

ONE HOTEL, SEVEN VERSIONS

How accommodation becomes understandable,
trustworthy and directly bookable
in AI-mediated travel



The risk is not that AI cannot read the page.
It is that AI can read several incomplete pages
and produce one confident mistake.

Publisher's note

This paper examines how one accommodation property is represented across the modern travel web. It combines official platform documentation, independent and vendor research, a public-source forensic review of Hôtel Mont-Blanc Chamonix, and a critical reading of the supplied white papers. It is not an audit commissioned by the hotel, and it should not be read as an assessment of the hotel's hospitality quality, commercial performance, legal compliance or technical implementation.

One of the seven versions studied is a JSON-LD proposal published in the supplied white paper *Devenez visible pour les IA*. That proposal is analyzed as a third-party representation of the property. It has not been verified as markup currently deployed by the hotel's website. The AI version is likewise based on the Gemini report reproduced in that publication, generated in December 2025. AI outputs, booking interfaces, prices and policies change; every time-sensitive observation in this edition is therefore dated and bounded.

Tamaga's editorial rule is simple: show what is sourced, show what is limited, and make corrections possible. The online Source Room accompanying a public release should carry the evidence register, figure data, methodology, corrections and version history.

Abstract

Accommodation discovery is often discussed as a contest for visibility in generative search. That framing is too narrow. A property no longer exists in one digital place. It exists simultaneously in its official website, direct booking engine, structured-data layer, local-search profiles, online travel agencies, destination and editorial sources, and the synthesized answers produced by AI systems. These versions can be individually plausible and collectively inconsistent.

This white paper names that condition the **Property Truth Gap**: the distance between what a property can actually deliver and what travelers, search engines, intermediaries and AI systems can discover, verify and act upon. Using Hôtel Mont-Blanc Chamonix as a public-source forensic case, the paper compares 16 traveler-critical facts across seven representations. The purpose is illustrative rather than statistical: to expose where identity, room configuration, access, timing, policy and commercial facts can disappear, flatten or contradict each other.

The paper then develops two further concepts. The **Omnichannel Decision Mesh** extends semantic architecture beyond internal links to the wider source environment that supports a travel decision. The **Source-to-Service Chain** distinguishes description, entity resolution, corroboration, comparison, live offer, confirmation, action and recovery. Together, the three concepts reframe AI visibility as a knowledge-governance and service-design problem rather than a new layer of acronym-driven optimization.

The central conclusion is deliberately modest: structured data is important because it makes facts and relationships explicit. It does not make them true, current, authoritative, preferred or bookable. Durable advantage comes from agreement between the public story, structured facts, distributed sources, live offer and human service.

Executive summary

The problem is not invisibility alone

A hotel can be visible and still be represented incorrectly. It can appear in search, maps, an OTA, a destination guide and an AI answer while the decisive detail remains wrong: a checkout time, bed configuration, late-arrival condition, accessibility statement, seasonal facility or cancellation rule. The risk is not simply that an AI system fails to read the official page. The risk is that it reads several incomplete sources and produces one confident synthesis.

The supplied white papers correctly identify a large structured-data implementation gap in hospitality. A 2026 crawl of 121,425 hotel homepages found that 36.3% of reachable sites had no structured data, while only 32.4% of the sites with JSON-LD used a lodging-specific primary type. Decision-relevant properties such as **geo**, **aggregateRating** and **amenityFeature** were uncommon.[5] But these findings establish an implementation problem, not a causal law of AI recommendation. A property without rich markup may still be extracted from ordinary text and third-party sources; a property with perfect syntax may still publish the wrong fact.

Seven versions of one property

This paper separates seven versions that are often collapsed into the phrase "the hotel online":

1. **Owned website** - the facts and claims a traveler can read.
2. **Direct booking engine** - live dates, rates, restrictions and transactional choices.
3. **Structured-data layer** - explicit entities and relationships expressed for machines.
4. **Google and local surfaces** - aggregated identity, amenities, reviews, location and prices.
5. **OTAs and review platforms** - commercial detail, policies, inventory and social proof.
6. **Destination and editorial web** - local context, independent corroboration and story.
7. **AI synthesis** - a compressed comparison assembled from several of the above.

One hotel, seven public and transactional versions

The strategic problem is not that the versions differ. It is that consequential differences are rarely governed.



Figure 1. One hotel, seven public and transactional versions.

Source: Tamaga synthesis.

No version is inherently complete. The official site may contain the richest qualitative explanation but scatter it across seasonal pages. The booking engine may hold the freshest commercial truth but expose little to a crawler. Google may compress a nuanced access description into the label "Accessible." An OTA may provide late-check-in conditions absent from an official FAQ. A destination office may independently confirm a spa or pool but not the room configuration. An AI system may combine all of them - and still recommend a different property.

Three concepts to remember

Property Truth Gap

The distance between operational reality and its public, structured, distributed and synthesized representations.

Omnichannel Decision Mesh

The connected set of pages, profiles, sources, partners, concepts and booking paths that allows a traveler to move from a need to a verified decision without losing meaning between channels.

Source-to-Service Chain

The progression from describing a property to resolving it, corroborating it, comparing it, exposing a live offer, confirming consequential details, completing an action and recovering when something fails.

Ten findings

- 1. SEO did not become obsolete.** Google explicitly states that existing SEO fundamentals remain foundational for its generative Search features and that no special Schema.org markup or new AI file is required.[1]
- 2. Structured data is a representation layer, not ground truth.** It can reduce ambiguity, but syntax cannot certify a claim.
- 3. Accommodation must be modeled as a graph.** Schema.org separates the lodging business, the accommodation unit and the offer; prices and conditions belong to the offer rather than to the abstract hotel or room.[4]
- 4. Different facts move at different speeds.** An address may be stable for years; restaurant opening, inventory, prices and reservation state require different systems and update frequencies.
- 5. The wider web matters.** Large-scale observational research shows that AI hotel answers draw heavily on OTAs, review platforms, brand sites, Wikipedia, social sources and editorial domains, with substantial differences between models.[6]
- 6. Citation is not recommendation, and recommendation is not booking.** Measurement must distinguish entity resolution, candidate inclusion, factual accuracy, link destination, live offer and completed action.
- 7. Direct booking is not automatically better.** OTAs provide reach, comparison, trust, payment and support. Direct channels create value when they offer parity, clarity, distinctive service and an accountable relationship.
- 8. AI may disintermediate or reintermediate.** A conversational interface can send a traveler directly to a property, or it can route the entire transaction through a connected intermediary.
- 9. Agent tools do not guarantee agent use.** An API or MCP tool must still be discovered, selected, called correctly, authorized, confirmed and recoverable.[14]
- 10. The durable advantage is agreement.** The property, page, markup, platform, offer and team should describe the same stay.

What accommodation providers should do first

Do not begin with an AI content factory or a custom agent. Begin with one decision-critical fact - for example late arrival, connecting rooms, step-free access or cancellation - and trace it through all seven versions. Identify the owner, source, conditions and update path. Then repeat.

The implementation sequence in this paper is:

Align the facts. **Prove** the consequential claims. **Distribute** coherent representations. **Connect** live operational truth. **Act safely** with confirmation and recovery.

How to read the evidence

The white paper uses four evidence labels.

ESTABLISHED - supported by official specifications, stable documentation or strong peer-reviewed evidence.

OBSERVED - supported by a documented dataset or public-source audit, but not necessarily causal or generalizable.

ILLUSTRATIVE - a case observation used to expose a mechanism, not a sector-wide estimate.

UNKNOWN - a plausible relationship for which current public evidence is insufficient.

The supplied preliminary papers are treated as research inputs. Their statistics and technical examples are retained only where the original source can be identified or the paper itself is the object of analysis. Unsupported claims - such as the idea that JSON-LD automatically improves recommendations - are not silently carried forward.

Contents

Executive summary.....

How to read the evidence.....

Contents.....

The plausible wrong answer.....

One hotel, seven versions.....

The Property Truth Gap.....

AI discovery is a source-selection system.....

What structured data can and cannot do.....

Model the stay, not the homepage.....

Put each fact in the system capable of keeping it true.....

Build the Omnichannel Decision Mesh.....

The wider destination web as evidence infrastructure.....

Direct booking is not a button.....

The Source-to-Service Chain.....

Measurement without theatre.....

A practical implementation roadmap.....

The agreement is the asset.....

Become the clearest source before trying to become the chosen answer.....

Case methodology and detailed observations.....

Illustrative accommodation JSON-LD pattern.....

Property Truth inventory.....

Data placement matrix.....

References.....

About Tamaga.....

2

4

5

8

9

12

16

22

25

28

31

34

36

40

43

46

49

49

50

52

54

55

56

57

PROLOGUE

The plausible wrong answer

A traveler asks:

Find a quiet independent hotel in Chamonix for two adults and two children in late October. We arrive by train after 20:00 on Sunday. We need two connecting rooms, step-free access, dinner that evening and flexible cancellation. Compare the direct offer with an OTA, identify what is confirmed, and tell us what still requires a human.

The answer engine has enough material to sound decisive. It finds a five-star hotel in the center of Chamonix. It reads that Standard and Superior rooms can connect. It sees an outdoor heated pool, a restaurant, a shuttle, parking and a spa. It sees a 24/7 reception label on some room pages. It finds policies and fees on an OTA. It finds the pool and spa independently described by the destination office. It may even find a long-form report that frames the hotel as a strong urban alternative for a family.

A fluent answer might therefore say:

Hôtel Mont-Blanc is a strong fit. Book connected Standard and Superior rooms, arrive at any time, dine at Le Matafan on Sunday evening, and use the hotel's step-free facilities. The outdoor pool is heated year-round. Book direct for the best service and flexible cancellation.

Every sentence is plausible. Not every sentence is established.

The official room pages say communicating rooms are available **on request**, not guaranteed.[20] The official materials reviewed confirm a reception service and the standard check-in time, but the OTA is the source that explicitly frames check-in as 15:00 to midnight and late check-in as subject to availability.[27] The official restaurant material reviewed makes Sunday lunch explicit, but does not by itself guarantee a late Sunday dinner for the traveler's dates. The destination office confirms a spa open every day and an outdoor heated pool, yet that does not prove the exact October operating conditions of every associated service.[29] Google compresses access into the generic label "Accessible," while the hotel's own information is more specific: two Prestige rooms are adapted for guests with reduced mobility.[20]

The answer is not necessarily false. It is **underqualified**. It turns several kinds of evidence into one promise.

That is the central failure examined in this paper.

AI did not create the contradiction. It exposed a representation system in which:

- marketing copy, policy pages and room pages are maintained separately;
- a booking engine holds live commercial choices behind an application interface;
- structured data may be incomplete, stale or wrong;
- local and intermediary platforms simplify nuance;
- editorial sources add useful context without transaction responsibility;
- AI synthesis removes the boundaries between them.

The problem is not merely visibility. It is the integrity of the path from **source to service**.

CHAPTER 1

One hotel, seven versions

1.1 The property chosen for the forensic case

Hôtel Mont-Blanc Chamonix is useful as an illustrative case because it is not digitally absent. It has a substantial official website, room-category pages, seasonal practical information, a direct booking application, Google Hotels representation, major OTA listings, destination-office pages, editorial coverage and an existing AI report in the supplied white paper. In other words, it has the kind of distributed public footprint that many "AI visibility" guides tell hotels to build.

The question is not whether the hotel exists online. The question is whether the versions agree where the traveler needs certainty.

This is a public-source study. It does not claim access to the hotel's PMS, CRS, internal accessibility documentation, current markup deployment, private booking rules or staff procedures. Where operational confirmation is unavailable, the paper says so.

1.2 Version 1 - The owned website

The owned site is the richest first-party narrative. It describes room size, occupancy, beds, views, services and facilities. The Standard room page states that a Standard and Superior room can connect and labels communicating rooms as available on request.[20] The Duplex Suite page states four-person occupancy, 75 square meters and two king-size beds.[21] Current practical information gives 15:00 check-in, 11:00 checkout, parking and electric-vehicle fees, pet conditions and adapted Prestige rooms.[19]

The official site also demonstrates a common structural problem: the truth is distributed across page families and seasons. A traveler may need to combine a room page, winter practical information, a spa page, a restaurant page and booking-engine output. The website knows much more than any one page proves.

1.3 Version 2 - The direct booking engine

The direct booking path is not merely another page. It is where dates, occupancy, availability, rate plans, payment conditions and cancellation choices should become transactional. In the text-only public audit, the engine was delivered through a JavaScript application and its live contents could not be fully inspected. That limitation is important.

"Not observable to this audit" does not mean "not available to the traveler." It means the booking engine is a separate representation surface with different access, rendering and update mechanisms. A crawler, browser-based agent, mobile user and human researcher may not see the same thing.

The direct engine should therefore be treated as its own version, not assumed to inherit the official site's content or structured data.

1.4 Version 3 - The structured-data layer

The supplied white paper presents a proposed JSON-LD block as a state-of-the-art representation for the hotel.[32] It is a useful demonstration of both the power and danger of structured data.

The proposal correctly attempts to create explicit entities and relationships: hotel, Suite Duplex, restaurant, address, coordinates, facilities and occupancy. But it also contains consequential errors or ambiguities:

- the name is rendered as "Hôtel du Mont-Blanc Chamonix" rather than the current official name;

- the email differs from the official address published by the hotel;
- the check-in value is malformed;
- checkout is stated as 12:00 rather than the current 11:00 published by the official site, Google and OTAs;
- the Duplex bed configuration is stated as "King + 2 single" rather than two king-size beds;
- a ski room is assigned **additionalType: SkiResort**, which confuses an amenity with a destination type.

The proposed code is machine-readable. That is exactly why the errors matter. Structured data can turn ambiguity into explicitness. It can also turn an uncertain or incorrect statement into an explicit, reusable error.

1.5 Version 4 - Google and local surfaces

Google Hotels combines the property identity, location, review-derived themes, amenities, prices and links from multiple sources. Its current listing states check-in at 15:00 and checkout at 11:00. It shows pool, spa, hot tub, restaurant, shuttle, paid parking, pet-friendly and the generic label "Accessible." Google itself notes that it collects hotel amenities from a variety of sources.[26]

That aggregation is convenient for comparison. It is also lossy. "Accessible" does not tell a traveler whether the entrance, reception, lift, bedroom, bathroom, restaurant and route from the station are step-free. The local surface is a powerful candidate-selection interface, not a complete accessibility statement.

1.6 Version 5 - OTAs and review platforms

OTAs often carry operationally useful detail because they need to support booking decisions and customer service. Expedia currently states 15:00 to midnight check-in, late check-in subject to availability, checkout before 11:00, dogs at EUR 25 per night and city tax conditions.[28] Other OTA representations may vary by market, translation and inventory source.

This version can be more explicit than the official site on some policies while being less nuanced on room fit or service context. It also has a commercial incentive and a customer-support obligation that editorial sources do not.

1.7 Version 6 - Destination and editorial web

The Chamonix destination site independently describes the 250-square-meter Clarins spa, a heated pool and outdoor whirlpool, an age minimum of six and daily opening.[29] Editorial guides and specialist travel sources may add heritage, restaurant, location and experience context.

These sources are valuable because they can corroborate and contextualize. They are not necessarily current systems of record for room availability, rate conditions or late arrival. Their authority is claim-specific.

1.8 Version 7 - AI synthesis

The Gemini report included in the supplied white paper is a vivid example of cross-source synthesis.[31] For a family seeking luxury, food, a heated outdoor pool and ski access, it recommended Hameau Albert 1er as the primary choice and presented Hôtel Mont-Blanc as an urban alternative. It successfully extracted connected-room information, pool, spa, shuttle, restaurant and location context from ordinary pages and third-party sources.

That appendix undermines the strongest claim made elsewhere in the same paper: that an AI will leave empty-handed without structured data. The AI clearly found and combined substantial information. But it also illustrates a second truth: extraction is not preference. A hotel can be understood and still lose the recommendation.

1.9 The strategic lesson

The seven versions should not be forced into identical copies. Each has a role:

- the official site explains;
- the booking engine transacts;
- structured data makes entities explicit;
- local surfaces compare;
- OTAs distribute and service;
- destination sources corroborate and contextualize;
- AI synthesizes.

The governance objective is not uniform text. It is agreement on consequential facts, visible ownership of uncertainty and a safe handoff to the system or person capable of confirmation.

CHAPTER 2

The Property Truth Gap

2.1 Definition

The **Property Truth Gap** is the distance between what an accommodation provider can actually deliver and what its public, structured, distributed and synthesized representations allow a traveler or machine to conclude.

The gap has four forms:

Absence - a decision-critical fact exists internally but is not publicly available.

Flattening - a nuanced condition becomes a generic label, such as "accessible" or "family-friendly."

Contradiction - two versions publish incompatible values.

Action failure - the information is understandable, but no safe route exists to confirm, book, modify or recover.

2.2 The illustrative audit

Sixteen traveler-critical facts were reviewed across the seven versions. Each observation was coded as:

- precise or confirmed;
- partial or generic;
- contradicted or erroneous;
- not stated or not observable.

This is not a rating of the hotel. It is a representation map based on public evidence available to the research process in July 2026. The booking engine's live contents, for example, were not fully observable in the text-only audit, so absence in that column is a limitation rather than a claim that the engine lacks the data.

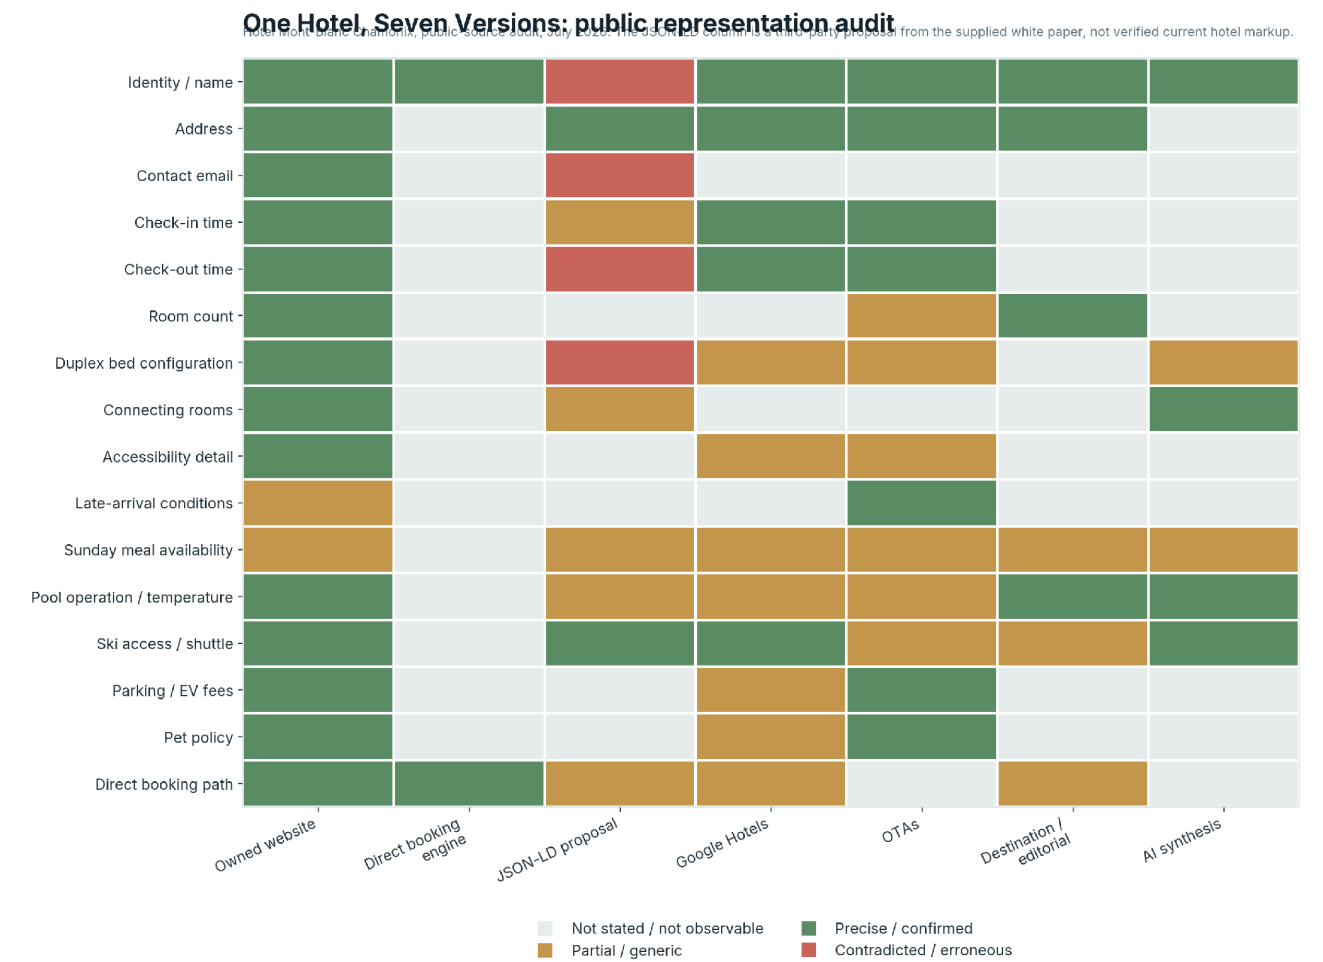


Figure 7. One Hotel, Seven Versions: public representation audit.

Source: Tamaga public-source audit, July 2026. The JSON-LD column analyzes a third-party proposal, not verified deployed markup.

2.3 What the heatmap reveals

The official site is the strongest public source for many stable and experiential facts. It is precise on identity, address, contact, room occupancy, the Duplex bed configuration, communicating rooms, parking, pet policy and direct booking. Yet some decision details remain fragmented or conditional: communicating rooms are on request; Sunday lunch does not establish late dinner; an adapted room does not answer the entire step-free journey.

The direct booking engine is central to commercial truth but comparatively opaque to a static text audit. This is not a defect unique to the case. Modern booking applications often require dates, scripts, cookies, inventory calls and user interaction. Their truth is dynamic by design.

The JSON-LD proposal contains the highest concentration of outright contradictions. This is analytically useful because it shows how easy it is to build an impressive graph from imperfect research. The problem is not JSON-LD. The problem is separating authoring convenience from fact governance.

Google and OTAs cover broad decision needs but often compress meaning. They are strong on check-in, checkout, amenities, fees and commercial conditions. They are weaker at distinguishing "available" from "guaranteed," or broad access labels from precise pathways.

Destination and editorial sources are strongest where independent context matters: spa description, location, reputation and local positioning. The AI synthesis is broad but selective. It extracts enough to compare the hotel, yet does not expose every uncertainty in the underlying sources.

How much of the same hotel survives in each version?

The direct booking engine is deliberately shown as opaque where live JavaScript content could not be inspected; absence here is not proof that the engine lacks the data.




Figure 8. How much of the illustrative case survives in each version?

Source: Tamaga public-source audit, July 2026. Direct booking-engine non-observation is a method limitation, not proof of missing data.

2.4 The most dangerous facts are conditional facts

Identity errors are visible and easy to correct. Conditional facts are harder:

- connecting rooms exist, but allocation is on request;
- late check-in may be possible, but an OTA says it is subject to availability;
- a pool exists, but access rules, operating hours and maintenance can vary;
- a restaurant exists, but a specific Sunday dinner is not guaranteed;
- adapted rooms exist, but suitability depends on the guest's precise access needs;
- a flexible rate may exist, but only for selected dates and room categories.

These facts combine a noun with a condition. Generic markup and generic content tend to preserve the noun and lose the condition.

2.5 The Property Truth Gap is an organizational signal

When versions disagree, the underlying cause is rarely "bad AI." It may be:

- no named owner for the fact;
- separate marketing and operations systems;
- manual re-entry across channels;
- old translations;
- a CMS field that cannot represent the condition;
- an OTA extranet updated independently;
- a dynamic offer incorrectly copied into static markup;
- a destination record with no correction workflow;
- a property change that propagated only to some channels.

The gap therefore measures the distance between expertise and infrastructure. It should be treated as a governance problem before it becomes a marketing problem.

2.6 A simple Property Truth Card

For every consequential fact, maintain:

- canonical statement;
- subject entity;
- applicable conditions;
- system of record;
- source or evidence;
- owner;
- last checked date;
- public wording;
- machine-readable representation;
- channels where it appears;
- guarantee status;
- correction path.

For a small property, this can begin as a spreadsheet. The value comes from ownership and propagation, not from software complexity.

CHAPTER 3

AI discovery is a source-selection system

3.1 The interface changed; the foundations did not disappear

Generative answer surfaces compress the number of visible candidates and synthesize information across sources. They can answer longer, constraint-rich questions without requiring a traveler to open ten tabs. This is a real interface change.

It is not a clean replacement of SEO by GEO. Google states that specific optimization is not required for AI Overviews and AI Mode, that existing SEO fundamentals remain worthwhile, and that websites do not need a special AI file or special Schema.org markup.[1] Crawling, internal links, visible text, page experience, accurate structured data, images, video, Business Profile and Merchant Center data remain part of the same system.

OpenAI likewise describes search visibility in terms of public accessibility and OAI-SearchBot access rather than a proprietary GEO schema.[3]

The useful distinction is not SEO versus GEO. It is between:

- being retrievable;
- being correctly resolved as an entity;
- entering the candidate set;
- being selected for the traveler's constraints;
- being supported by credible sources;
- offering a safe next action.

3.2 The hotel source ecosystem is heterogeneous

A large observational study of 19,579 AI hotel-recommendation runs found that models used markedly different source mixes.[6] In GPT-5.2 runs, Booking.com appeared in 53.9%, Hotels.com in 31.9%, Marriott.com in 30.6%, Wikipedia in 30.0%, Expedia in 28.9% and Tripadvisor in 20.5%. Grok and Perplexity showed very different dependencies.




Figure 5. Domains cited in GPT-5.2 hotel-recommendation runs.

Source: Nicolas Sitter, *AI Hotel Landscape 2026*. Observational model snapshot; not a ranking-factor study.

These figures should not be interpreted as ranking factors. They show that:

- no single source controls the answer;
- intermediaries can influence recommendations even when the final link is direct;
- source preferences differ by model and version;
- public representation must be coherent beyond the official domain;
- visibility tests are snapshots.

The same study reported that 75% to 91% of hotel links in its model outputs pointed to hotel-owned domains even though OTAs were heavily consulted. That finding suggests a direct-referral opportunity, but it does not guarantee that a hotel will be selected or that a direct booking will follow.

3.3 Search traffic is not disappearing, but click behavior is changing

Pew Research Center analyzed 68,879 Google searches from 900 U.S. adults. Users clicked a traditional result in 8% of visits when an AI summary appeared, versus 15% without one; links inside the summary were clicked in 1% of such visits.[7]





Figure 6. Observed click behavior on Google search pages with and without an AI summary.
Source: Pew Research Center, 68,879 Google searches from 900 U.S. adults, March 2025.

The evidence is specific to the study period, population and Google interface. It does not mean every travel query becomes zero-click. It does mean that a property should evaluate representation inside the answer as well as referral traffic after it.

Adobe's analysis of more than eight million visits to U.S. travel sites found AI-source traffic up 194% year over year in May 2026 and 2,215% since October 2024, while conversion remained 28% below non-AI traffic but the gap had narrowed substantially.[8] The absolute share remains much smaller than search for most sites. The rational response is measurement, not panic.

3.4 Structured-data adoption remains weak

The 2026 hotel schema crawl provides the strongest public baseline currently available.[5]



*Derived as the remainder after JSON-LD and no-structured-data categories; non-JSON-LD formats may coexist with JSON-LD in the source study.

Figure 2. Structured-data adoption on reachable hotel homepages, 2026.

Source: Nicolas Sitter, Hotel Schema.org Adoption Study 2026. The 7.9% remainder is derived; formats may overlap in the underlying crawl.

Among 105,002 reachable homepages, 55.8% contained JSON-LD and 36.3% contained no structured data. The remainder used only other formats or fell outside those two categories.

Among JSON-LD implementations, Organization was more common as the primary type than Hotel. Only 32.4% used a lodging-specific primary type.




Figure 3. Primary JSON-LD type used by hotel homepages.

Source: Nicolas Sitter, Hotel Schema.org Adoption Study 2026. Among 58,625 homepages with JSON-LD.

The most frequently present field was **name** at 71%. More decision-relevant fields were much rarer.




Figure 4. Adoption of selected decision-relevant Schema.org properties.

Source: Nicolas Sitter, Hotel Schema.org Adoption Study 2026. Presence does not establish truth, freshness or correct modeling.

The study's custom completeness score evaluates field presence, not truth, freshness or correct relationships. The right conclusion is therefore:

Hospitality has a large machine-representation gap.

The wrong conclusion is:

The missing fields prove that those hotels cannot be understood or recommended by AI.

3.5 Candidate selection is not a schema lookup

A traveler asks for "a quiet hotel for an older parent near the station with dinner on Sunday." The answer may depend on:

- location and transport;
- review evidence about noise;
- room and lift access;
- restaurant hours;
- date-specific availability;
- editorial understanding of the neighborhood;
- budget and cancellation;
- brand or host trust;
- the model's source retrieval and synthesis.

Schema.org can express pieces of this. It cannot perform the entire judgment, and many of the most useful signals are conditional, qualitative or external.

3.6 Visibility should be measured as a vector

A hotel may have:

- high entity accuracy;
- low recommendation frequency;
- strong direct-link share;
- weak policy accuracy;
- excellent source diversity;
- poor booking completion.

Reducing these into one "AI visibility score" conceals the problem that needs repair. Chapter 11 proposes a multidimensional measurement model.

CHAPTER 4

What structured data can and cannot do

4.1 What it can do

Structured data is valuable because it makes a publisher's claims explicit. It can:

- identify the type of entity;
- assign a stable identifier;
- connect a hotel to rooms, offers, restaurants, people and places;
- distinguish property-level and room-level facts;
- represent occupancy, beds, coordinates and selected amenities;
- connect an offer to price, currency, validity and conditions;
- support validation and reuse by systems that consume the vocabulary;
- reduce the need to infer simple relationships from prose.

These benefits existed before generative AI. The new interfaces increase their strategic relevance because more systems now synthesize information across sources.

4.2 What it cannot do

Structured data cannot independently establish:

- factual truth;
- current availability;
- service quality;
- popularity;
- authority;
- traveler fit;
- recommendation preference;
- legal or operational guarantees;
- a successful transaction.

A false checkout time remains false inside JSON-LD. A stale price remains stale. A property can correctly state that it has adapted rooms without proving that a particular traveler can independently navigate every part of the stay.

4.3 Three layers of validation

Syntax

Does the JSON parse? Are the types and properties recognized?

Semantics

Is the entity modeled correctly? Does an offer point to a room? Is a ski room being incorrectly typed as a ski resort? Does the property belong at hotel level or room level?

Truth

Does the value match visible content and current operational reality? Who owns it? When was it checked? Under which conditions does it apply?

Validators are strongest at the first layer. Human and operational governance are necessary for the third.

4.4 Visible-content parity

Google explicitly advises that structured data should match visible text.[1] This is not only a policy issue. Parity allows:

- travelers to inspect the claim;
- editors to notice errors;
- support teams to explain the offer;
- systems to share one governed source;
- corrections to propagate from a visible record rather than a hidden script.

The worst structured-data implementations are hand-authored shadow websites: detailed hidden blocks that are more ambitious than the page and disconnected from the CMS.

4.5 Stable identifiers matter more than decorative completeness

A stable `@id` allows entities to be referenced consistently across pages. For example:

- `https://example.com/#hotel`
- `https://example.com/rooms/duplex/#room`
- `https://example.com/offers/flexible-duplex/#offer`
- `https://example.com/restaurant/#restaurant`

The identifier should persist even when the visual template changes. Names, URLs, languages and channel listings can then be mapped to the same conceptual object.

4.6 `sameAs` is not a link dump

`sameAs` should identify pages that unambiguously represent the same entity. It should not be used to attach every mention, directory or partner. An OTA property page may be an identity reference, but editorial articles and review pages usually have different roles. Treating all external links as equivalent erases the distinction between identity, corroboration and commentary.

4.7 Ratings require provenance

`aggregateRating` is frequently recommended because it appears decision-relevant and is rare in the crawl. But a publisher should not simply copy a third-party rating into owned markup without satisfying platform policies and clearly establishing the source. A rating is not just a number; it has a population, date, scale and owner.

4.8 FAQ markup is not a universal AI tactic

Hotels should publish clear answers to genuine questions. That is good content and service design. It does not follow that FAQPage markup is an AI ranking tactic. Google has substantially restricted FAQ rich-result eligibility, and its generative Search guidance says no special markup is needed.[1]

The best question page is one that exposes a maintained operational answer, not one manufactured to increase schema count.

4.9 The right strategic claim

Structured data should be positioned as:

a public interface to a governed knowledge model.

Not as:

a machine-readable spell that makes the hotel recommendable.

CHAPTER 5

Model the stay, not the homepage

5.1 Three core objects

Schema.org's hotel guidance identifies three core objects:[4]

1. the lodging business;
2. the accommodation unit;
3. the offer to rent that unit under particular terms.

This distinction is essential because travelers do not book an abstract hotel. They book a dated right to use a particular category of accommodation under a particular policy.

Model the stay, not the homepage

Commercial terms belong to the offer. Transaction state belongs to the reservation system.




Figure 9. Model the stay, not the homepage.

Source: Tamaga synthesis based on Schema.org hotel modeling guidance.

5.2 The lodging business

The property entity can carry relatively stable facts:

- official name;
- canonical URL and identifier;
- address and coordinates;
- operator and brand;
- contact points;
- check-in and checkout policy at a general level;
- property-wide facilities;
- contained places such as restaurant and spa;
- images and general description.

Use the most specific applicable type - for example **Hotel**, **Resort**, **Hostel** or **BedAndBreakfast** - rather than treating **Organization** as the complete representation.

5.3 The accommodation unit

A room or suite should be its own entity when it has a meaningful category page or bookable identity. It can carry:

- occupancy;
- bed configuration;
- floor size;
- view;
- contained amenities;
- room-level accessibility;
- smoking policy;
- connection to the property;
- images.

Schema.org uses multi-typed entities so that a **HotelRoom** can also be a **Product** when it is the item offered commercially.[4]

5.4 The offer

Prices, validity and commercial conditions belong to the **Offer**, not to the room as an eternal fact. The offer can represent:

- rate plan;
- price and currency;
- date validity;
- occupancy basis;
- availability;
- breakfast or other inclusions;
- eligible audience;
- advance-purchase condition;
- booking URL;
- cancellation or payment references where the vocabulary and consuming platform support them.

Google's hotel pricing systems are specialized and separate property content from live price and inventory mechanisms.[12][13] Static website markup should not be used as a fragile substitute for a live rate feed or booking engine.

5.5 The reservation

A reservation is transaction state:

- requested;
- held;
- confirmed;
- paid;
- modified;
- cancelled;
- refunded.

This state belongs in authenticated systems. It should not be exposed as public structured data simply because a public vocabulary exists.

5.6 Property facts versus room facts versus offer facts

A common error is putting everything at hotel level.

Property level: outdoor pool, parking, restaurant, spa, address.

Room level: two king-size beds, 75 square meters, private sauna, communicating-room relationship.

Offer level: breakfast included, refundable until a date, minimum stay, member price.

Reservation level: room allocated, late arrival noted, payment received.

The level determines what can safely be inferred. "The hotel has connecting rooms" is not equivalent to "this booking guarantees two connecting rooms."

5.7 Guaranteed, available and request-only

Accommodation data needs a status vocabulary more precise than a Boolean.

- **Guaranteed by the selected unit or offer**
- **Available as a property capability**
- **Request-only and subject to confirmation**
- **Seasonal or conditional**
- **Unavailable**
- **Unknown**

This status may need to live in the content model even when Schema.org lacks a perfect property. The visible page should explain it, and an agent should not silently upgrade request-only to guaranteed.

5.8 An illustrative corrected graph

Appendix B provides an illustrative JSON-LD pattern for the case property. It deliberately omits live prices and availability because those values should be generated from the current commercial system. The example uses:

- one canonical hotel entity;
- one Standard room and one Duplex Suite entity;
- room-level occupancy and beds;
- a request-only communicating-room statement in visible description rather than a false guarantee;
- an independent restaurant entity;
- offers only where the values can be kept current.

The point is not to maximize field count. It is to make every public claim defensible.

CHAPTER 6

Put each fact in the system capable of keeping it true

6.1 Fact velocity

Travel facts do not all age at the same rate. Governance should begin by matching each fact to its volatility and system of record.

Put each fact in the system capable of keeping it true

Static JSON-LD is useful for public description; it is a poor substitute for live rates, restrictions or reservation state.

Stable identity	Contextual truth	Commercial truth	Transaction state
Changes rarely	Changes seasonally	Changes frequently	Changes per action
Name, address, operator, room size, permanent facilities	Restaurant opening, pool season, transport, temporary access	Rates, inventory, fees, minimum stay, restrictions	Held, booked, paid, modified, cancelled, refunded
SYSTEM OF RECORD	SYSTEM OF RECORD	SYSTEM OF RECORD	SYSTEM OF RECORD
CMS / entity record	Governed content / calendar	PMS / CRS / feed / API	Reservation / payment system

Figure 10. Put each fact in the system capable of keeping it true.
Source: Tamaga synthesis.

6.2 Stable identity facts

Examples:

- official name;
- address;
- coordinates;
- operator;
- room dimensions;
- permanent room categories.

Best home:

- governed CMS or entity record;
- visible page;
- structured data generated from the same fields.

Review when the property changes or at a scheduled low frequency.

6.3 Contextual facts

Examples:

- seasonal pool operation;
- restaurant opening days;
- local transport;
- construction or temporary access constraints;
- partner availability;
- event dates.

Best home:

- operational calendar;
- maintained content record;
- destination or partner feed where appropriate;
- explicit review date.

A generic amenity field is rarely enough.

6.4 Commercial facts

Examples:

- price;
- inventory;
- minimum stay;
- room availability;
- taxes and fees;
- rate restrictions;
- package inclusions.

Best home:

- PMS;
- CRS;
- channel manager;
- booking engine;
- hotel price feed or API.

These values should not depend on manual edits to a static page or JSON-LD block.

6.5 Transaction facts

Examples:

- room held;
- connecting allocation confirmed;
- payment accepted;
- late arrival acknowledged;
- cancellation processed.

Best home:

- reservation and payment systems;
- authenticated agent tools;
- audit log;
- human service workflow.

6.6 One source of truth is usually a myth

A hotel rarely has one database that can authoritatively maintain every fact. "Single source of truth" is therefore better interpreted as:

- one authoritative owner per fact type;
- known synchronization rules;
- explicit precedence when sources conflict;
- a record of where the fact is published;
- a correction path.

The property CMS may own room descriptions; the PMS may own inventory; the CRS may own rate plans; the front office may own late-arrival exceptions; the DMO may own local event records.

6.7 Propagation matters more than publication

When checkout changes from 12:00 to 11:00, the work is not complete when the official page changes. The organization should know whether the update reaches:

- structured data;
- booking engine;
- Google Business Profile;
- Google Hotels feed;
- OTA extranets;
- DMO record;
- printed confirmation;
- pre-arrival messaging;
- chatbot or agent knowledge.

A fact without a propagation map becomes tomorrow's contradiction.

6.8 Multilingual drift

Travel facts can change strength in translation. "Accessible rooms" may become "fully accessible hotel." "Connecting rooms on request" may become "connecting rooms available." "Dogs welcome" may become "pets allowed." Controlled concepts and named review are therefore operational requirements, not editorial niceties.

CHAPTER 7

Build the Omnichannel Decision Mesh

7.1 From page architecture to decision architecture

Laurent Bourrelly's Semantic Cocoon is a practitioner method built around audience analysis, search intent, offer-demand alignment, site architecture and deliberate internal linking.[15] Its useful contribution is not a secret LLM factor. It is the insistence that content architecture should begin with the user's intended outcome rather than with a pile of page formats.

His more recent omnichannel framing extends coherent authority across web pages, video, audio, newsletters and social channels.[16] Again, the relevant insight is architectural: a subject is represented across an ecosystem.

Christian Méline's Metawords work treats a subject as a semantic fingerprint rather than a keyword or synonym list, and emphasizes intention, lexical relationships and semantic continuity between pages.[17][18]

The **Omnichannel Decision Mesh** is a travel-specific extension inspired by these methods. It does not claim that either method is a proven LLM ranking mechanism.

7.2 The traveler decision at the center




Figure 11. The Omnichannel Decision Mesh.

Source: Tamaga synthesis, inspired by semantic-architecture methods from Laurent Bourrelly and Christian Méline.

The mesh begins with a decision:

Can I trust and book this stay for this traveler, date, route and constraint?

It connects six practical dimensions:

- **party fit** - occupancy, beds, connecting needs;
- **arrival** - transport, reception, late access;
- **access** - entrance, lift, room and bathroom suitability;
- **season** - facility operation, transport and local rhythm;
- **place** - neighborhood, evening food, noise and walking;
- **offer** - price, fees, cancellation, inclusions and support.

The human boundary is part of the mesh: which detail still requires confirmation?

7.3 Internal links as decision transitions

A useful internal link should move the traveler to the next unresolved question.

Examples:

- Standard room → communicating-room conditions;
- room → occupancy and beds;
- arrival page → station route and late-arrival procedure;
- October guide → seasonal facility status;
- restaurant page → service days and reservation;
- accessibility page → room-specific access;
- flexible rate → exact cancellation terms;
- uncertainty → named human contact.

"Learn more" is not a decision transition.

7.4 Beyond the domain

The mesh also includes external nodes:

- Google Business Profile and Maps;
- Google Hotels;
- OTA listings;
- destination office;
- transport provider;
- partner restaurant or guide;
- editorial coverage;
- review evidence.

The goal is not duplicate text. The goal is semantic and factual continuity. A destination office can provide local context; the hotel should provide room truth; the OTA can provide bookable policy; the transport operator owns the schedule.

7.5 Metawords as governed concepts

A keyword is a query expression. A governed concept is a maintained piece of meaning.

Consider **family-friendly**. For a hotel, the concept might resolve into:

- maximum occupancy;
- bed types;
- connecting-room status;
- cot and extra-bed policy;
- pool age rules;
- restaurant service times;
- lift and stairs;
- road crossing;
- transfer burden;

- child pricing;
- quiet-room options.

Consider **accessible**. It should resolve into:

- arrival route;
- entrance;
- reception;
- lift dimensions and coverage;
- bedroom circulation;
- bathroom configuration;
- restaurant and spa access;
- parking;
- assistance;
- limitations.

The controlled concept can guide:

- visible content;
- fields;
- filters;
- internal links;
- translation notes;
- structured-data choices;
- API attributes;
- AI prompts and review.

Not every nuance belongs in Schema.org. The model should not force a false Boolean where a human explanation is needed.

7.6 The mesh should include omission

Good travel knowledge states what is not guaranteed or not suitable. Examples:

- rooms connect only on request;
- the property is not ski-in/ski-out;
- pets are not allowed in the restaurant;
- a facility is seasonal;
- a room type is not adapted;
- a direct rate is non-refundable;
- late dinner is not available.

Omission reduces conversion theatre and increases decision quality.

7.7 How to audit the mesh

For one high-value traveler scenario:

1. List the decisions required.
2. Find the canonical page or source for each.
3. Record the internal and external transitions.
4. Identify missing conditions.
5. Compare terminology across languages and channels.
6. Test whether a human and an AI system reach the same qualified conclusion.
7. Test the route to booking or confirmation.

The output is not a keyword map. It is a decision continuity map.

CHAPTER 8

The wider destination web as evidence infrastructure

8.1 The official site cannot credibly prove everything

A hotel can authoritatively state its room configuration and policies. It is less independent when claiming that it is "the quietest," "best located," "authentic" or "ideal for families." External sources serve three distinct functions:

Identity - confirming that the entity is the same property.

Corroboration - independently supporting a claim.

Context - supplying destination knowledge outside the property's operational scope.

8.2 Destination organizations as information stewards

A DMO or tourism board can maintain:

- canonical place identity;
- events;
- transport context;
- seasonal openings;
- attractions;
- local businesses;
- accessibility information;
- media rights;
- correction and update channels.

Its value is not merely a backlink. It can reduce source poverty for independent properties that cannot create a large editorial footprint.

The Chamonix destination site adds independent detail about the hotel's spa and pool that is useful to both travelers and systems.[29] But it should not be expected to own room allocation or rate rules.

8.3 Partners should be visible as entities

A hotel experience may depend on:

- restaurant;
- host;
- transfer operator;
- guide;
- ski shop;
- museum;
- producer;
- event organizer.

When these actors are hidden in prose or private supplier records, the wider web cannot understand why the stay works. At the same time, public representation requires permission, accurate roles and review dates.

8.4 Reviews are evidence, not property facts

Reviews can support patterns about noise, service, cleanliness or family experience. They should not be converted into static facts without context. Review volume, recency, selection bias and platform moderation matter.

The useful question is:

Which decision does this evidence help a traveler make?

Not:

How many review keywords can be copied into the page?

8.5 Contradictions should be managed, not hidden

A hotel cannot control every external source. It can maintain a contradiction register:

- source;
- conflicting fact;
- current official value;
- commercial or safety impact;
- correction request date;
- response;
- fallback wording on owned channels.

This turns distributed authority from a vague PR objective into an operational workflow.

8.6 Visibility concentration

AI systems may overrepresent chains, central districts and digitally mature properties because those actors have more reviews, stronger feeds and richer source coverage. Destination organizations can improve the evidence environment for smaller businesses by maintaining precise, current records rather than publishing only campaign pages.

That is not a guarantee of recommendation. It is the minimum condition for fair consideration.

CHAPTER 9

Direct booking is not a button

9.1 A fair account of OTAs

OTAs provide genuine value:

- global reach;
- comparison;
- familiar payment;
- standardized policies;
- reviews;
- customer support;
- multilingual interfaces;
- fraud and chargeback infrastructure;
- mobile conversion;
- connected inventory.

The objective is not to erase them. It is to avoid a distribution system in which the intermediary becomes the only complete source and service interface for the property.

9.2 Why direct can be commercially valuable

SiteMinder's 2025 dataset reported an average booking value of US\$516 through hotel websites, compared with US\$445 through wholesalers, US\$392 through GDSs and US\$312 through OTAs.[9]




Figure 12. Average booking value by channel in SiteMinder data.

Source: SiteMinder Hotel Booking Trends 2026; 2025 booking data. Vendor panel, not a causal estimate.

Cloudbeds' independent-hotel panel reported OTA share at 63.4% and cancellation rates of 21.8% for OTA bookings versus 10.6% for direct bookings.[10]

Cancellation rates in Cloudbeds independent-hotel data

Observed across the vendor panel. Channel mix, rate conditions, market and property type may explain part of the difference.




Figure 13. Cancellation rates by channel in Cloudbeds independent-hotel data.

Source: Cloudbeds State of Independent Hotels 2026; 2025 booking data. Vendor panel.

These are vendor datasets, not universal causal estimates. Direct bookings may be higher value because different guests, room types, stay lengths and extras select the direct channel. A property should measure its own net contribution rather than copy the headline.

9.3 Direct booking does not mean free booking

Direct acquisition can involve:

- website and content cost;
- booking-engine fees;
- metasearch or paid media;
- payment processing;
- CRM;
- support;
- loyalty benefits;
- fraud and chargeback risk;
- technology integration.

A commission comparison alone is incomplete.

9.4 The channel choice is a trust decision

Experimental research by Lee and Sharma found that price parity can shift preference toward the direct hotel site, while cancellation flexibility and OTA credibility influence choice when offers differ.[11]

This matters for AI-mediated discovery. An AI answer may find the hotel, but the traveler will still compare:

- total price;
- room equivalence;
- refundability;

- payment terms;
- support;
- perceived reliability;
- ability to change the stay.

A "book direct" slogan cannot compensate for a worse or less intelligible offer.

9.5 The direct knowledge advantage

The direct channel can carry truth that an OTA listing flattens:

- which room combination actually fits;
- the difference between guaranteed and request-only;
- late-arrival procedure;
- exact accessibility conditions;
- neighborhood and local rhythm;
- host or concierge support;
- seasonal facility operation;
- what to omit;
- how the property handles disruption.

This is a knowledge advantage before it is a pricing advantage.

9.6 The direct service advantage

A direct relationship can support:

- named human confirmation;
- pre-arrival planning;
- special requests;
- modification;
- recovery;
- consented preferences;
- repeat communication;
- local partner coordination.

These advantages exist only if the property has the process and staff to deliver them.

9.7 AI can disintermediate or reintermediate

Two plausible paths coexist.

Direct path: the AI uses OTAs and editorial sources to research, recommends the hotel and links to the official site.

Intermediated path: the AI invokes a connected OTA, completes the transaction inside that ecosystem and leaves the intermediary as merchant, data holder or service owner.

The hotel should therefore track not only mentions, but:

- link destination;
- transaction channel;
- merchant of record;
- customer-data access;
- service owner;
- correction owner.

9.8 The direct path must be serviceable

A direct booking journey should make clear:

- what is being booked;
- total price and fees;
- applicable policy;
- room and feature guarantees;
- confirmation state;
- how to contact the property;
- how to modify or cancel;
- what happens after an error.

The commercial objective is not "remove the OTA." It is "carry more truth and responsibility than the alternative."

CHAPTER 10

The Source-to-Service Chain

10.1 From description to recovery

The Source-to-Service Chain

Visibility creates value only when the representation can survive comparison, live offer creation, confirmation and recovery.




Figure 14. The Source-to-Service Chain.

Source: Tamaga synthesis.

The **Source-to-Service Chain** has eight steps:

1. **Describe** - public facts and narrative.
2. **Resolve** - identify the correct property, room and offer.
3. **Corroborate** - compare independent and operational evidence.
4. **Compare** - evaluate fit against the traveler's constraints.
5. **Offer** - retrieve live price, availability and rules.
6. **Confirm** - surface consequential conditions and uncertainty.
7. **Act** - book, hold, pay or submit a request.
8. **Recover** - modify, cancel, correct or reach a human.

Visibility creates commercial value only when enough of this chain survives.

10.2 Description is not live truth

HTML and JSON-LD can describe that a Duplex Suite exists and sleeps four. They should not be assumed to prove that it is available on the traveler's dates.

A price feed or availability API can state that inventory is open. It may still not guarantee a connecting-room request or specific access requirement.

A booking API can create a reservation. It may still fail to provide an understandable recovery path.

10.3 The specialized interfaces

Accommodation distribution commonly separates:

- property content;

- room and rate plans;
- price and availability;
- restrictions;
- reservation actions;
- payment;
- service messaging.

Google's hotel stack likewise separates hotel-list content from rates, availability and inventory mechanisms.[12][13] OpenTravel standards provide messages and models for hospitality distribution. The web page is one interface among several.

10.4 Agent tools

MCP and other agent protocols can expose resources and executable tools. The MCP specification describes tools as model-controlled and recommends user visibility, confirmation and a human in the loop for consequential operations.[14]

An agent-ready hotel service therefore requires more than publishing a server. The tool must be:

- discoverable by the client;
- authorized;
- appropriate to the task;
- selected by the model;
- called with valid arguments;
- connected to live truth;
- confirmed by the user;
- logged;
- reversible or recoverable.

10.5 The promise boundary

For each action, the property should define what a machine may:

- explain;
- compare;
- quote;
- prepare;
- hold;
- book;
- modify;
- cancel;
- refund;
- promise.

For the recurring Chamonix scenario, a system may be able to show live room availability and rate conditions. A human may still need to confirm the exact connected-room allocation or nuanced access fit.

10.6 Read-only before write

A conservative progression is:

1. expose current facts;
2. retrieve live availability;
3. prepare a quote or booking draft;
4. require explicit human confirmation;
5. execute;
6. support modification and recovery.

The hardest problem is not creating the booking. It is preserving responsibility when the journey changes.

CHAPTER 11

Measurement without theatre

11.1 Why one prompt is not evidence

AI answers vary by:

- model and version;
- date;
- geography;
- language;
- logged-in state;
- conversation history;
- connected tools;
- search mode;
- source availability.

A screenshot is an anecdote. A useful visibility study needs repeated runs, defined prompts, recorded conditions and separate metrics.

11.2 The measurement vector

Entity resolution

Does the system identify the correct property rather than a similarly named hotel?

Candidate inclusion

Does the property enter the considered set for the scenario?

Recommendation frequency

How often is it actually recommended?

Fact accuracy

Are identity, room, policy, access, season and offer claims correct?

Qualification quality

Does the answer preserve conditions such as request-only, seasonal or subject to availability?

Source diversity

Which owned, intermediary, destination, review and editorial sources support the answer?

Citation and link destination

Is the hotel cited? Where does the link send the traveler?

Offer accuracy

Does the live price, availability and policy match the booking interface?

Action completion

Can the traveler complete the intended direct or intermediated action?

Recovery

Can the traveler correct or reverse it?

11.3 A practical prompt panel

Test scenario classes rather than only brand-name prompts:

- family room configuration;
- older traveler or mobility constraint;
- late arrival;
- pet policy;
- seasonal facility;
- quiet cultural stay;
- direct-versus-OTA comparison;
- cancellation and modification.

For each query, record:

- exact wording;
- date and time;
- model and mode;
- locale and language;
- session conditions;
- sources;
- hotels mentioned;
- recommendation order where meaningful;
- factual claims;
- links;
- tool calls;
- uncertainty language.

11.4 Control the denominator

"Share of voice" is meaningless unless the prompt universe is defined. A hotel's visibility for luxury couples in Paris cannot be compared with family ski queries in Chamonix without weighting assumptions.

Publish the prompt panel and weighting, or do not publish a single score.

11.5 Separate association from intervention

A hotel with rich schema may also have a stronger brand, better site, more reviews and broader distribution. Observational correlation cannot isolate the effect of markup.

A better test uses staged interventions:

- correct structured data;
- improve decision content;
- align distributed sources;
- improve direct handoff;
- allow propagation time;
- repeat the prompt panel.

Report what changed and what did not.

11.6 Business metrics

AI representation should be connected to:

- qualified direct referrals;
- direct conversion;
- net booking value;
- cancellation;
- support load;
- modification success;
- first-party consent;
- repeat booking;
- source-correction time.

The purpose is not to prove that AI traffic is valuable. It is to find where the source environment improves or damages the traveler's decision.

11.7 The proposed Property Truth Audit

A future full benchmark should audit a defined panel of independent properties across several destination types. For each, it should compare approximately 30 decision-critical facts across owned pages, markup, local profiles, OTAs, destination sources, booking interfaces and repeated AI answers.

This edition provides the method and one forensic case. It does not pretend that one hotel establishes sector prevalence.

CHAPTER 12

A practical implementation roadmap

A practical path from representation to accountable action

The sequence is intentionally conservative: correct public truth before automating consequential actions.




Figure 15. A practical path from representation to accountable action.

Source: Tamaga synthesis.

12.1 Stage 1 - Align (0-30 days)

Choose one canonical value and owner for:

- property name;
- address and contact;
- check-in and checkout;
- room names;
- occupancy and beds;
- parking and pet policies;
- accessibility wording;
- direct booking URL.

Compare the official site, markup, maps, OTAs, destination records and confirmation messages. Correct the highest-risk contradictions first.

Deliverable

A Property Truth Card and propagation map for the first 20-30 facts.

12.2 Stage 2 - Prove (30-90 days)

For consequential claims:

- attach the source;
- record conditions;
- distinguish guaranteed, available and request-only;

- add owner and review date;
- publish clear visible explanations;
- create a correction process;
- remove unsupported adjectives.

Deliverable

A maintained claim and source register, even if it begins as a spreadsheet.

12.3 Stage 3 - Distribute (60-120 days)

Build coherent public representation:

- correct property type and stable @id;
- model room categories;
- connect offers only where maintainable;
- improve room and policy pages;
- build decision-oriented internal links;
- align Business Profile, maps, OTAs and DMO records;
- govern multilingual concepts.

Deliverable

An Omnichannel Decision Mesh for the highest-value traveler scenarios.

12.4 Stage 4 - Connect (3-9 months)

Connect changing facts to their systems of record:

- room and rate plan data;
- availability;
- taxes and fees;
- cancellation;
- seasonal services;
- partner status;
- booking analytics.

Generate public structured data from governed CMS fields. Use feeds and APIs for live commercial truth rather than manual JSON-LD.

Deliverable

A documented source-to-publication architecture with monitoring.

12.5 Stage 5 - Act safely (6-12 months)

Begin with read-only capabilities:

- property facts;
- live availability;
- rate conditions;
- policy retrieval.

Then add:

- quote preparation;
- booking draft;
- explicit confirmation;
- audit logging;
- human escalation;

- modification and cancellation;
- recovery testing.

Deliverable

A bounded Source-to-Service Chain, not an open-ended "AI agent."

12.6 Priorities by organization size

Small independent property

Start with fact ownership, visible decision pages, correct basic markup, channel alignment and a clear human handoff. Do not build a custom knowledge graph or MCP server unless the operational need justifies it.

Small hotel group

Standardize identifiers, room and offer models, page templates, multilingual concepts, validation and propagation while preserving property-specific truth.

DMO or tourism organization

Maintain canonical local entities, event and transport data, small-business visibility, update rights, media rights and correction channels. Do not overwrite the property's responsibility for live room and rate facts.

Technology provider

Do not label a product "AI-ready" merely because it outputs JSON-LD or adds a chatbot. Demonstrate source ownership, synchronization, action boundaries, auditability and recovery.

12.7 The first test after implementation

Run the original traveler scenario again.

A good answer should state:

- which property and room configuration fit;
- which facts are established;
- which are request-only;
- which facilities and services are current for the dates;
- whether the direct and OTA offers are equivalent;
- which option has the relevant cancellation condition;
- what still requires human confirmation;
- how to proceed and recover.

That is more meaningful than an abstract visibility score.

CONCLUSION

The agreement is the asset

A modern accommodation property is represented by a network of systems and sources. Its official website is only one version. The booking engine owns a different kind of truth. Structured data makes selected facts explicit. Maps and OTAs aggregate. Destination and editorial sources corroborate. AI systems compress the fragments into an answer.

The versions do not need identical language. They need agreement on consequential facts and clear boundaries around uncertainty.

The evidence reviewed in this paper supports several firm conclusions:

- structured data is underused and frequently modeled poorly in hospitality;
- generative interfaces are changing how sources are selected and how clicks occur;
- external sources and intermediaries remain central to hotel discovery;
- direct booking can carry higher commercial and relationship value, but only when the offer and service justify trust;
- live action requires operational systems, permissions and recovery beyond public markup.

It does **not** support the claim that a new GEO discipline replaces serious SEO, or that a JSON-LD block automatically creates citation or recommendation.

The more useful ambition is harder:

Make the property's public story, structured facts, distributed representation, live offer and human service agree.

When they agree, machines make fewer guesses. Travelers can see what is confirmed. Direct channels can compete on more than commission. Staff can honour what the digital system promises.

The hotel is not a page.

It is an agreement.

Become the clearest source before trying to become the chosen answer.

APPENDIX A

Case methodology and detailed observations

A.1 Scope

Property: Hôtel Mont-Blanc Chamonix. Audit type: public-source forensic review. Primary review period: July 2026. Historical AI snapshot: Gemini report generated December 20, 2025 and reproduced in the supplied white paper. Structured-data version: third-party proposed JSON-LD from the supplied white paper; not verified as deployed markup. Direct engine: application interface present; live contents not fully inspectable in the text-only audit.

A.2 Status definitions

Precise / confirmed - explicit and supported by the reviewed source.

Partial / generic - present but insufficiently qualified for the decision.

Contradicted / erroneous - conflicts with a stronger current source or is modeled incorrectly.

Not stated / not observable - absent from the reviewed surface or inaccessible to the method.

A.3 High-value observations

Identity

The official property name is consistently represented by the owned site, Google and OTAs. The proposed JSON-LD uses a variant name that should not become the canonical identity without evidence.

Contact

The official site publishes info@hmbchamonix.com. The proposed JSON-LD uses a different address.

Check-in and checkout

Official site, Google and OTAs align on 15:00 check-in and 11:00 checkout. The proposed JSON-LD contains malformed time formatting and a 12:00 checkout value.

Duplex Suite

The official room page states four people, 75 square meters and two king-size beds. The proposal states a different bed configuration.

Connecting rooms

The owned site states that Standard and Superior rooms can connect and labels the feature on request. The AI report preserves the availability but does not foreground the request-only condition.

Accessibility

The owned information identifies two adapted Prestige rooms. Google and OTAs use broader labels. None of the reviewed public sources alone resolves the full step-free journey for a specific guest.

Late arrival

Some room pages state 24/7 reception. Expedia frames check-in to midnight and late check-in subject to availability. A traveler should not infer an unconditional guarantee without checking the selected rate and arrival procedure.

Sunday meal

The reviewed official restaurant information establishes Sunday lunch. It does not by itself establish a late Sunday dinner for the recurring scenario.

Pool and spa

The official site and destination office corroborate the heated outdoor pool and spa. The destination office adds daily opening and minimum age for the spa. Date-specific access and maintenance remain operational questions.

Direct booking

The official site exposes a direct booking path. This study did not perform a live date-specific parity transaction, so it does not claim that the direct rate is cheaper, more flexible or available for the recurring scenario.

A.4 Supplementary data

The detailed 16-fact by seven-version coding is provided as a companion CSV. The codes are an audit aid, not a quality score.

APPENDIX B

Illustrative accommodation JSON-LD pattern

The following pattern demonstrates structure rather than a production-ready deployment. It uses public case facts only where supported and deliberately avoids a live price. In production, values should be generated from governed records and validated against visible content.

```
{
  "@context": "https://schema.org",
  "@graph": [
    {
      "@type": "Hotel",
      "@id": "https://www.hotelmontblancchamonix.com/#hotel",
      "name": "Hôtel Mont-Blanc Chamonix",
      "url": "https://www.hotelmontblancchamonix.com/",
      "telephone": "+33 4 50 53 05 64",
      "email": "info@hmbchamonix.com",
      "checkinTime": "15:00",
      "checkoutTime": "11:00",
      "address": {
        "@type": "PostalAddress",
        "streetAddress": "62 Allée du Majestic",
        "postalCode": "74400",
        "addressLocality": "Chamonix-Mont-Blanc",
        "addressCountry": "FR"
      },
      "containsPlace": [
        { "@id": "https://www.hotelmontblancchamonix.com/rooms-and-suites/standard-room/#room" },
        { "@id": "https://www.hotelmontblancchamonix.com/rooms-and-suites/duplex-suite/#room" },
        { "@id": "https://www.hotelmontblancchamonix.com/#restaurant" }
      ]
    },
    {
      "@type": ["HotelRoom", "Product"],
      "@id": "https://www.hotelmontblancchamonix.com/rooms-and-suites/standard-room/#room",
      "name": "Standard Room",
      "occupancy": {
        "@type": "QuantitativeValue",
        "value": 2,
        "unitCode": "C62"
      },
      "floorSize": {
        "@type": "QuantitativeValue",
        "value": 19,
        "unitCode": "MTK"
      },
      "description": "Standard and Superior rooms can be connected on request; confirmation is required.",
      "containedInPlace": { "@id": "https://www.hotelmontblancchamonix.com/#hotel" }
    },
    {
      "@type": ["Suite", "Product"],
      "@id": "https://www.hotelmontblancchamonix.com/rooms-and-suites/duplex-suite/#room",
      "name": "Duplex Suite",
      "occupancy": {
        "@type": "QuantitativeValue",
        "value": 4,
        "unitCode": "C62"
      },
      "bed": {
        "@type": "BedDetails",
        "numberOfBeds": 2,
        "typeOfBed": "King-size bed"
      },
      "floorSize": {
        "@type": "QuantitativeValue",
        "value": 75,
        "unitCode": "MTK"
      },
      "containedInPlace": { "@id": "https://www.hotelmontblancchamonix.com/#hotel" }
    },
    {
      "@type": "Restaurant",
      "@id": "https://www.hotelmontblancchamonix.com/#restaurant",
      "name": "Le Matafan",
    }
  ]
}
```

```
"containedInPlace": { "@id": "https://www.hotelmontblancchamonix.com/#hotel" }  
}  
]  
}
```

B.1 Production cautions

- Do not add prices unless the generation process can keep dates, currency, taxes and conditions current.
- Do not mark a request-only feature as guaranteed.
- Do not copy third-party ratings without provenance and policy compliance.
- Do not place room-level facts at hotel level.
- Do not use **additionalType** to turn an amenity into an unrelated entity class.
- Validate syntax, modeling and factual parity separately.

APPENDIX C

Property Truth inventory

A practical audit should cover at least the following fields.

Identity

Official name; canonical URL; property type; operator; brand; address; coordinates; telephone; email; languages.

Accommodation

Room category; occupancy; bed configuration; floor size; connecting relationship; guarantee status; extra-bed policy; cot policy; smoking status; room-level accessibility.

Property access and arrival

Entrance; lift; adapted room; accessible bathroom; parking; EV charging; station route; shuttle; reception hours; late-arrival procedure.

Facilities and services

Restaurant; service days; breakfast; pool; spa; age rules; seasonality; pet policy; room service; concierge; ski storage.

Commercial offer

Direct booking URL; room availability; rate plan; total price; taxes; mandatory fees; cancellation; payment; minimum stay; inclusions; member conditions.

Service and recovery

Human contact; modification path; cancellation path; refund process; special-request confirmation; after-hours support; correction owner.

APPENDIX D

Data placement matrix

Fact type	Visible page	Structured data	CMS / content record	PMS / CRS / feed	API / agent tool	Human confirmation
Property identity	Yes	Yes	System of record	Reference	Read	Rare
Room size and beds	Yes	Yes	System of record	Map to inventory	Read	If unusual allocation
Connecting rooms	Explain request status	Limited	Capability record	Allocation state	Read / request	Usually yes
Check-in / checkout	Yes	Yes	Policy record	Booking rules	Read	Exceptions
Restaurant schedule	Yes	Possible	Operational calendar	Optional	Read	Special service
Pool operation	Yes	Amenity + conditions where supportable	Operational calendar	Optional	Read	Maintenance / exception
Price and availability	Summary only	Only when reliably generated	Reference	System of record	Live read	No, unless exception
Cancellation	Selected offer	Offer reference where supported	Policy library	Rate plan	Live read	Edge cases
Reservation state	Confirmation only	No public state	Reference	System of record	Authenticated read/write	Consequential changes
Accessibility fit	Detailed visible explanation	Selected stable facts	Audited access record	Room mapping	Read	Individual fit often yes

References

1. Google Search Central. "AI Features and Your Website." 2026. <https://developers.google.com/search/docs/appearance/ai-features>
2. Google Search Central. "Introducing Search Generative AI performance reports in Search Console." June 3, 2026. <https://developers.google.com/search/blog/2026/06/gen-ai-performance-reports>
3. OpenAI. "Publishers and Developers FAQ." 2026. <https://help.openai.com/en/articles/12627856>
4. Schema.org. "Markup for Hotels." 2026. <https://schema.org/docs/hotels.html>
5. Nicolas Sitter. "Hotel Schema.org Adoption Study 2026." 2026. <https://www.nicolassitter.com/research/hotel-schema-adoption-study-2026>
6. Nicolas Sitter. "AI Hotel Landscape 2026." 2026. <https://www.nicolassitter.com/research/ai-hotel-landscape-2026>
7. Pew Research Center. "Google users are less likely to click on links when an AI summary appears in the results." July 22, 2025. <https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/>
8. Adobe. "AI travel traffic surges as engagement hits new highs." June 2026. <https://business.adobe.com/uk/blog/adobe-report-ai-traffic-travel-sites-surges-200-percent>
9. SiteMinder. "Hotel Booking Trends 2026." 2026. <https://www.siteminder.com/hotel-booking-trends/>
10. Cloudbeds. "The 2026 State of Independent Hotels." 2026. <https://www.cloudbeds.com/hospitality-industry-report/>
11. Lee, S. W., and Sharma, A. "Beyond rate parity: Examining offer uniqueness and channel credibility in hotel pricing." *Tourism Economics* 31(2), 2025. <https://doi.org/10.1177/13548166241273881>
12. Google for Developers. "Hotel List Feed." <https://developers.google.com/hotels/hotel-prices/xml-reference/hotel-list-feed>
13. Google for Developers. "Availability, Rates, and Inventory (ARI) Overview." <https://developers.google.com/hotels/hotel-prices/xml-reference/ari-overview>
14. Model Context Protocol. "Tools." <https://modelcontextprotocol.io/specification/2024-11-05/server/tools>
15. Laurent Bourrelly. "Formation Cocon Sémantique." <https://www.laurentbourrelly.com/blog/boutique/cocon-semantique>
16. Laurent Bourrelly. "Cocon Sémantique OmniCanal." <https://www.laurentbourrelly.com/blog/boutique/cocon-semantique-omnicanal>
17. Christian Méline. "Définitions - Métamots, lexies et signature sémantique." <https://seo.metamots.xyz/cours-1-semantique-seo-definitions/>
18. Christian Méline. "Cocon sémantique et glissement sémantique." <https://seo.metamots.xyz/>
19. Hôtel Mont-Blanc Chamonix. "Winter Practical Information." Accessed July 2026. <https://en.hotelmontblancchamonix.com/winter/information>
20. Hôtel Mont-Blanc Chamonix. "Standard Rooms." Accessed July 2026. <https://en.hotelmontblancchamonix.com/rooms-and-suites/standard-room>
21. Hôtel Mont-Blanc Chamonix. "Duplex Suites." Accessed July 2026. <https://en.hotelmontblancchamonix.com/rooms-and-suites/duplex-suite>
22. Hôtel Mont-Blanc Chamonix. "Spa by Clarins." Accessed July 2026. <https://en.hotelmontblancchamonix.com/summer/spa>
23. Hôtel Mont-Blanc Chamonix. "Le Matafan / Restaurant information." Accessed July 2026. <https://en.hotelmontblancchamonix.com/>
24. Hôtel Mont-Blanc Chamonix. Official website. <https://www.hotelmontblancchamonix.com/>
25. Hôtel Mont-Blanc Chamonix. Direct booking application. Accessed July 2026 via the official website.
26. Google Hotels. "Hôtel Mont-Blanc." Accessed July 2026. https://www.google.com/travel/hotels/entity/CqsltL-rwM_ZqPiOARAB
27. Booking.com. "Hôtel Mont-Blanc Chamonix." Accessed July 2026. <https://www.booking.com/>
28. Expedia. "Hôtel Mont Blanc Chamonix." Accessed July 2026. <https://www.expedia.com/Chamonix-Mont-Blanc-Hotels-Hotel-Mont-Blanc-Chamonix.h425247.Hotel-Information>
29. Chamonix-Mont-Blanc Tourist Office. "Spa by Clarins Hôtel Mont-Blanc." Accessed July 2026. <https://en.chamonix.com/en-cas-de-mauvais-temps/spa-by-clarins-hotel-mont-blanc>
30. Chamonix-Mont-Blanc Tourist Office. Accommodation and pool listings. Accessed July 2026. <https://en.chamonix.com/>
31. Morand, Jean-Claude, and Schegg, Roland. *Devenez visible pour les IA: Les clés de l'optimisation sémantique et du GEO pour les hôtels*. Version 0.4.6, February 19, 2026. Supplied research file.
32. Ibid., proposed Hôtel Mont-Blanc JSON-LD case, pp. 43-46.
33. Ibid., Gemini Pro family-hotel report appendix, pp. 61-74.
34. Google Search Central. "A new resource for optimizing for generative AI in Google Search." May 15, 2026. <https://developers.google.com/search/blog/2026/05/a-new-resource-for-optimizing>
35. Schema.org. "HotelRoom." 2026. <https://schema.org/HotelRoom>
36. Nicolas Sitter. "The ChatGPT Direct-Traffic Explosion for Hotels." May 2026. <https://www.nicolassitter.com/research/chatgpt-hotel-direct-traffic-explosion-2026>

About Tamaga

Tamaga is a travel knowledge infrastructure company. It helps travel organizations structure what they know, prove what they claim, and build platforms that can be discovered, trusted, cited, updated, translated and operated.

Clearer sources for better travel.

