's appointment availability. Use authorised system queries for current state and mark stale or missing information explicitly.

READINESS GATE

Do not launch a use case whose critical answer depends on a source with no owner, unreliable identity matching or uncontrolled access. Repair that specific dependency and retest the workflow.

Survey prompt: Can the team trace one AI answer to its exact source, access rule and effective date?

12 / A CONTROLLED PATH FROM CONTEXT TO ACTION

Integrations and platform architecture

A suitable architecture separates the user experience, AI orchestration and systems of record. The diagram is a conceptual design, not a required product stack. A consolidated suite and a modular platform can both work if ownership, reliability and portability are demonstrated.

EXPERIENCE

Guest channels / colleague workspace / visible human handover

AI CONTROL LAYER

Model routing / approved knowledge / evaluation / policy checks

INTEGRATION AND ACTION

Identity / scoped tools / validation / approval / transaction logs

SYSTEMS OF RECORD

PMS / CRS / CRM / POS / revenue / service operations

Conceptual framework. Security, observability and recovery span every layer. The AI layer does not replace the authoritative hotel record.

Maintain an interface catalogue with business owner, data objects, permissions, service objectives and change responsibilities. Test unavailable systems, rate limits, duplicate requests, cancelled stays and delayed events. Confirm business completion in the destination record rather than relying on an API response alone.

Use durable task state, unique request identifiers and reconciliation for actions with side effects. Apply approval checks outside the language model. Preserve data export, knowledge portability and a model-replacement path so a supplier change does not require rebuilding every workflow.

ARCHITECTURE GATE

Demonstrate the full request-to-fulfilment chain and its manual fallback. A successful model demonstration without the hotel integrations is not production acceptance.

PMS: property management; CRS: central reservations; CRM: customer relationship management; POS: point of sale.

13 / GIVE EVERY USE CASE AN OWNER

Governance and accountability

Governance should be proportionate to the consequence of an error. A staff draft and an autonomous refund do not require identical controls. Maintain an AI use-case register covering purpose, owner, users, data, suppliers, authority level, evaluation results, approval status and retirement plan.

Accountable role	Decision responsibility
Executive sponsor	Approves intended value, risk tolerance, investment and stop/scale decisions.
Operational owner	Owns the service procedure, exception queue, staffing and guest outcome.
Data / technology owner	Maintains sources, interfaces, access, monitoring and change control.
Privacy / security leads	Review data use, threat exposure, testing and incident procedures.
Finance / commercial	Validate benefit attribution, cost assumptions and financial authority.

An operating cycle, not a one-time sign-off

NIST organises AI risk management around Govern, Map, Measure and Manage [2]. Use that structure to assign accountability, understand the service context, evaluate failure modes and act on results. Its generative AI profile provides additional risk considerations [3]. These voluntary references inform the approach; they do not certify a hotel or this scoring model.

Before launch, approve an evaluation set covering routine requests, difficult languages, incorrect premises, unavailable data and attempts to cross permissions. Define quality and incident thresholds locally. Repeat evaluations when models, prompts, sources, tools or permissions change; retain the version and decision record.

STOP CONDITIONS

Suspend the affected workflow after unauthorised disclosure, an unapproved financial action or loss of trustworthy transaction state. Assign an incident lead, preserve evidence, restore the manual process and approve restart only after remediation and retesting.

14 / PRESERVE GUEST TRUST

Privacy and cybersecurity

Treat guest conversations, preferences and operational access as sensitive assets. Determine applicable requirements for each jurisdiction and deployment. The controls below are a management review framework; they are not a complete legal assessment. ICO guidance addresses AI and data protection [6]; verify current local obligations with the responsible advisers before launch.

Risk	Required design and evidence
Excessive personal data	Document purpose, applicable lawful basis and necessity; minimise model inputs and assess whether an impact assessment is required.
Cross-guest or cross-property disclosure	Test access enforcement in retrieval and tools. A guest identity claim in a prompt must not grant access.
Prompt injection	Treat external text as untrusted. Separate it from instructions, restrict tools and test indirect attacks embedded in documents.
Unapproved action	Apply least privilege, transaction limits and explicit approval for consequential changes. Log who authorised the action.
Supplier exposure	Confirm retention, training use, subprocessors, location, incident support and deletion behaviour in contractual terms.
Outage or compromise	Maintain isolation, credential revocation, tested restoration, manual service continuity and incident escalation.

OWASP describes how excessive functionality, permissions or autonomy can turn model errors into damaging actions [4]. Its prompt-injection guidance supports layered defences [5]. Test controls against concrete hotel scenarios, such as a malicious itinerary attachment requesting guest records or a forged message trying to obtain a refund.

Use redacted or synthetic test records wherever possible. Restrict and retain logs deliberately: debugging traces can contain personal information or secrets. Provide a route to correct inaccurate data and explain AI involvement where appropriate to the interaction and applicable rules.

NO SINGLE CONTROL IS ENOUGH

Combine access enforcement, data minimisation, action validation, monitoring and human escalation. Passing one demonstration does not establish continuing security.

15 / BUY AN OPERATING CAPABILITY

Vendor and platform strategy

Compare suppliers against the same hotel scenarios and acceptance criteria. A capable model can still sit inside a weak product, and a strong product can be unsuitable for a property's integration constraints. Evaluate the complete service, including human review, implementation and exit.

Evaluation area	Evidence to request
Hotel workflow fit	Demonstrate room moves, group bookings, multilingual requests, cancellations, exceptions and colleague handover.
Quality and control	Show source traceability, evaluations, version changes, approval controls and configurable action limits.
Data and security	Provide data-flow documentation, tenant isolation evidence, role permissions and incident responsibilities.
Integration and operations	Verify supported APIs, sandbox access, reconciliation, service support, failure handling and named owners.
Economics	Price usage, connectors, implementation, evaluation, storage, support, overages and internal review effort.
Exit and continuity	Demonstrate export of owned data and content, deletion confirmation and a transition plan for live workflows.

Choose a sourcing posture deliberately

Embedded capability: may reduce integration effort but can limit cross-system scope or portability.

Specialist application: may fit a service need well but adds supplier and interface dependencies.

Shared platform: can support reusable controls and knowledge across use cases, but requires stronger internal product ownership. Validate each claim in the actual environment.

CONTRACT AND ACCEPTANCE

Tie acceptance to observable hotel outcomes and controls. Agree responsibility for model changes, data incidents, service continuity and failed transactions. Do not accept a roadmap promise as a delivered production feature.

Survey prompt: Could the hotel replace the model or supplier while retaining its approved knowledge, records and service procedures?

16 / FUND THE CONSTRAINTS THAT MATTER

Barriers and investment priorities

Treat barriers as hypotheses to investigate, not assumed industry findings. A problem described as low AI adoption may actually be unreliable source data, limited devices on shift, an unclear process owner or an integration contract that prevents the required action.

Observed barrier	Investment response	Evidence before release
Fragmented information	Assign owners and repair critical source content.	Traceable, current answers on the pilot test set.
Legacy integrations	Instrument and stabilise the required workflow.	Successful retries, reconciliation and outage recovery.
Limited staff confidence	Provide task-based practice and shift support.	Demonstrated verification and safe escalation.
Uncertain value	Establish baseline and bounded comparison.	Quality-adjusted results with complete costs.
Unclear authority	Agree permissions and decision ownership.	Approved action limits and tested enforcement.
Supplier dependence	Negotiate portability and rehearse transition.	Usable export and documented continuity plan.

A staged allocation model

First fund the minimum data, security and integration work required for chosen use cases. Next fund evaluation, operational training and a bounded pilot. Reserve expansion funding for evidence of sustainable service quality and positive economics. Avoid fixed allocation percentages until the property-specific business case is understood.

Annual net benefit = verified incremental contribution + genuinely avoidable costs - recurring operating costs. Show one-time implementation separately. Test a conservative case with lower usage, higher review time and higher supplier cost. Capacity released without a cost reduction belongs in a service-capacity measure, not cash benefit.

INVESTMENT DECISION

Expand, redesign or stop. A pilot should end with one of those decisions and its evidence. Continuing indefinitely without a decision converts experimentation into unmanaged operating cost.

17 / PROPOSED INTERNAL ASSESSMENT MODEL

AI maturity scoring

Assess eight equally weighted domains. For each, score three criteria: **A. ownership and documented practice; B. deployment and control; C. measurement and improvement.** Use the domain evidence on the indicated pages and the anchors below. This proposed model is not externally validated or a certification.

Domain	Evidence pages	Weight
Strategy and value	2, 5, 17	12.5%
Guest and commercial application	6, 9	12.5%
Operations and workforce	7, 11	12.5%
Data and analytics	8, 12	12.5%
Architecture and integration	10, 13	12.5%
Governance and evaluation	14	12.5%
Privacy, security and resilience	15	12.5%
Vendor and delivery capability	16, 20	12.5%

0: confirmed absent. **1:** ad hoc activity, inconsistently applied. **2:** documented owner and repeatable practice in a bounded scope. **3:** managed across the declared scope with monitoring and exception ownership. **4:** sustained outcomes, documented improvements and repeated failure testing. Award the highest fully evidenced anchor for each criterion.

Domain D = (A + B + C) / 3. Overall index = sum of the eight domain scores / 32 x 100. Round only the final result to one decimal. Proposed bands: **0-<25 Initial; 25-<50 Developing; 50-<75 Managed; 75-100 Improving.** The index measures assessed capability, not an adoption percentage or peer percentile.

Record unknown evidence as **U**, not zero. Coverage = assessed criteria / 24. If any criterion is unknown, publish coverage and index bounds using 0 and 4 for unknowns; withhold a single index and band. Do not reweight missing or inapplicable domains; narrow the scope explicitly and retain a separate profile.

Use a declared 90-day evidence window as a starting convention; retain longer records where needed to demonstrate sustained outcomes. Two reviewers resolve differences against evidence. Any unresolved critical access, privacy, action-control or recovery gap blocks dependent scaling regardless of score.

18 / FICTIONAL WORKED EXAMPLE

Illustrative maturity profile

The values below describe an invented property solely to demonstrate calculation. They are not survey responses, industry averages or a forecast. All 24 criteria are assumed assessed; the evidence descriptions are fictional examples.

Domain	A / B / C	Domain score / 4
Strategy and value	3 / 3 / 2	2.67
Guest and commercial	2 / 2 / 1	1.67
Operations and workforce	3 / 2 / 2	2.33
Data and analytics	2 / 2 / 1	1.67
Architecture and integration	3 / 2 / 1	2.00
Governance and evaluation	2 / 2 / 1	1.67
Privacy, security and resilience	3 / 2 / 2	2.33
Vendor and delivery	3 / 2 / 2	2.33

ILLUSTRATIVE RESULT: 52.1 / 100

The criterion total is 50 out of 96 possible points: $50 / 96 \times 100 = 52.1$. Coverage is 24 / 24. The proposed band is Managed. Displayed domain scores are rounded; the index uses the original criterion values.

Read the profile before the headline

In this fictional scenario, procedure ownership is stronger than measurement. The immediate priorities would be to improve source quality, establish outcome baselines and test end-to-end recovery. A Managed overall band does not imply that every domain is managed; several are still developing.

Unknown-evidence variation: if the architecture C score of 1 were unknown, the known total would be 49. Coverage would be 23 / 24. The lower bound would be $49 / 96 \times 100 = 51.0$ and the upper bound $53 / 96 \times 100 = 55.2$. Publish the interval, not a single score or band.

Framework safeguard: an untested action permission or recovery control would still block the affected deployment. A numerical average cannot authorise a use case.

19 / A GATED PATH TO PRODUCTION

Implementation roadmap

This is an indicative sequence, not a promised delivery schedule. The starting plan spans roughly one year; procurement, integration and operating conditions may require a different pace. Move between phases when evidence passes the gate, not when a calendar date arrives.

Phase / lead	Deliverables	Exit gate
Days 1-30 Sponsor + operations	Define scope; inventory AI use; map two workflows; establish baseline and data/control gaps.	Approved charter, owners, eligible tasks and stop criteria.
Days 31-60 Data + technology	Prepare sources and interfaces; configure limited access; agree vendor terms and evaluation set.	Verified permissions, source quality and manual fallback.
Days 61-90 Pilot owner	Run bounded trial across representative shifts; record usage, quality, costs and exceptions.	Evidence supports expand, redesign or stop.
Months 4-6 Operations + IT	Harden support, monitoring, training and change control; extend only approved scope.	Production owner accepts workload and tested recovery.
Months 7-12 Portfolio sponsor	Replicate proven workflows; adapt locally; reassess value and supplier dependence.	Sustained benefits and controls confirmed per property.

Pilot acceptance record

Before launch, set a minimum quality threshold, acceptable review burden, service-response objective and incident triggers for the particular task. Define who measures each and how many cases or operating cycles are sufficient for the decision. Avoid choosing a success threshold after seeing results.

Run routine and failure scenarios: ambiguous identity, missing source, contradictory policy, unavailable tool, repeated request and colleague handover. Record false answers, unsafe suggestions and incomplete actions, including cases where the system declines. Compare with the current process using the same definitions.

RELEASE PACKAGE

Retain the business case, test results, source and model versions, support rota, action permissions, rollback procedure and named sign-off. Repeat acceptance checks after material changes.

20 / PREPARE FOR SEVERAL FUTURES

Outlook and executive recommendations

The scenarios below are planning assumptions, not predictions or empirical survey findings. Their purpose is to identify investments that remain useful if adoption develops differently from expectations.

Plausible scenario	Executive response
AI is embedded in core hotel software	Review each new feature through the same register and access controls; do not assume existing vendor approval covers new AI behaviour.
Guests arrive through AI planning channels	Maintain accurate, consistent property information and verifiable offers; examine how discovery becomes a confirmed booking.
Agents coordinate more transactions	Strengthen tool permissions, approvals, reconciliation and cross-supplier incident ownership before increasing autonomy.
Trust, cost or quality constrains expansion	Retain a useful human service path and focus on bounded colleague assistance with demonstrated value.

Five decisions for the next executive meeting

- 1. Choose the service ambition.** Name where personal attention must increase and where automated assistance is appropriate.
- 2. Approve a small portfolio.** Select owned, measurable workflows and establish explicit conditions for stopping.
- 3. Fund the dependencies.** Repair critical data, identity, integration and content gaps within the chosen scope.
- 4. Set authority boundaries.** Document who approves messages, financial actions and guest commitments; test enforcement.
- 5. Review evidence regularly.** Examine service quality, colleague experience, economics and incidents together, then expand only what works.

THE LONGER-TERM ADVANTAGE

Build a repeatable ability to introduce useful AI while maintaining trusted information, human judgement and dependable service. Those capabilities remain valuable as models and suppliers change.

Recommended cadence: operational review during the pilot; a formal executive gate at its end; then a periodic value and control review tied to the hotel's management cycle.

21 / PART A - SCOPE AND ADOPTION

Executive survey instrument

Use this questionnaire for an internal workshop or a properly designed future study. It is an uncompleted instrument. Record one response per defined operating unit and assessment date; reconcile cross-functional input into a single authorised record.

Field / question	Response format
A1. What is being assessed?	Property or business ID; country; service model; room count; ownership/operating model; included departments.
A2. Who owns this response?	Sponsor and accountable respondent role; contributing functions; assessment date and evidence window.
A3. What is the AI scope?	Named use cases and systems; supplier; eligible properties, users and tasks; internal or guest-facing.
A4. What is the status of each use case?	Not in use / exploring / pilot / production / scaled. Record start date, scope and supporting evidence.
A5. What authority does it hold?	Retrieve/draft / recommend / prepare for approval / execute within bounds / coordinate across tools.
A6. How much is it actually used?	Eligible task count and assisted task count for the same period; define exclusions and source. Do not substitute licences for usage.
A7. Which outcomes are measured?	Baseline, current value, units, dates, denominator, comparison method and known confounders.
A8. What are the main barriers?	Rank three: data / integration / skills / ownership / cost / trust / supplier / other. Supply a concrete example.
A9. What is planned?	Use case, intended outcome, investment range if authorised, owner, dependency and expected decision date.

Interview follow-through

Ask the respondent to demonstrate one completed task and one exception. Compare the demonstration with usage logs, service records and the documented procedure. Record unsupported answers as unknown rather than inferring readiness from confidence or job title.

For research use, obtain an appropriate participant notice and data-handling protocol. Do not publish identifying property information or commercially sensitive responses without the required authority.

22 / PART B - MATURITY AND ACTION

Evidence and assessment record

Complete the three criteria for each domain using page 18. For every score retain the evidence reference, date, assessor and any unresolved issue. An interview can identify evidence to collect; it does not automatically demonstrate managed capability.

Survey domain	Evidence question
B1. Strategy and value	Who owns the use case, which outcome is measured, and what decision did the latest review produce?
B2. Guest and commercial	Can the team demonstrate correct information, approved commitments and a completed guest handover?
B3. Operations and workforce	Are colleagues trained across shifts, exceptions owned and net workload measured?
B4. Data and analytics	Are sources current, access-controlled and reproducible; are decision metrics reconciled?
B5. Architecture and integration	Can actions be traced, reconciled and recovered after timeouts or partial completion?
B6. Governance and evaluation	Are changes tested, authority documented and stop/restart decisions assigned?
B7. Privacy and security	Have access boundaries, data handling, incident response and recovery been tested?
B8. Vendor and delivery	Are support, cost, data terms, portability and acceptance responsibilities evidenced?

Record for each domain: A ___ / B ___ / C ___; evidence IDs ___; window ___; assessor ___; confidence high / medium / low; unresolved risk ___. High confidence requires direct, current operating evidence; medium indicates partial corroboration; low indicates mainly self-report. Confidence does not replace the scoring anchor.

ASSESSMENT CLOSE-OUT

Coverage ___ / 24. Index or bounds ___. Critical risk flag ___. Priority actions ___. Accountable owners ___. Due dates ___. Executive decision: expand / redesign / stop / collect evidence. Next review ___.

If conducting a real multi-property survey, predefine recruitment, eligibility, duplicate handling and analysis. Publish fieldwork dates, achieved sample, non-response, question wording, weighting, missingness and limitations. A convenience sample cannot be assumed representative of luxury hospitality.

23 / REFERENCE REGISTER

Sources and publication notes

Public sources below support the specific context or control concepts cited in the report. They do not supply an AI adoption dataset for this publication. All links were reviewed on 8 September 2026. Company statements remain attributed claims; technical guidance is not evidence of hotel outcomes.

[1] Hilton. Hilton Introduces the Hilton AI Planner, Advancing the Future of Curated Travel Discovery. March 2026.

Public example of a generative AI discovery and planning application; no results or luxury-market adoption rate inferred. [Open source.](#)

[2] NIST. AI Risk Management Framework (AI RMF 1.0), 2023; framework resource page.

Voluntary risk-management structure supporting governance, context, evaluation and response. [Open source.](#)

[3] NIST. Artificial Intelligence Risk Management Framework: Generative Artificial Intelligence Profile. NIST AI 600-1, July 2024.

Generative AI risk considerations; does not validate the proposed maturity score. [Open source.](#)

[4] OWASP. LLM06:2025 Excessive Agency.

Specific 2025 guidance on excessive functionality, permissions and autonomy; cited by edition, not as a claim about the latest ranking. [Open source.](#)

[5] OWASP. LLM Prompt Injection Prevention Cheat Sheet.

Application threat context and layered prevention practices. [Open source.](#)

[6] Information Commissioner's Office. Guidance on AI and data protection.

UK data-protection reference; check current guidance and jurisdiction-specific obligations. [Open source.](#)

Publication integrity

Prepared for the Liamol Resource Library, September 2026. The assessment model, scoring bands, survey instrument and roadmap are proposed frameworks. The worked property profile is fictional. No proprietary fieldwork, respondents, market share, sample size or measured industry finding is claimed. Reassess sources, legal requirements and product capabilities before applying the report to a live deployment.

Liamol Hospitality Intelligence | Luxury Hospitality AI Adoption Survey | Executive survey framework, edition 1.0