WORKING PAPER

Carrying creative semantics across the travel consideration window

Chris LaColla

24 July 2026

ABSTRACT

Discovery for considered travel purchases has migrated to social feeds, while purchase resolves days to months later at a booking surface. Across hotels, vacation rentals, tours, and cruise, the consideration window exceeds the lifetime of browser-resident identity, and the semantics of the initiating creative are discarded at the click. This paper assembles the booking-window and distribution evidence across verticals, corrects several platform-behavior claims that circulate in stale form, locates the proposal against existing feed-driven ad products and CDP architectures, and specifies a trip-scoped, append-only intent record, a thread, written by every surface and read at render time, with the email address as the identity carrier.

Consideration windows across travel verticals

The window between first interest and purchase varies by an order of magnitude across travel verticals, and the variation determines where the architecture applies.

Hotels. Global average booking lead time on SiteMinder's platform is 36.4 days, up from 35.7, within a typical industry range of 20–60 days; market-level averages run from 23.8 days (Malaysia) to 53.7 days (Ireland), and lead time compressed to 22.7 days in the year following the pandemic shutdown. The research window is longer than the booking lead time: Anderson and Han's comScore panel analysis found consumers make roughly 25 visits to travel-related sites in the 60 days preceding a hotel reservation.

Vacation rentals. AirDNA's median U.S. booking lead time was just under 40 days in 2019 and has remained under 30 days since the pandemic, with strong seasonality: over half of July–August bookings are placed more than 30 days out, while half of January–February bookings fall within two weeks of the stay.

Cruise. Only 11% of cruise travelers book a year or more before sailing, but more than 40% of North American travel advisors report client booking windows have lengthened to 9–18 months. Carnival Corporation characterized its 2025 booking curve as the furthest out on record against $8.5 billion in customer deposits. Cohort-clean fare tracking of Royal Caribbean sailings departing 15 April–4 June 2026 shows interior fares rising from a mean of $142 per person per night at 91–180 days out to $182 at 0–7 days, and balconies from $191 to $294, making early commitment economically rational.

The architecture below is parameterized by this window. At 30–60 days, browser-resident identity survives only with repeat visits; at six months, it does not survive at all. New-to-category share compounds the problem in cruise: 31% of cruisers over the past two years are first-timers, up from 24% in 2019, a discovery audience with no existing loyalty or profile relationship.

Identity mechanics

Safari. WebKit's current tracking prevention policy blocks all third-party cookies without exception, deletes all script-writable storage after seven days of browser use without user interaction with the site, and caps JavaScript-set cookies on a landing page at 24 hours when the arrival URL carries click identifiers, which is the shape of every tracked ad click. Cookies set through CNAME-cloaked third-party responses are capped at seven days, closing the server-side proxy workaround.

> The frequently repeated claim of a flat seven-day expiry on JavaScript cookies describes the 2019 ITP 2.1 rule. The current rule keys off interaction: storage survives while the visitor keeps coming back. That changes nothing at multi-month windows, since a shopper researching in February and depositing in August will go seven days without a visit.

Chrome. Google announced in April 2025 that it will not deprecate third-party cookies and will not ship the planned user-choice prompt. This removes the industry's forcing function but does not alter the analysis: third-party cookies never solved cross-device identity, and the display retargeting they enable decays over exactly the interval that matters. A large-scale randomized experiment found retargeting lifts four-week return rates by 14.6%, but a third of the first week's effect materializes on day one and effectiveness declines with time since the site visit.

Email. The operator-held address is the only identifier that survives the window and a device change. Cookie-based identity does not survive channel handoffs, which is why the token in the link must resolve server-side to the record, with no dependence on client state.

Why the discontinuity persists

Three factors, of which only the third is engineering. Paid social, lifecycle messaging, and the web property are separately owned, budgeted, and often separately outsourced; a mechanism spanning all three has no natural owner. The semantics of creative were never stored. Platforms return a creative identifier, but what the unit depicted lived in a brief and a filename. And the available tooling assumes either a SKU with a short window or trait-shaped profiles, and neither carries inbound-creative semantics to an owned surface.

Distribution and the direct channel

Hotels. Direct channels accounted for 50.9% of European overnight stays in 2023, still the largest channel group, but down from 57.6% in 2013, while OTA share rose from 19.7% to 29.6%. Skift Research projects direct digital channels overtake OTAs as the largest hotel distribution channel by 2030. The channels interact as much as they compete: Anderson's field experiment found OTA listing raised a hotel's own-site reservations 9–26%, and the 2017 follow-up found 65% of direct bookers had visited an OTA in the prior 60 days.

Vacation rentals. U.S. channel mix: Airbnb ~43%, direct ~28%, Vrbo ~21%, Booking.com ~8%. Europe inverts it: Booking.com ~48%, Airbnb ~40%, direct ~11%.

Cruise. The widely cited figure that advisors account for 72% of cruise bookings originates in CLIA leadership statements reported in trade press; CLIA's published consumer research instead shows 43–50% of cruisers self-reporting they booked through an advisor. The advisor share is therefore best stated as a 43–72% range.

Across all three verticals the pattern is the same: the direct channel is large, under continuous pressure from consolidated intermediaries, and the object of sustained operator investment, while remaining the least instrumented surface in the stack, because discovery moved to the feed and the direct funnel did not follow.

The inventory feed as catalog, and the prior art

A travel promotion is a structured set: for cruise, on the order of 10^2 sailings per promotional cycle, each row carrying ship, itinerary, departure, lead fare, availability, and expiry. Feed-driven creative over travel inventory is not novel as an ad product.

- **Meta Travel Ads** is a current product supporting hotel, flight, and destination catalogs with price and date fields, serving inventory-specific creative off cross-device intent signals. Cruise catalogs are absent from the product.
- **Feed-driven DCO platforms** generate creative variations from a product feed across social, display, and CTV, with travel as a supported vertical.

What no existing product does is persist the semantics of the engaged unit beyond the click. Meta's travel ads optimize delivery and render the catalog row into the creative; the click still lands on a surface that knows only that a click occurred. DCO systems personalize the ad, not the destination. The contribution proposed here is the continuity layer.

Architecture

The thread

A single durable record scoped to a trip intent, written by every surface and read at render time. It accumulates three signal classes: the engaged creative's semantics stored structurally, as fields and not as a bare identifier, volunteered zero-party responses, and behavioral events. Scoping to the trip rather than the person keeps concurrent intents from contaminating one another; a person record holds only durable attributes, and a new intent inherits selectively: embarkation preference and physical constraints carry, prior destinations and stale budgets do not.

The append-only design follows event-sourcing practice, with its documented costs adopted deliberately. Deriving state by replaying an event stream at request time is expensive; production systems materialize read-optimized projections instead, and those projections are eventually consistent. For this workload the lag is acceptable: the thread's reads are send-time and render-time, not transactional. Append-only storage conflicts with deletion rights, so personal data held outside the event store by reference, or crypto-shredding via per-subject keys, is a requirement. Event history is retained in full, because recurrence is signal: three views of one sailing across two weeks differs materially from one.

Position against CDP architectures

The customer data platform is trait- and audience-shaped: a repository of attributes and segment memberships. It has no native representation of which specific creative initiated this intent and what that creative depicted, which is the datum the entire mechanism turns on. Segment's own engineering documentation states its Profiles API is subject to response times unsuitable for the request path. The thread is that representation, plus the token discipline that re-establishes identity from an email click without depending on device state.

Presentation constraints from the personalization evidence

A meta-analysis of 53 experiments finds personalized advertising more effective than generic on persuasion, attitudes, and intent, with perceived relevance as the operative mechanism, and personalization built on real user data considerably outperforming inferred personalization. The same body of evidence carries two warnings. Covert personalization tends to outperform overt. And field experiments in email find that high personalization levels can reduce purchase probability, and that personalized placeholders in cart-abandonment messages depress opens precisely at the highest-intent moment.

The design consequence is a rule. Stated data determines what the surface *is*: which sailing leads, which imagery assembles, which alternates rank. It is not announced.

Measurement and falsification

Long-window attribution is where architectures of this class most commonly overclaim. eBay's randomized shutdown of brand-keyword search found 99.5% of paid clicks substituted by organic traffic, and conventional attribution implied an ROI above 4,100% where the experiment measured −63%. The pilot design accordingly treats attributed deposits as unmeasurable within a first cycle and specifies endpoints by latency: cost per zero-party capture in days, message-to-surface engagement and token return-visit rate in weeks, deposit attribution over one full window minimum and credible only against a holdout.

Falsification criteria. Capture cost no better than the broad-creative control means specificity does not drive response. Token return rate indistinguishable from baseline direct traffic means the email link is not functioning as the re-identification mechanism. Validation-gate rejection high enough that generated coverage falls materially below the bucket means the library's tagging cannot support the assembly step.

Limitations

Read access to live inventory and pricing is a contractual dependency without which the architecture has no input. The cruise fare curve is one operator over a seven-week departure range; the hotel lead-time figures are platform-specific samples; the advisor-share range spans trade-cited and self-reported figures that have not been reconciled. Platform behavior changes without notice. The cookie-cap correction and the April 2025 Chrome reversal both postdate widely circulated guidance. The retargeting and personalization experiments cited are retail-dominant; their transfer to travel windows is an assumption of this paper, not a measured result.

SOURCES

1. SiteMinder, Booking Lead Time: Hotel Trends https://www.siteminder.com/r/booking-lead-time/
2. Anderson & Han, The Billboard Effect: Still Alive and Well, Cornell SHA (2017)
3. AirDNA, Vacation Rental Metrics: Booking Lead Time https://www.airdna.co/blog/vacation-rental-metrics-booking-lead-time
4. CLIA, State of the Cruise Industry Report 2025
5. Carnival Corporation & plc, 2025 Q1 earnings release
6. All Aboard Analytics, The Booking Curve
7. WebKit, Tracking Prevention Policy; CNAME Cloaking and Bounce Tracking Defense (2020) https://webkit.org/tracking-prevention/
8. Google Privacy Sandbox, Next steps for Privacy Sandbox and tracking protections in Chrome (22 April 2025) https://privacysandbox.com/news/privacy-sandbox-next-steps/
9. Journal of Marketing Research 56(3), randomized field experiment on retargeting (2019)
10. HOTREC, European Hotel Distribution Study 2024
11. Skift Research, Hotel Distribution Outlook 2024
12. Meta for Developers, Travel Ads, Marketing API v25.0
13. Microsoft Azure Architecture Center, Event Sourcing pattern
14. Twilio Segment Engineering, Segment Profiles to the Edge for Real-time Personalization
15. Yeo, Chu & Li, How Persuasive Is Personalized Advertising? A Meta-Analytic Review, JAR 65(4) (2025)
16. Blake, Nosko & Tadelis, Consumer Heterogeneity and Paid Search Effectiveness, NBER WP 20171 (2014) https://www.nber.org/system/files/working_papers/w20171/w20171.pdf

HOW THIS WAS MADE

AI is part of how the work here gets made: finding sources, pulling figures out of long documents, drafting and restructuring. The arguments and the judgement are mine. Figures are checked against primary sources where a primary source exists, and the source is named next to the number so you can check it too.

That reduces errors and does not eliminate them. This is an exploratory surface and some of what is here will turn out to be wrong. If you find something, email **chris@lacoven.com** and I will correct it.

LaCo Ventures LLC · 407 Lincoln Rd, Suite 6H PMB 1277, Miami Beach, FL 33139