

A PRELUDE & TWELVE ESSAYS

On the Experiences Industry

*A point of view on the structural
forces shaping the tours, activities
and experiences market.*

MARC WIELAND

On the Experiences Industry

On the Experiences Industry

*A point of view on the structural forces
shaping the tours, activities and
experiences market.*

BY MARC WIELAND

On the Experiences Industry collects, in a single volume, a prelude and twelve essays on the structural forces shaping the tours, activities and experiences market.

All claims, numbers and historical references are traceable to cited sources or to the author's direct professional knowledge. Where the author's organisation has contributed to a source in an advisory capacity, the contribution is disclosed in the Author and Disclosures section at the end of the book.

CONTENTS

00	The Experiences Industry, From Where I Sit	6
01	Why \$270 Billion Still Runs on Spreadsheets <i>The digitisation lag in tours and activities is not a cultural lag. It is a unit-economic one.</i>	11
02	Fragmentation Is Not a Transitional State <i>Why fragmentation has held its shape through every wave that was supposed to consolidate it.</i>	21
03	The Shadow Infrastructure of the Experiences Industry <i>The infrastructure problem the industry has been solving in the dark, one company at a time.</i>	31
04	Anatomy of a Booking <i>A booking is a polite fiction. Trace what it actually has to do and the industry's shape becomes legible.</i>	46
05	What Amadeus and Sabre Don't Teach Us About Experiences <i>The GDS analogy is the one the industry reaches for most often. It does not transfer.</i>	57
06	Channel Managers Are a Partial Answer <i>The channel-manager category solves for shared integration cost. It does not solve for coherence.</i>	68
07	AI Doesn't Remove Infrastructure, It Reveals It <i>What looked like connectivity functioning at scale was humans absorbing the gap. That subsidy is being withdrawn.</i>	80
08	The Agent-Era Supply Problem <i>An agent does not compensate for incoherence. The supply problem the industry already had, revealed.</i>	89
09	Where Value Migrates When Interfaces Die <i>Search, social, cloud, retail. The same structural move each time. Experiences is next.</i>	101
10	The Experiences Industry Is Hitting Its Inflection Point <i>Five structural conditions that held the industry's shape are weakening simultaneously.</i>	111
11	What the Experiences Industry Looks Like in 2030 <i>Two paths, read side by side. Structural reasoning forward, not forecasting.</i>	123
12	The Question the Industry Is Not Asking Itself <i>The question is being answered either way. In the industry's voice, or in its silence.</i>	134

The Experiences Industry, From Where I Sit

I have spent many years inside the experiences industry, building products, running companies, watching it change from up close. In recent months I have moved into a role with more room to think than to operate, and it has given me time to do something I have been meaning to do for a while: set down, in public, how I have come to see this industry and where I think it is going.

The timing is not incidental. The industry is entering a structural transition that will shape how discovery, distribution and booking work over the next decade.

This is the first of what I expect to be a longer body of writing. I am not announcing a programme, and I am not leading up to anything. I am writing because I have a point of view, and because after enough time inside a market, one develops an obligation to share what one sees, especially when the market is changing as quickly, and as unevenly, as

this one is. These essays are pointed by design, but they are not casual opinion. Each essay is grounded in industry data, public sources, academic work where useful, and my own professional experience.

So let me put down three things I have come to believe. Each of them is a long argument in compressed form; I expect to spend some of my future writing taking them apart properly.

One. This is the most fragmented industry in travel, on both sides.

Almost everyone who works in travel experiences can name the supply side fragmentation. Some six hundred thousand operators worldwide, the majority of them small, roughly a third meaningfully digitised, the rest running on email, spreadsheets, paper diaries, and phone calls. Phocuswright described the market as “fantastically fragmented” in 2011. Fifteen years on, the underlying shape has barely changed.

What is less often named is that the sales side is also fragmented, in ways that distinguish experiences from every other travel vertical. Hotels have two online travel agencies that, between them, account for the majority of online distribution. Flights consolidate through three global distribution systems. Cars through a handful. Experiences have none of that. The visible top of the market is a long list: Viator, GetYourGuide, Klook, Expedia, Headout, Civitatis, Fever, KKday, Trip.com, TUI Musement, Groupon, the more recent entrants Booking.com and Airbnb, and countless regional players. Beneath them sits an enormous and largely

invisible layer of hotel concierges, tour desks, destination marketing organisations, reseller networks, cruise shore-ex teams, bus-tour brokers and local agents who still sell much of the product that moves through any given destination. No single online travel agent commands a clearly dominant share of the industry's bookings; even the top ten combined do not approach the concentration that hotels and flights show.

This matters more than it is usually made to matter. When both sides of a market are fragmented, no part of the system has the gravity to pull the rest into shape. Supply does not organise around a dominant distributor, because no distributor is dominant enough to force organisation. That is the condition this industry has been in for a long time, and it is worth understanding why it has persisted.

Two. The connectivity layer has been the most consistently underestimated layer of this industry.

One answer to why fragmentation has held its shape is that the layer that would make it navigable at industry-wide scale never matured. Connectivity does not end fragmentation. It makes a fragmented market machine-legible, distributable and bookable. That is a different kind of solution from consolidation, and in this industry it may be the more important one. In practical terms, this is the plumbing that moves availability, pricing, bookings and content between operators, resellers, distributors and channels. When I look at why this industry does not consolidate in the ways adjacent markets have, I keep ending up in the same place.

Not at the level of the operator, and not at the level of the marketplace, but at the layer between them, which, in this industry, does not exist in any mature form.

This is unglamorous ground. Nobody pitches connectivity as a headline. But the industries that have been transformed over the last half-century (airlines in the 1970s, hotels through the 1990s, retail in the 2000s, banking now) were transformed from the plumbing up, not from the shop window down. I think experiences will follow the same path, and I think the participants who are paying serious attention to the plumbing will be the ones best placed when it does.

Three. The rules of discovery, distribution, and booking are about to be rewritten faster than most participants realise.

Discovery is moving into generative interfaces. A new layer is forming between the traveller and the operator: not an app, not a website, but a conversational agent that answers travel questions and, increasingly, executes on them. Booking pathways are being rebuilt around those agents. The review-layer economics that have carried marketplaces for more than a decade are beginning to shift under their feet.

The argument does not depend on agents becoming the majority booking path immediately. It depends on them becoming commercially material enough that supply which is incoherent, incomplete or unreachable to machines becomes structurally disadvantaged. AI does not remove the infrastructure problem this industry has been carrying. It

exposes it. None of this is yet visible in the headline numbers, which is precisely why it is worth writing about now, while there is still time to think clearly about what it means rather than react to what it has already done.

I do not know exactly how this plays out. I do know that the industry is not yet looking at it with the seriousness it deserves.

Those three observations are where I plan to start. Each deserves its own essay, and will have one. I will write when I have something considered to say rather than on a fixed cadence. I will write often enough to be worth following if this is your industry too. I will not be promoting a company. I will not be selling a product.

This book does not argue that any existing company, including mine, already holds the answer. It argues that the question itself is becoming unavoidable, and that the industry should ask it directly rather than have it answered by default.

I will be trying to think clearly in public about a market I care about, and inviting the people who work in it to think alongside me.

A closing thought, to set the tone for what follows. Almost every participant in this industry can describe the problem. Almost none can describe the shape of the solution. That gap is what I want to write about.

ESSAY ONE

Why \$270 Billion Still Runs on Spreadsheets

The digitisation lag in tours and activities is not a cultural lag. It is a unit-economic one.

The global in-destination experiences market, covering tours, activities, attractions and events, is approximately \$270 billion in size as of 2024, and is forecast to reach approximately \$342 billion by 2029 at a compound annual growth rate of roughly 8%.¹ By travel-industry standards this is a large, well-documented and visibly growing category. By the standards of any comparably sized category in which digital systems have had twenty years to take hold, it is extraordinarily under-digitised.

The headline figures are well-rehearsed within the industry but worth stating plainly. Approximately 33% of the industry's bookings flow through online channels, against roughly 64% for broader travel.² Among operators themselves, around 39% still run without any vendor booking

system, with another small share on in-house tools; their availability calendars, reservation lists and post-booking communications live across a mix of email, spreadsheets, paper diaries and telephone calls.³ A separate Arival finding, attached to a 3,500-plus operator-response base, describes this as the industry's "dirty secret."⁴ McKinsey, surveying the same operator layer for a different metric, puts the share of operational data that remains unstructured at 40-60%, a related readout of the same digitisation gap.⁵

The reflexive explanation for this state of affairs is cultural. The experiences industry is said to be slow to adopt software because its operators are small, locally rooted and proudly analogue; because the people running these businesses are guides and curators rather than technologists; because the category has always lagged other travel verticals and will continue to lag them for the foreseeable future.

This framing is widespread, partially true, and ultimately the wrong level of analysis. The assertion that operators "just need to catch up" presupposes that there is a structurally similar rung they could step onto, that the software stack appropriate to a boutique hotel or a regional airline is, with minor adjustments, also appropriate to a ten-person kayak operator or a family-run vineyard tour. That presupposition does not survive a look at transaction volume: a mid-sized hotel will typically handle more bookings in a month than a typical tour operator handles in a year, and a regional airline will handle more in a day than a typical tour operator handles in a month. The same software stack does

not face the same per-booking economics across those scales. The reason is simple: most software costs arrive as fixed or semi-fixed costs, while the operator's ability to absorb them depends on booking volume.

The more useful reframe is to ask a quieter but harder question: what is the unit economics of enterprise software adoption for a tour operator, and how does it compare with the unit economics that drove digitisation in hotels and airlines? Hannes Werthner and Stefan Klein, in their 1999 study of information technology and tourism, and Dimitrios Buhalis, in *eTourism*, both treat tourism as an information-intensive sector: one in which the product itself is shaped by, delivered through, and marketed by information flows.⁶ On this description, tours and activities are at least as information-intensive as hotels or flights; arguably more so, because the product is more heterogeneous and less standardisable. If information intensity alone drove digitisation, the experiences category would be as digitised as the adjacent categories, or more so.

It is not. The reason, on closer inspection, is that the economic surface area over which software costs can be amortised is fundamentally different. Across the operator population, approximately 86% of businesses run below \$250,000 in annual revenue, and approximately 50% have been operating for less than five years.⁷ At those scales, the per-booking cost of a conventional booking-system, channel-manager and revenue-management stack is not a marginal operating expense; it is a proportionally significant fraction of total revenue. The software decision is therefore not a

question of willingness. It is a question of whether the software exists in a form whose unit economics work at the operator's scale.

A worked sketch makes this concrete. Consider an operator typical of the long tail: roughly two hundred thousand dollars in annual revenue, around four thousand bookings a year at an average ticket value of fifty dollars, two employees and a founder-operator. A booking-system and channel-connectivity stack that costs five to eight thousand dollars a year before implementation, configuration and ongoing administration is not a rounding error. It is two to four percent of revenue and a meaningful share of operating margin. The same annual software cost, deployed by a larger attraction or a hundred-room hotel handling tens of thousands of bookings a year, amortises across an order of magnitude more transactions. The software is not the variable. The denominator is.

When this frame is applied, the 33% digitisation figure stops being a puzzle of cultural lag and starts being a predictable equilibrium. Operators that can make the economics work, such as higher-volume attractions, larger multi-site tour companies, and the upper tail of the operator distribution, appear from the same operator data to digitise at rates closer to the rest of travel, even if the cohort cut is not always made cleanly in surveys. Operators below a threshold that the industry has not cleanly quantified, but which hovers somewhere around the long-tail line, adopt patchier, cheaper, often consumer-grade tools that they do not themselves describe as "booking systems" in operator

surveys. There is also a measurement artefact at work: a spreadsheet used systematically, a shared messaging channel used as a dispatch tool, or an off-the-shelf calendar tool wired into a payment provider by the operator themselves can all be functionally digital operations, even if they do not register as such in an Arival or Phocuswright response. The cultural-lag reading double-counts the gap.

The important consequence of this reframe is that the usual remedies do not follow from it. If digitisation is a cultural matter, the remedy is patience: eventually operators will modernise. If digitisation is a unit-economic matter, the remedy is structural. Either the software has to become an order of magnitude cheaper per booking (which is bounded by the underlying cost of building and maintaining it), or it has to be bundled with a distribution mechanism that carries the cost (which introduces its own distortions in distribution economics), or the economic frame itself has to change. The direction in which current indicators point is the third.

Myskiv and Nycz-Wojtan, in their 2022 treatment of reservation systems in small tourism enterprises, observe that in the SME operator segment, booking systems are often adopted primarily as marketing surfaces rather than as operations tools.⁸ The booking form on a website is the public-facing reason for the software; the internal use is patchier and weaker. This is not a failing of the operator. It is a rational response to a system whose per-booking cost is higher than its per-booking operational value. The marketing surface provides a return; the operational depth, at the operator's scale, often does not.

Nine years ago, Phocuswright's Global Travel Activities report made the same observation in a slightly different register. Most operators were either not digitised or were "minimally" digitised, and the structural reason was the mismatch between the cost of enterprise-grade tooling and the scale at which most operators run.⁹ The 2017 reading was, by the standards of the time, cautious. The 2026 reading is less cautious on the scale, since the market is half again as large and growing faster than broader travel, but arrives at structurally the same description of the operator layer.¹⁰ That nine-year continuity is the most important single fact about the digitisation question in this industry. It is not that the catch-up has been slower than expected. It is that the catch-up frame was the wrong frame to hold the question in to begin with.

What, then, is the right frame? A category whose unit economics do not support conventional enterprise software adoption will either remain under-digitised indefinitely, or will digitise through some mechanism whose cost is not borne one-operator-at-a-time. Two candidate mechanisms, each of which moves the cost off the operator's per-booking budget but by a different route, have been visible in the industry for close to a decade. The first is distribution-bundled software, in which a distribution partner absorbs the fixed cost of the stack in exchange for the downstream booking flow: the model that the industry's major online travel agents pursued through a cluster of reservation-technology acquisitions in 2018. The second is infrastructure-layer software, in which a neutral layer

provides the functions below the operator-level stack and removes them from the operator's budget line altogether.

Of the two, only the first has been materially tested at scale, and the test is instructive. The distribution-bundled model was initially built on a zero-operator-fee pricing structure, with economic return captured through consumer-side booking fees: a design that made commercial sense if adoption scaled fast enough for the distribution flow to cover the cost of maintaining the platform. That original design is beginning to shift. Connectivity fees have been introduced in selected regions by OTA-owned reservation-technology providers, and the public pricing profile of the category has started to resemble that of a conventional paid-for software product. That shift is, in itself, testament to one of two underlying realities: either adoption did not scale to the volume the original economics required, or the unit economics of distribution capture alone did not carry the platform at the scale the acquirers had modelled.

There is a second structural issue alongside the unit economics. A reservation system owned by a distribution partner is, at the margin, optimised to route bookings through that partner's channels, which is not the same thing as a reservation system whose commercial interests are aligned with the operator's. Operators who adopt OTA-owned reservation technology have, whether the choice is framed this way or not, made a distribution-channel decision at the same time. The distribution-bundled model therefore resolves the unit-economic problem by transferring rather

than removing it, and introduces an incentive structure that does not align neatly with the operator's own economics.

None of this is an indictment of the model. It remains live. But the proposition that the distribution-bundled approach has cleanly solved the unit-economic problem, and that the industry should expect to converge on it, is harder to defend in 2026 than it would have been in 2018.

The cultural-lag reading would have the industry wait. The unit-economic reading does not. If the gap is economic rather than attitudinal, it will close only when the economic frame changes: when the cost of the layer operators need is no longer carried one operator at a time. Distribution-bundled software changed that frame in one direction. Whether an infrastructure layer can change it in another, and in what form, is the question the next essays in this sequence take up.

ENDNOTES

1. Phocuswright × Arival, *The Outlook for Travel Experiences 2019–2029*, February 2026.
2. Phocuswright × Arival, *The Outlook for Travel Experiences 2019–2029*, February 2026, cross-referenced against the 2025 Arival Experiences Traveler Outlook 2026 and Phocuswright global travel data.
3. Arival, *Global Operator Landscape 4th Edition: The State of Experiences*, January 2026, which reports approximately 39% of operators with no booking system, with most of the remainder relying on a vendor platform and a smaller share on in-house tools. The “mix of email, spreadsheets, paper diaries and telephone calls” describes the operational layer beneath the booking surface for the no-system cohort.
4. Arival, *The Dirty Secret of the Tours and Activities Industry, 2025*, summarising the GOL 4th Edition operator-layer findings.
5. McKinsey & Company, *The Evolving Role of Experiences in Travel*, September 2024. The 40–60% range refers to operational data unstructuredness across the category and is presented here as a complementary rather than identical readout to the booking-system adoption gap.
6. Buhalis, D., *eTourism: Information Technology for Strategic Tourism Management*, Financial Times Prentice Hall, 2003, particularly chapters 1–3 on information intensity as a sector-defining feature. The earlier foundational treatment is Werthner, H. and Klein, S., *Information Technology and Tourism: A Challenging Relationship*, Springer-Verlag, Vienna, 1999, which framed tourism as among the most information-intensive sectors and is the antecedent to Buhalis’s later argument.
7. *Tourpreneur* coverage of Arival operator-size statistics, reconfirmed in the Arival *Global Operator Landscape 4th Edition*, 2026. For the broader long-tail framework drawn on here, see Anderson, C., *The Long Tail: Why the Future of Business is Selling Less of More*, Hyperion,

- 2006, and the earlier empirical foundation in Brynjolfsson, E., Hu, Y. J., and Smith, M. D., "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, 49(11), 2003, pp. 1580-1596.
8. Myskiv, G. & Nycz-Wojtan, A., *Reservation Systems as Tools of Marketing Communication in Small Tourism Enterprises*, 2022. DOI 10.56318/eem2022.02.007. Consistent with earlier work on the structural mismatch between SME scale and productivity returns from enterprise IT in tourism: Sigala, M., "The Information and Communication Technologies Productivity Impact on the UK Hotel Sector," *International Journal of Operations & Production Management*, 23(10), 2003, pp. 1224-1245.
9. Phocuswright, *Global Travel Activities*, 2017.
10. Arival, *Global Operator Landscape 4th Edition: The State of Experiences (2026)*; Phocuswright × Arival, *Outlook for Travel Experiences 2019-2029 (2026)*.

ESSAY TWO

Fragmentation Is Not a Transitional State

Why fragmentation has held its shape through every wave that was supposed to consolidate it.

In a 2011 research report, Cathy Schetzina and Douglas Quinby, writing for Phocuswright, opened their analysis of the in-destination tours and activities sector with a chapter titled “A Market Fantastically Fragmented.”¹ The chapter title was not rhetorical flourish. The evidence assembled in the report (a supply side composed overwhelmingly of independent operators, a distribution layer without a dominant platform, and no standard architecture for how bookings moved between the two) fit the description exactly.

Fifteen years later, the underlying shape of the industry is broadly the same. Current estimates put the global operator count at approximately 600,000, the large majority running businesses of under \$250,000 in annual revenue.² Roughly half of all operators have been trading for fewer than five

years.³ By current readouts of online distribution share, no single online travel agent commands a dominant share of the industry's bookings, and the top ten of them together do not. The market's fragmentation, on both the supply side and the sales side, has been a durable structural feature, not a passing stage of commercial development.

This essay takes up a question the industry routinely mishandles. Is the fragmentation observed in 2011 and still observed in 2026 transitional, a stage the market is slowly working through, or structural, a permanent feature of the market's supply shape? The answer matters. If fragmentation is transitional, the industry's strategic posture should be to wait it out and adapt to the consolidated market that will eventually arrive. If it is structural, the industry has been preparing for the wrong event for fifteen years, and should be building for a different future altogether.

The argument this essay makes, on the evidence, is that fragmentation in this category is structural. This is not a prediction. It is a description of a pattern visible over decades, grounded in the specific economic and experiential properties of the product, that the common industry framing consistently underweights.

The "consolidation is coming" framing

The framing assumed by most industry observers is that the experiences market will eventually consolidate in the pattern set by adjacent travel verticals. Hotels consolidated into chain brands through the 1980s and 1990s; airlines consolidated through the 2000s; online distribution in both

categories concentrated around a small number of online travel agents. By this account, experiences are at an earlier stage of the same trajectory, and the long tail will contract into a smaller number of larger operators as the category matures.

Concrete evidence is often cited for this view. A visible wave of top-down merger, acquisition and capital-markets activity has unfolded across the distribution layer in 2024–2025: American Express Global Business Travel’s acquisition of CWT, Expedia’s acquisition of Tiqets in late 2025, Klook’s filing for a New York Stock Exchange listing in November 2025, and a \$100 million Series F round for Klook earlier in the same year.⁴ The observations are correct; the inference from them is not. The consolidation currently visible is real, and consequential at the layer where it is occurring. But it is occurring at the distribution and aggregator layer, not within the operator base, and the long-tail supply structure is the question this essay is taking up. The long-tail supply side is not consolidating, and it is not showing the structural signals that precede consolidation in other categories. New operators continue to enter the market at a rate that has, if anything, accelerated since the pandemic.² The long tail is not contracting; it is regenerating.

The product cannot be standardised in the way services can

Two structural features of the experiences category explain why this is so, and the first is the nature of the product itself. Joseph Pine and James Gilmore, in the 1998 Harvard

Business Review article that named the experience economy, and at greater length in the 1999 book of the same name, argued that experiences sit above services on a commoditisation ladder that runs from commodities through goods to services and then to experiences.⁵ Each step on the ladder is more personalised, more memorable, less fungible than the one below. A hotel room is a service: differentiated at the margin, but fundamentally comparable on a small number of service attributes such as room quality, location, and price. An airline seat is similar. An experience, by contrast, is defined by its heterogeneity. A guided sunset kayak tour through a Costa Rican mangrove and a pastry-making class in a Parisian bakery share the word “experience” in a way that a double room in a business hotel and a double room in a resort hotel share the word “room,” but the comparability ends much earlier.

The consequence of this difference is that the consolidation pressure that produced chain hotels and consolidated airline groups does not act on experiences in the same way. Consolidation in hotels was enabled by the fact that a traveller could credibly substitute one hotel for another on a short list of comparable dimensions; the chain extended that substitutability across a network of locations, and the network effect did the rest. In experiences, substitutability is much weaker. The kayak tour is not the pastry class; a global kayak-tour brand does not offer the pastry traveller anything. The product heterogeneity limits the distance that consolidation can travel.

This is not a claim that no consolidation can ever occur in experiences. It is a claim that the unit of consolidation is different. Consolidation can occur at the aggregator layer, where platforms such as Viator and GetYourGuide consolidate access to experiences while leaving the supply side fragmented, and at the destination-specific operator layer, where a single operator in a single city may grow large. The pattern, in which a multi-sided platform grows by aggregating access to a fragmented supply base it does not own, is the canonical platform-economy form described by Parker, Van Alstyne and Choudary in *Platform Revolution*.⁶ What this pattern cannot easily do is produce a chain-equivalent structure that spans the category, because there is no category-wide comparable product to chain.

Demand is compound-segmented, not axis-segmented

The second structural feature is on the demand side. Travellers do not choose experiences along a single dimension. Destination, activity type, budget, timing, motivation, group composition and personal preference combine into specific demand pockets. Swarbrooke and Horner, in *Consumer Behaviour in Tourism*, and the more recent review by Cohen, Prayag and Moital, describe this formally as a compound set of motivators and determinants acting together,⁷ with demographic, psychographic, behavioural and benefit-seeking axes combining into segments that cannot be collapsed onto a single dimension. The well-worn industry segmentations (leisure versus business, domestic versus outbound, adventure versus

cultural) are axis-level simplifications of a segmentation that is in fact compound. The consequence for supply is that a market serving compound-segmented demand cannot collapse into a small number of mass-market products without leaving the long-tail demand unserved. In hotels, compound segmentation produced brand families. Marriott now operates several dozen brands, each aimed at a distinct compound segment, but the products are still drawn from a manageable vocabulary of service attributes. In experiences, the same compound-segmented demand produces something more granular: an effectively infinite catalogue of specific products, each serving specific compound segments, and each resistant to absorption into a larger vehicle.

Analogies in other categories

This pattern is not unique to experiences. Several adjacent categories show the same structural behaviour, and the analogies are instructive.

The restaurant industry has been described as “about to consolidate” for several decades. Chain restaurants exist, and some are large, but the independent-restaurant base has been remarkably stable over time. Most cuisine, most dining occasions and most local restaurant economies are still delivered by independents.⁸ Consolidation has occurred at the franchise layer and at the supplier layer, not within the operator base. The structural feature responsible is the same one Pine and Gilmore name: restaurants sell an experience, not a standardised service.

Independent retail survived the rise of Amazon not by consolidating against it but by specialising under it, with the long tail of specialist independent retailers stable or growing in many segments.⁹ In each case, a category with irreducible product heterogeneity has resisted the consolidation that analysts predicted and that seemed structurally comparable to the consolidation occurring in adjacent categories. Experiences sit cleanly within this pattern. It is not coincidental.

The right question

If the long tail is the permanent shape of the supply side in this industry, the question the industry has been asking for fifteen years is the wrong question. “When will fragmentation end?” has no useful answer, because fragmentation at the supply base is not going to end in the way the question implies. A more useful question is: how does an industry serve its market at scale through fragmentation, rather than around it? The two questions imply very different architectural decisions. The first points towards waiting for a dominant consolidation event that will not arrive. The second points towards infrastructure that makes a fragmented market navigable (connectivity, distribution, discovery, trust, settlement) without requiring the supply side to collapse into a smaller number of operators. That is a different kind of problem. It is also a harder one. It does not resolve itself by waiting. It requires deliberate design, and it requires the industry to accept a set of constraints it has so far preferred not to accept.

Fifteen years is long enough to notice a pattern. The 2011 Phocuswright description of a “fantastically fragmented” market is not a snapshot from an earlier stage of the industry. It is a description of a structural equilibrium the industry has been in for at least that long and shows every sign of remaining in. The interesting work is no longer to argue for or against consolidation; it is to describe what operating at scale through fragmentation actually requires. The future of this industry is not a consolidated supply base; it is the infrastructure that makes a fragmented one navigable. That is what the rest of this book is about.

ENDNOTES

1. Schetzina, C., and Quinby, D., *When They Get There (and Why They Go)*, Phocuswright, February 2011. Chapter 1, *A Market Fantastically Fragmented*.
2. Arival, *Global Operator Landscape 4th Edition: The State of Experiences*, January 2026. The 600,000 figure is an Arival industry estimate rather than a precise enumeration. The same report tracks the rate of new operator entry into the category, which has continued to rise post-2020 across most regions.
3. *Tourpreneur* coverage of Arival operator-size statistics, reconfirmed in the *Arival Global Operator Landscape 4th Edition*, 2026.
4. Phocuswire, *The Coming Wave of Consolidation in Tours and Activities*, 2024–2025 coverage; American Express Global Business Travel investor releases announcing (March 2024) and completing (September 2025) the CWT acquisition; Expedia Group, *Agreement to Acquire Tiqets*, 10 December 2025; Klook Technology Limited, Form F-1 registration statement filed with the U.S. Securities and Exchange Commission, 10 November 2025 (NYSE listing under the proposed ticker “KLK”); Klook, *Closes \$100 Million Investment to Fuel Next Decade of Travel Experiences Growth*, February 2025.
5. Pine, B. J., II, and Gilmore, J. H., “Welcome to the Experience Economy,” *Harvard Business Review*, July–August 1998. The expanded book-length treatment is Pine, B. J., II, and Gilmore, J. H., *The Experience Economy: Work Is Theatre & Every Business a Stage*, Harvard Business School Press, 1999.
6. Parker, G. G., Van Alstyne, M. W., and Choudary, S. P., *Platform Revolution: How Networked Markets Are Transforming the Economy and How to Make Them Work for You*, W. W. Norton, 2016, particularly the chapters on multi-sided platform aggregation across fragmented supply bases.

7. Swarbrooke, J., and Horner, S., *Consumer Behaviour in Tourism*, Butterworth-Heinemann, particularly the chapters on motivator/determinant segmentation and the compound nature of tourism decision-making. For a more recent academic review of the same demand-side framework, see Cohen, S. A., Prayag, G., and Moital, M., "Consumer Behaviour in Tourism: Concepts, Influences and Opportunities," *Current Issues in Tourism*, 17(10), 2014, pp. 872-909.
8. National Restaurant Association, *State of the Restaurant Industry*, 2024 edition, which tracks the share of independent operators across the U.S. restaurant industry and the long-running stability of that share against repeated forecasts of chain consolidation.
9. Hortaçsu, A., and Syverson, C., "The Ongoing Evolution of US Retail: A Format Tug-of-War," *Journal of Economic Perspectives*, 29(4), 2015, pp. 89-112, on the parallel growth of platform aggregators and specialist independents in US retail. Supplementary policy-side coverage: Brookings Institution and National Retail Federation reporting on the long-tail share of specialist independent retailers across categories.

ESSAY THREE

The Shadow Infrastructure of the Experiences Industry

The infrastructure problem the industry has been solving in the dark, one company at a time.

The 2023 Arival Guide to Channel Management reports a single load-bearing number. The average experiences operator works with fourteen distributors. The average large operator works with forty-seven. The average attraction works with more than forty, and many larger attractions report working with hundreds of distribution partners.¹ Each of those relationships, in its raw form, is a one-to-one connection that must be built and maintained on both sides.

Per-operator, fourteen is a large number. At industry scale, it is larger still. The Global Operator Landscape 4th Edition puts the global operator count at approximately 600,000.² The Technology & Connectivity survey, on a 3,500-plus operator response base, reports that roughly half of

operators run some form of booking system: 44% use vendor res-tech and 9% run an in-house system, with the remainder on no formal system at all.³ The arithmetic of distribution connectivity therefore runs against a denominator on the order of hundreds of thousands of operators and a numerator on the order of the tens of thousands of resellers that collectively sell their product. The position of this essay is that the aggregate engineering labour devoted to building, maintaining and replacing those one-to-one connections is the experiences industry's shadow infrastructure. Shadow, because it never appears on any single balance sheet. Infrastructure, because the industry does not function without it. The framing has its theoretical home in Star and Ruhleder's account of infrastructure as a layer that becomes visible only on breakdown, embedded inside the working systems that depend on it.⁴ The breakdown surface in this industry is the part operators see daily: manual extranets, stale availability, mapping errors, failed API calls, reconciliation work between calendars and manifests, the same integration written twice on either side of the same connection.

The mesh is conserved, not eliminated

The structural logic of channel management is straightforward. Without a channel manager, an operator's booking system fans out to each reseller individually: one operator multiplied by N resellers equals N connections to build and maintain. With a channel manager, the operator, or more precisely the operator's reservation system, connects

once to the channel manager, and the channel manager holds the mesh of reseller integrations on behalf of many suppliers and many reservation systems at once. Arival's 2023 channel-management guide visualises this distinction directly.⁵

A channel manager, in the precise sense of the term, is not itself a reservation system. Many reservation systems route bookings and availability to third-party channels as one feature among many; that functional overlap is not what a channel manager is. A channel manager is a tool whose singular purpose is to route content, availability, bookings and cancellations between suppliers and distribution channels. What it offers in exchange for its position in the stack is twofold. It consolidates access to resellers that would otherwise have to be integrated one by one; and it lets suppliers and reservation-technology vendors share the cost of building and maintaining those reseller integrations as requirements change. The value proposition is not a new sales channel, nor a new booking engine. It is shared infrastructure for a task that would otherwise be paid for many times over.

The important move is to notice what the second picture does not do. It does not eliminate the mesh. It shifts the mesh's locus. The reseller-side integrations still exist; they have been moved one layer up, into the channel manager's own engineering backlog. The operator has been spared a per-reseller integration; the channel manager has not. At industry scale, the mesh is conserved. The cost has moved. It has not disappeared.

The cost the industry under-reports

Per-integration engineering cost, in the publicly available Arival data, is expressed indirectly. The Technology & Connectivity survey records how operators actually manage distributor bookings: extranet (47% of tour operators, 41% of activities, 31% of attractions); email (37%, 39%, 27%); booking-system API (33%, 46%, 27%); channel manager (8%, 18%, 27%). Extranet and email are, by any definition, manual. Arival's own characterisation of the finding reads: "distribution connectivity remains fragmented and in many cases manual, despite significant efforts by booking systems and distributors to advance direct API connectivity with operators."⁶

The implication of those shares is not that API connectivity is impossible. It is that, for most operators below a certain size, the per-connection engineering cost exceeds the per-connection operational value. At the scales at which most operators trade (Essay 1 established the benchmarks: 86% under \$250,000 in annual revenue, 50% with fewer than five years of operating history), direct API connectivity is not affordable. The cheaper substitutes, hand-managed extranets and email, take its place. API penetration correlates strongly with operator size: the same Arival survey reports 59% of large operators (above 10,000 annual guests) using an API, against 29% of small operators.⁷ The integration tax is not evenly distributed. It is concentrated in the upper tail.

Integration is rent, not capex

Integrations are not permanent. APIs version. Authentication schemes evolve. Rate limits change. Endpoints deprecate. Every integration is a live relationship, subject to re-work on a multi-year cadence on both sides. The familiar operator-side mental model, that the integration is a capital expenditure to be made once and amortised over the life of the distribution relationship, is wrong. The integration is not capex. It is rent. The framing is consistent with the long-established finding from Brooks and the software-engineering literature that maintenance and evolution dominate the lifecycle cost of any non-trivial software system, often by an order of magnitude over initial build.⁸ Integrations are not exempt from this; if anything, they are a sharper case of it, because both sides of the connection are evolving in parallel.

The rent is paid on both sides of every connection: the operator-side cost (or, more commonly, the booking-system vendor's cost, internalised into the vendor's SaaS fee), and the reseller-side cost (or the channel-manager provider's cost, internalised into distributor-side fees and platform-level engineering budgets). The aggregate is genuinely large. The reason no research firm has published an industry-wide total is not that the total does not exist; it is that the cost is distributed across thousands of engineering backlogs, none of which is labelled "integration tax." That atomised incidence is the distinguishing feature of shadow

infrastructure. It is not absent from the industry's economics. It is invisible in the industry's accounting.

A conservative order-of-magnitude construction, clearly labelled as construction rather than a sourced total, makes the scale legible. Take the 600,000-operator base from the latest Global Operator Landscape; apply the roughly 54% booking-system-adoption rate; apply the 14-average-distributors-per-operator figure (which likely understates, because larger operators pull the mean sharply upward); assume a mix of direct-API, extranet, and channel-manager-mediated connections; apply multiple engineering-weeks per full integration on each side; apply a multi-year half-life for maintenance, version churn, and replacement. The arithmetic produces a figure in the low-to-mid hundreds of millions of US dollars per year at industry level, and that range is best read as a lower bound. The construction does not capture the reseller-side and aggregator-side engineering labour in full, nor the opportunity cost of the engineering time itself. A more complete accounting would land higher. This is not a sourced number; it is an order-of-magnitude argument built from per-operator data points that are. The point is not the specific figure. The point is that the figure is material at industry scale, and that the industry has never formally added it up.

Why standardisation has not settled it

In airlines, a standard settled the distribution-connectivity problem in the 1970s through the GDS architecture and IATA's settlement infrastructure. In hotels, partial

standardisation through the property-management-system-to-switch-to-distributor pathway reached sufficient adoption to contain the problem over the 1990s and 2000s. In experiences, standardisation has been attempted twice, and a neutral specification exists today.

The first attempt came from outside the category. The OpenTravel Alliance was founded in 1999 by a consortium of airlines, hotels, car-rental companies and travel-software vendors, with the stated purpose of building an XML-based messaging specification for electronic communication between disparate travel systems.⁹ Its first release appeared in 2000; its object-model 2.0 release, adding JSON support, followed in 2016. OpenTravel progressively expanded its coverage into rail, cruise, golf and, from 2008, tours and activities, with the tour-operator messaging built into a live implementer's system in 2009 and described at the time as a first-of-its-kind standards-based connection integrating tour operators into the broader online travel market.¹⁰ The commentary surrounding that release explicitly identified the tour-and-activity segment as the long tail of online travel, most of which had little or no ability to connect online with other suppliers or distributors.¹¹

Tens of thousands of OpenTravel messaging structures remain in use in the upstream segments (airline, hotel, car rental) the specification was built for. In tours and activities, adoption never crossed the threshold that would eliminate the bespoke integration layer. The reasons reward close attention. OpenTravel's governance and technical agenda reflected its founding constituencies, whose own distribution

problems were already substantially solved; tours and activities were an addition to the agenda, not its centre of gravity. And the supply side of the experiences segment, then as now, was too atomised for the OpenTravel membership to sustain a coordinated push into it.

The second attempt originated inside the category. OCTO (Open Connectivity for Tours, Activities and Attractions) was formally established as an industry association in October 2022.¹² Its first specification (OCTO 1.0) was published in mid-2023; its next major version (OCTO 2.0, with enhanced pricing models and improved booking flow) has been in community review through early 2026.¹³ OCTO is a 501(c)(3) non-profit, governed by an industry-led board drawn from across the category, and the specification and its documentation are free. It does not operate an API, a channel manager, or connectivity infrastructure itself.¹⁴ Its technical coverage is comprehensive for the category, spanning products, availability, bookings, pricing, supplier management, pickup and drop-off, and authentication.

OCTO is, by any reasonable standard, a credible standardisation effort. By a figure OCTO itself reported in 2023, 114 trading partners worldwide were working through the specification within its first year.¹⁵ And yet, four years after its formal launch, most of the industry is not yet connected through it. The specification exists; adoption has not reached the threshold that would eliminate the bespoke integration layer.

This is not a criticism of OCTO, which is doing everything a standards body can do. It is an observation about the

structural conditions that produce industry-wide standardisation in adjacent categories, and the extent to which those conditions hold in this one. Historically, industry-wide standardisation has required at least one of three forces: a dominant player whose scale compels compliance (the airline-GDS case); a regulator with standing over the interchange (the banking ISO 20022 case); or a cooperative adoption threshold, at which a critical mass of major participants converges simultaneously and the remaining holdouts are forced to adopt or be excluded.¹⁶ Experiences has none of the three in the shape the adjacent industries had. The supply side is too fragmented for any single participant to force compliance, for the reasons Essay 2 established. No regulator has standing over industry-wide distribution APIs. And the cooperative-threshold condition has not been met, because the individual incentive for each participant to maintain its own integration relationships has so far been stronger than the collective incentive to converge.

There is a sharper version of that last observation worth making directly. The category's largest online travel agents and attraction-side platforms have, over the past decade, invested substantially in proprietary connectivity stacks of their own. Those stacks are, from the platform's own vantage, competitive assets rather than neutral pipes: the engineering investment is sunk, the reseller-side and supplier-side integration portfolio is extensive, and in several cases the connectivity layer is actively positioned as a source of advantage over smaller distributors. In that frame, a

neutral specification is not merely redundant; it is strategically undesirable, because it lowers the distribution-integration cost for competitors who have not made the same investment. The participants whose adoption would most plausibly trigger convergence are, on their own unit economics, the participants with the strongest individual incentive not to converge. This is incentive analysis, not a charge of bad faith: rational actors with sunk investment in proprietary stacks face a different convergence calculus from actors without it. Until that calculus changes, mass adoption of a neutral specification will lag, whatever the technical merits of the specification itself.

That is the structural reason the integration tax persists. A neutral specification exists and is available. The forces that would otherwise produce industry-wide adoption are weaker in this industry than in the adjacent industries a reader might assume comparable, and the actors with the most leverage to move the market toward convergence have the clearest reason not to.

The two counter-arguments

Two counter-arguments deserve to be named directly.

The first is that every industry has integration overhead, and that this is simply the experiences-industry version of a normal cost of doing business. The response is a question of order of magnitude. The average experiences operator maintains an order of magnitude more distribution relationships than a typical hotel; it does so without the convergent standard that hotels reached over two decades;

and the median operator sits so far below the scale at which conventional integration economics work that the cheaper substitutes (email, extranet, manual posting) remain in broad daily use. The problem is not that experiences have integration overhead. The problem is that the scale, churn rate, and distribution of that overhead are non-normal.

The second counter-argument is that integration is just the normal IT cost of running a software business, absorbed into a SaaS vendor's P&L like any other engineering expense. This is partially true for the individual vendor. What it misses is that the cost is not contained inside one vendor's P&L. It is spread across operator-side vendors, reseller-side vendors, channel-manager intermediaries, aggregator engineers, in-house operator teams, and distributor-side IT departments. The atomised incidence is the distinguishing feature. Atomised incidence is not the same thing as small cost.

What comes next

The useful question, accepting the shadow-infrastructure frame, is no longer how any individual operator can reduce its integration overhead. That question is bounded by the per-connection engineering cost, which is itself determined by a standardisation environment the operator does not control. The useful question is what the industry stack looks like if the shadow infrastructure is priced honestly, engineered once, and maintained at scale by actors whose economics do not depend on capturing the adjacent distribution flow. What the shape of such a layer would be,

and whether it is emerging, is the question the essays that follow take up.

Before that question can be addressed well, the reader needs a concrete picture of what actually happens in a single booking: the systems it touches, the API calls it generates, the data-format translations it requires, the latency budget it operates within. The next essay traces one. The shadow infrastructure becomes legible when you watch it work.

The industry has been paying for this plumbing out of thousands of separate engineering budgets, for more than a decade, without ever adding up the bill.

ENDNOTES

1. Arival, *Guide to Channel Management: Strategy, Connectivity & Technology*, November 2023, p. 14 (the 14-distributor average and the larger-operator and attraction figures are drawn from Arival's own 2023 Technology & Connectivity survey and reproduced in the channel-management guide).
2. Arival, *Global Operator Landscape 4th Edition: The State of Experiences*, January 2026. The 600,000 figure is an Arival industry estimate rather than a precise enumeration.
3. Arival, *Technology & Connectivity: The state of booking system adoption and satisfaction, distribution connectivity, and tech priorities for tours, activities & attractions*, January 2023, p. 7 (adoption breakdown: 44% vendor res-tech, 9% in-house, 46% no system, with minor rounding).
4. Star, S. L., and Ruhleder, K., "Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces," *Information Systems Research*, 7(1), 1996, pp. 111-134, on infrastructure as a layer that becomes visible only on breakdown and is otherwise embedded inside the working systems that depend on it. The "shadow infrastructure" framing here is an applied case of that broader theoretical point.
5. Arival, *Guide to Channel Management*, November 2023, pp. 9 and 11 (Figure 1 and Figure 2: the without-and-with-channel-management diagrams).
6. Arival, *Technology & Connectivity*, January 2023, p. 12 (quoted verbatim), with the underlying channel-breakdown figures on the same page.
7. Arival, *Technology & Connectivity*, January 2023, p. 12.
8. Brooks, F. P., *The Mythical Man-Month: Essays on Software Engineering*, Addison-Wesley, 1975 (anniversary edition 1995), on the dominance of maintenance and evolution costs over initial-build cost in

- any non-trivial software system. See also Lehman, M. M., "Programs, Life Cycles, and Laws of Software Evolution," *Proceedings of the IEEE*, 68(9), 1980, pp. 1060-1076, which formalises the requirement of continual change for software systems to retain operational value.
9. OpenTravel Alliance, About Us (opentravel.org/about-us), accessed April 2026; see also OpenTravel Alliance, Wikipedia (accessed April 2026), for the 1999 founding date, the 2000 first release, and the 2016 2.0 release adding JSON to the XML schema.
 10. PEAK 15 Systems press release, OpenTravel Member PEAK 15 Systems Connects Tour Operators and Travel Agents, 2009 (archived at travelindustrywire.com), reporting that OpenTravel's tour-operator messaging was developed and published in 2008 and built into PEAK 15's next-generation system in 2009.
 11. Adventure Travel Trade Association, OpenTravel Standards Connect Tour Operators and Travel Agents for The First Time (adventuretravelnews.com), reporting the tour-operator-and-outfitter segment as the long tail of online travel, with the vast majority of the then-8,000-plus operators and outfitters in the U.S. and Canada having little or no ability to connect online with other suppliers or distributors.
 12. OCTO announcement and Travel Weekly Asia coverage of OCTO's first members' meeting on 10 October 2022 at Arival 360 | Las Vegas (both October 2022).
 13. OCTO, Specification (octo.travel/specification), accessed April 2026. OCTO 1.0 was published in mid-2023; OCTO 2.0 is reported as being in community review through early 2026.
 14. OCTO, governance overview (octo.travel), accessed April 2026.
 15. OCTO, OCTO is a year old and growing, 2023, reporting 114 trading partners working through the specification within its first year.
 16. The three-forces account of industry-wide standardisation draws on the network-economics literature, principally Shapiro, C., and Varian, H. R., *Information Rules: A Strategic Guide to the Network Economy*, Harvard Business School Press, 1999, particularly chapters 7-9 on networks, standards and lock-in, and Farrell, J., and Saloner, G., "Standardization, Compatibility, and Innovation," *RAND Journal of Economics*, 16(1),

1985, pp. 70–83. For the airline-GDS case specifically, see Copeland, D. G., and McKenney, J. L., “Airline Reservations Systems: Lessons from History,” *MIS Quarterly*, 12(3), 1988, pp. 353–370. For the hotel-distribution case, see O’Connor, P., *Electronic Information Distribution in Tourism and Hospitality*, CABI Publishing, 1999. For the collective-action lens on why industry-wide standardisation requires conditions beyond individual incentive, see Olson, M., *The Logic of Collective Action: Public Goods and the Theory of Groups*, Harvard University Press, 1965.

Anatomy of a Booking

A booking is a polite fiction. Trace what it actually has to do and the industry's shape becomes legible.

At ten past nine on a Friday evening, a family of four in a hotel room in Munich chooses a day tour to Neuschwanstein for the following morning. One of them opens a distributor's app, scrolls the results, taps the tour with the best reviews, enters names and a card number, and receives a voucher in around twenty seconds. From the outside, this looks like a small act of consumer commerce, not meaningfully different from ordering a taxi.

The purpose of this essay is to describe what actually happens during those twenty seconds, and during the hours and days on either side of them. The previous essays in this series have traced the industry's structural position (digitisation lag, fragmentation, shadow infrastructure) at a general level. The picture is still abstract. Before the arc can turn to what an alternative industry stack might look like, the existing stack needs to be legible. One booking, traced

from consumer tap to operator manifest, is the cheapest way to make it legible.

The boundaries a booking crosses

A consumer booking in this industry typically crosses six boundaries between systems that belong to at least three different organisations.

The first is between the customer's device and the distributor they are buying from. The second is between the distributor's customer-facing layer and its back-end, where cart and pricing logic live. The third is between the distributor's back-end and whatever it uses to reach supply (a direct operator connection, a neutral specification such as OCTO, a channel manager, or a large marketplace's proprietary supplier interface). The fourth is between that connectivity layer and the supplier-side reservation system. The fifth is between the reservation system and the operator's own database and daily manifest. The sixth is between the operator's system and the voucher that eventually reaches the customer.

Each boundary has its own authentication, its own data model, and its own ways of failing. A successful booking is the concurrent success of all of them. A failed booking is the failure of any one.

The sequence

Underneath the family's four screens (product, calendar, checkout, confirmation), the conversation between systems proceeds in roughly the following order.

Long before the family opens the app, the distributor already has a product catalogue for this supplier in its own systems. In a fully automated topology, that catalogue is pulled from the supplier's reservation system on a continuous background cadence. In practice, for a substantial share of suppliers in this industry, the catalogue is maintained by the supplier directly inside a distributor extranet, a web portal where prices, photographs and descriptions are entered and updated by hand.¹ The consumer cannot tell the difference between the two routes; the operating cost and the staleness risk behind each are not the same. For the rest of the sequence, it does not matter which route the content took, only that the distributor now has something to show.

When the family chooses tomorrow's early-morning departure, the distributor asks the supplier's system whether that exact slot is still open, with that exact mix of passengers, at the price the customer is seeing. This is the first synchronous conversation in the sequence, and it happens inside the twenty-second envelope the customer is willing to tolerate. Assuming the answer is yes, the distributor places the seats on hold. The hold is a temporary reservation that protects those seats while the customer enters card details. Holds carry a time limit, typically measured in minutes.² If the customer walks away, the seats return to inventory.

The card authorisation happens on a separate conversation with the payment processor. This is a different organisation, a different data model, a different set of failure

modes. It is on the critical path of the booking, but not on the booking connectivity path.

Once payment is authorised, the hold is converted into a firm booking. The supplier-side system returns a confirmation identifier and a voucher or scan-ready token. That voucher propagates back through the chain to the customer, and independently into the operator's daily manifest. The operator is now expecting the family at the Munich pick-up point the following morning, and the family has a record they can present on arrival. Between ten and twenty separate conversations between systems, in most topologies, across the six boundaries above (and sometimes more, when the supplier itself calls further upstream), all inside the time it takes a consumer to type a card number.³ The booking, as the consumer experiences it, is a polite fiction stitched together in real time across systems that do not share an owner.

The translation burden

The same booking is described differently at every boundary. At the first boundary, it is four names and a card. At the distributor's back-end, it is a cart line item, a pricing plan, and four roster entries. At the connectivity layer, it is reshaped into whatever data model the supplier-side system expects: units, age bands, ticket categories, unit types. The supplier's reservation system may map those categories again to whatever the operator has configured locally, which may be "adult," "student," "family," with its own pricing

rules attached. On the way back, the voucher carries a pax count in a format the distributor can display.

Each translation is a place where data can drift, and each is a place where a version change somewhere else in the chain can break something here. When the operator introduces a new pricing category, the mapping on the reservation system changes, the connectivity layer's translation changes, and the distributor's catalogue must pick up the change. Each of those changes is an engineering ticket somewhere.

In the aggregate, the translation layer is where the permanent cost lives. The first integration build is a one-off expense. The translation layer is paid every time any one of the data models in the chain changes, forever. That is why Essay 3 argued for integration as rent rather than capex. This is where the rent is paid.

The speed constraint

A checkout experience that takes longer than a few seconds loses customers. The thresholds are well established: about one second before a user's flow of thought is interrupted, ten seconds before attention is lost altogether.⁴ All of the synchronous conversations above have to complete within that window, or the distributor has to build fallback behaviour that hides the latency.

Most of the industry's published guidance to partner developers accommodates upstream response variance measured in tens of seconds on individual calls; in practice the budget at the customer's end is much tighter than that.

The slower the supplier-side system, the faster everything else in the chain has to be. If a channel manager sits between the distributor and the supplier, each call traverses an additional hop on the way out and another on the way back, consuming more of the same budget. If the supplier must itself call further upstream, for example to a partner for a combined-ticket product, the depth of the chain increases again.

The operational reality is that the latency the customer perceives is the slowest boundary in a chain, compounded by the hops between them. The faster boundaries are doing work no customer will ever attribute to them. The slowest boundary is absorbing reputational cost for the whole stack.

Where the interesting weight lives

The happy path is not where the weight of the system sits. The weight sits in the failure modes.

Two customers, in different cities, attempt to book the same last seat in the same second. Hold semantics exist to prevent this, but holds themselves temporarily consume inventory on bookings that may never confirm, which creates an inventory-accuracy problem of its own. The trade-off, between accurate inventory and serving every customer who tries to book, is the everyday face of Brewer's consistency-versus-availability problem in distributed systems.⁵

The family enters card details, gets distracted, closes the app. Fifteen minutes later, the hold expires and the seats return to inventory. If the payment processor has already authorised the card, the distributor is now holding an

authorisation for a booking that no longer exists. Reconciling authorisations to bookings is a full-time job at any distributor at scale.

The price the family saw at 21:09:30 came from the availability check. Between then and booking confirmation, the supplier republished prices for tomorrow. The distributor must decide whether to honour the original price, reprompt the customer, or absorb the difference. Each of those choices has an operational cost and a customer-experience cost.

The booking confirms on the supplier side, but the voucher email does not reach the customer because a mail provider marks it as spam. The operator has the family on tomorrow's manifest; the family has no record. The operator learns of the booking when the family does not appear at the pick-up point. Recovering the situation requires three organisations to reconstruct state from logs.

The supplier's system is unreachable for six minutes. During those six minutes, every availability check for that supplier fails. The distributor either caches stale availability, and risks overbooking, or returns "no availability" to every customer who searches, and loses bookings the operator would gladly have taken. Either way, the cost is real and distributed across organisations that never see each other's logs.

Every working distributor and every working supplier-side system maintains internal tooling to detect and respond to these situations.⁶ None of that tooling shows up on a balance sheet as infrastructure. It is, again, the shadow

infrastructure of Essay 3, paid for engineer by engineer, ticket by ticket.

Why vertical integration matters structurally

When a distributor and a reservation system sit inside the same corporate perimeter, the third, fourth and fifth boundaries collapse. They stop being cross-organisation interfaces, and become internal function calls inside a single company, against a single data model controlled by a single product team. The consumer-facing experience is indistinguishable from the non-integrated version. The underlying operation is materially cheaper to run, and the failure modes are fewer and easier to diagnose.

This is the structural advantage the category's largest online travel agents have acquired by buying reservation-system vendors in the layer beneath them. It is rational on their individual economics. It also compounds asymmetry across the rest of the industry, because the distributors that have not vertically integrated, and the suppliers that distribute through them, continue to pay the cross-boundary cost on every booking. Essay 3's argument about a neutral specification lands here concretely: the participants with the most leverage to move the industry toward a neutral layer are precisely the participants who have already solved the boundary problem inside their own walls.

What the picture suggests

Accept the topology as described. Six boundaries, a dozen conversations, a translation layer at every interface, a

continuously operating failure-response stack that nobody accounts for as infrastructure. The natural question is not how to make any one boundary faster. It is whether the topology itself is inevitable, and whether a neutral layer could emerge that collapsed some of these boundaries for participants who cannot buy their way into vertical integration.

The next essay takes up the analogy the industry reaches for most often when trying to think about exactly that possibility. It does the Amadeus-and-Sabre comparison honestly, and shows how much of what made a neutral layer possible in aviation is absent in experiences, and what that absence actually implies.

The system works. That it works at all is the accomplishment. That it works this way, for everyone except the vertically integrated, is the problem.

ENDNOTES

1. Arival, Technology & Connectivity: The state of booking system adoption and satisfaction, distribution connectivity, and tech priorities for tours, activities & attractions, January 2023. The description of distribution connectivity as “fragmented and in many cases manual” (p. 12) is the principal field reference for this essay, and the underlying channel-breakdown figures in the same report quantify the share of supplier content maintained via distributor extranets rather than via direct API.
2. Public partner documentation published by the category’s largest marketplaces describes hold-and-confirm booking flows and reservation-expiration windows measured in minutes, consistent with the sequence described in this essay.
3. OCTO, Specification overview (octo.travel), accessed April 2026. The core specification covers products, availability, bookings and associated operations; optional capabilities extend the specification to pricing, pickups, drop-offs, content, and notifications. The topology described in this essay is implementation-agnostic and holds whether the underlying connectivity uses OCTO, a proprietary marketplace specification, or a direct operator-to-distributor integration.
4. Nielsen, J., Usability Engineering, Academic Press, 1993, particularly the chapter on response times, which establishes the 0.1-second, 1-second, and 10-second thresholds that bound user perception of system responsiveness. The framing has carried forward into all subsequent web-latency research.
5. Brewer, E. A., “Towards Robust Distributed Systems,” keynote, ACM Symposium on Principles of Distributed Computing (PODC), 2000; formalised in Gilbert, S., and Lynch, N., “Brewer’s Conjecture and the Feasibility of Consistent, Available, Partition-tolerant Web Services,” ACM SIGACT News, 33(2), 2002, pp. 51–59. Together these establish the consistency-versus-availability trade-off that any distributed booking layer must make.

6. Beyer, B., Jones, C., Petoff, J., and Murphy, N. R. (eds.), *Site Reliability Engineering: How Google Runs Production Systems*, O'Reilly Media, 2016. The canonical practitioner reference for operating distributed systems at scale, including failure-mode detection, on-call response, and the engineering practices that sit beneath any working booking infrastructure.

ESSAY FIVE

What Amadeus and Sabre Don't Teach Us About Experiences

The GDS analogy is the one the industry reaches for most often. It does not transfer.

Having traced what the existing topology of a booking looks like in Essay 4, the natural next question is whether a different topology could emerge. Every serious conversation about infrastructure in the experiences industry eventually produces some version of the same sentence. “Who will be the Amadeus of experiences?” Or “this industry needs its own Sabre.” Or “we are still waiting for a GDS.” The sentence arrives with the weight of analogy, which is a powerful kind of weight, because it lets the speaker import a conclusion (a neutral infrastructure layer eventually emerges and extracts real value) without importing any of the structural work that conclusion depends on.

The analogy is worth taking seriously. Not to settle who will win a race no one can plausibly name the finish line of, but to do the aviation comparison honestly, and to notice how little of what made the airline clearinghouse layer possible is present in the market we actually operate in. What transfers is instructive. What does not transfer is more instructive still.

What a GDS actually was

The contemporary reader encounters “GDS” as a category of travel-technology company. The historical reality is narrower and more specific. SABRE, launched by American Airlines and IBM in the early 1960s, was built as the carrier’s internal reservation system before it was externalised to travel agents in the mid-1970s.¹ Apollo, at United, and later Galileo, followed similar paths from the US side. Amadeus was assembled in 1987 as a European consortium between Air France, Lufthansa, Iberia and SAS, explicitly to put European carriers on equal footing with American ones that had already productised their reservation systems. Worldspan followed in 1990.² The pattern is recognisable once named. The infrastructure emerged inside airlines, became external infrastructure over the course of a decade, and was then, through a combination of regulation and commercial pressure, divested into neutral ownership.

Three structural facts about this trajectory matter, and are usually glossed in casual analogy. The layer was not invented from outside the industry it eventually served; it was extruded from the carriers that needed it. Its neutrality was an institutional outcome, not a design choice. And it solved a

specific pre-internet problem (aggregated airline inventory visible to travel agents) that was bounded by the commercial shape of aviation in that period.

What transfers

Three structural lessons from that period do carry over to any market trying to develop a comparable layer. They are worth stating plainly, because they are almost the only things in the GDS story that do.

Neutrality. A clearinghouse functions as a clearinghouse only to the degree that its operator is not also a distributor. Once the clearinghouse competes with its own buyers, its buyers route around it. This constraint is why the airline GDSs ultimately had to be divested from carrier ownership, under the regulatory rules that established neutrality requirements as a precondition of operating the layer at all.³ The comparable-layer question in any other market is always, at root, a question about ownership structure.

Network effects. Supply and demand concentrate on one another, once concentration starts. The network effect is unglamorous and difficult to seed, but once a layer has enough sellers on one side that buyers find it efficient to start there, the process runs on its own.⁴ The GDSs benefited from a period in which travel agencies were commercially indispensable and the airline commission structure rewarded aggregation on their side.

Transaction-fee economics. The unit economics of a winning infrastructure layer are transaction fees, not software licences. The layer that charges by the booking has

a cost base and a revenue base that scale together. The layer that charges by the seat, the user, or the module eventually finds itself misaligned with the volume passing through it. This lesson is structural and survives the specific commission rates of aviation in the 1980s.

These three points transfer. Most of the rest does not.

What does not transfer

Four asymmetries stand between aviation in the GDS era and the experiences industry today. Each is structural enough that the analogy cannot be imported wholesale past it.

Inventory cardinality. An airline in 1980 sold a bounded vocabulary of products. A schedule, a route, a flight number, a published fare, a class of service. Tens of thousands of distinct units existed, but the schema describing them was small and well understood.⁵ A neutral layer could be built on a finite taxonomy because the taxonomy was finite. An experience does not have a bounded vocabulary. A single operator may sell the same underlying product in three languages, from four meeting points, with two pick-up variants, in private and shared configurations, combined with transfers, combined with meals, with and without a kid seat, on seasonal hours, under weather-dependent operations, and with crew-dependent capacity. Across the industry, the permutations are effectively open. Standardisation in this market is not a taxonomy problem; it is a translation problem. Those are different disciplines, and the second is harder.

Buyer structure. The aviation GDSs formed in a world where demand aggregated through a small and commercially coherent set of intermediaries. Travel agencies operated on predictable commission economics, corporate travel departments routed through a handful of managed providers, and the airline commission cut produced strong incentives for consolidation on the agent side. Experiences has no equivalent buyer-side concentration. The demand side today is an uneasy mix of large online marketplaces, destination-marketing organisations, hotel concierges, ground handlers, tour resellers, and hundreds of thousands of direct bookings mediated by search.⁶ None of these produce the buyer-side centre of gravity that made the airline clearinghouse economically coherent. A neutral layer in aviation had ten doors to knock on. A neutral layer in experiences has a continuous distribution of them.

Institutional apparatus. The aviation industry had, and still has, an institutional layer that predated the GDS and made many of its choices possible. IATA sets schedule standards, ticketing conventions, and settlement rules. National regulators and route authority govern who can fly where. Airline alliances produce shared codes and interline agreements.⁷ Much of what appears in retrospect to be technical standardisation in aviation is actually institutional standardisation, enforced by bodies with real authority over market participants. Experiences has no equivalent. There is no IATA, no route authority, no universal product taxonomy, and no institution with the standing to impose one. A

clearinghouse cannot be designed by assuming an institutional scaffold that does not exist.

Settlement infrastructure. The Billing and Settlement Plan on the international side, and the Airlines Reporting Corporation on the US side, clear money between agents and carriers under rules negotiated by the industry through IATA.⁸ This is not a minor administrative detail. It is a neutral, industry-operated piece of financial plumbing without which the commercial logic of a neutral inventory layer would not hold together. Experiences settles peer-to-peer. Commission terms vary by channel. Payment timing varies by operator. Refund policies vary by product. Voucher redemption practices vary by destination. A neutral inventory layer is easier to imagine than a neutral settlement layer. Together they made the aviation clearinghouse load-bearing.

Each of these asymmetries on its own is enough to weaken the analogy. Together they change what kind of problem the experiences infrastructure question actually is.

The false confidence the analogy produces

The temptation of the analogy is that it smuggles in a conclusion. A neutral layer exists in adjacent travel categories; therefore, the reasoning runs, it must eventually exist here. The smuggled part is the preconditions. Aviation had finite inventory vocabulary, buyer-side concentration, an institutional standardiser, and a settlement plan before it had a clearinghouse. Experiences has none of these in comparable form. Any reasoning that imports the conclusion

without importing the preconditions is doing less work than it appears to.

What the actual industry conversation reveals is that the analogy smuggles two opposing conclusions, not one. At the Arival Executive Summit in Valencia in 2026, a participant whose background was in the hotel sector described the GDS layer as exactly the kind of solution the experiences industry needs. A senior executive of a major online travel agent in the same conversation described the same layer as exactly what their organisation would resist, on the grounds that a non-fragmented supply side would dissolve the structural advantage their position depends on. The analogy was doing the work of two opposing arguments at once. Both speakers were right about the implications for their own position.

This matters because the industry's infrastructure debate is mostly a debate about conclusions. Who will build it. Which company is positioned to win it. When it will consolidate. The more useful conversation begins from the preconditions and reasons forward.

What the analogy, taken seriously, actually suggests

If the structural conditions that produced the aviation clearinghouse do not exist in experiences, the interesting question is not who will be the Amadeus of this industry. It is under what conditions a neutral layer could emerge in a market with open-cardinality inventory, no buyer-side concentration, no institutional standardiser, and no settlement infrastructure.

Some of those conditions may be substitutable. A translation-based standardisation model may do some of the work a fixed taxonomy would have done in aviation. A specification with enough voluntary adoption across the category may produce a soft institutional apparatus over time, without a regulator ever convening. A settlement layer may be assembled commercially rather than institutionally, by participants who agree to route payments through a shared piece of plumbing because it is cheaper than the alternative.

None of these substitutions is automatic, and each of them carries its own set of structural trade-offs that the GDS history does not help us reason about. The industry is not owed the aviation shape. It will develop, if it develops at all, a shape of its own.

Close

The analogy is tempting because it offers a familiar outline for a market whose shape is still being drawn. Familiar outlines are worth something, but not more than the preconditions that drew them. The experiences industry needs infrastructure (Essays 3 and 4 made that case at the operational level), but it does not need a clone of the aviation clearinghouse. It needs a layer shaped to its own supply shape, buyer structure and institutional reality. The clearest partial answer to the infrastructure question that the experiences industry has produced so far is the channel manager, which began as a routing tool for a narrower problem and has been pushed, over two decades, to carry weight it was not built to carry. The next essay takes that layer on its own terms.

ENDNOTES

1. Copeland, D. G., and McKenney, J. L., "Airline Reservations Systems: Lessons from History," *MIS Quarterly*, 12(3), 1988, pp. 353-370. Sabre historical filings and ATPCO fare-publication documentation underpin the chronology of the SABRE programme from American Airlines and IBM in the early 1960s through its externalisation to travel agents in the mid-1970s.
2. Amadeus historical filings on the 1987 European consortium between Air France, Lufthansa, Iberia and SAS; Worldspan formation history (1990); Galileo and Apollo historical chronologies, drawn together with Copeland and McKenney 1988.
3. United States Department of Transportation, Computer Reservation System rules, 14 CFR Part 255, originally promulgated 1984. The regulation that required GDS divestiture from airline ownership and established the neutrality requirements that made the layer load-bearing.
4. Rochet, J.-C., and Tirole, J., "Platform Competition in Two-Sided Markets," *Journal of the European Economic Association*, 1(4), 2003, pp. 990-1029. Foundational two-sided-market theory; supports the argument that supply and demand concentrate on each other once a threshold is reached.
5. OAG (Official Airline Guide) historical schedule data and ATPCO published-fare history together document the bounded product space of pre-deregulation and early-deregulation aviation: a finite combinatorial of schedules, routes, flight numbers, published fares and classes of service, on the order of tens of thousands of distinct units globally at any given point in the period the GDSs were built to serve.
6. Arival, *Technology & Connectivity: The state of booking system adoption and satisfaction, distribution connectivity, and tech priorities for tours, activities & attractions*, January 2023, on the variety of distribution channels and intermediaries operators use.

7. Doganis, R., *The Airline Business*, Routledge, 2nd edition 2006, on the institutional scaffold of aviation including IATA standards, route authority, alliance interlining, and the regulatory framework within which the GDS layer operates.
8. IATA, Billing and Settlement Plan documentation; Airlines Reporting Corporation, ARC system documentation. Together these clear money between agents and carriers under industry-negotiated rules.

Channel Managers Are a Partial Answer

The channel-manager category solves for shared integration cost. It does not solve for coherence.

Essay 5 ended by naming the channel manager as the clearest partial answer the experiences industry has produced. Partial is the operative word. The channel manager is not glamorous. It does not appear in industry keynotes under its own name. It does not attract the attention that consumer-facing marketplaces do, or the attention that reservation systems attract when they touch the operator's daily work directly. And yet much of the industry's non-direct booking flow passes through this layer, and has done so since the category emerged as a distinct piece of infrastructure.

If the infrastructure question is going to be asked properly, the category that sits in the middle of most current traffic is the place to start. The purpose of this essay is to describe,

carefully, what the channel manager is, what problem it solves, what problems it was not designed to solve, and why it has ended up carrying many of those additional problems anyway. None of this is criticism of the category. A solution that fits its original problem well has no obligation to fit every problem that shows up afterwards. The question is how much of the weight the infrastructure of this industry now needs to carry can continue to be carried by a category built for a narrower purpose.

What the channel manager is, precisely

A channel manager is a routing tool. It is not a reservation system. A reservation system holds inventory, takes bookings, manages refunds, and produces the operator's daily manifest. A channel manager does none of those things. It routes content, availability, bookings, and cancellations between a reservation system on the supplier side and a set of resellers on the distribution side. Its purpose is singular,¹ though its commercial neutrality is not: some channel managers operate as standalone routers, others sit inside reservation-system or distribution-side organisations. Three kinds of participants use it, and the distinction matters for what the category is asked to do.

Suppliers. Operators running their own reservation system use a channel manager to expose inventory to many resellers without building and maintaining a direct integration to each one. The alternative is point-to-point integration: the operator's engineering team builds one connection to one reseller, another to another, and so on. Each integration is a

line item that consumes engineering time, carries a deprecation lifecycle, and breaks when either side changes. The channel manager collapses this cost by replacing many operator-to-reseller connections with one operator-to-channel-manager connection, which then fans out to many resellers.

Res-techs. Other reservation systems that serve a book of operators use a channel manager for a structurally similar reason. A reservation system used by a hundred or a thousand operators does not want to maintain direct integrations to every reseller either. By routing its operator base through a channel manager, it shares the integration-build and maintenance cost across the aggregate and gains, by aggregation, easier access to resellers that might not otherwise negotiate individually.

Resellers. A large reseller that wants access to thousands of operators does not want to integrate with every one of them directly. A channel manager offers the reseller a consolidated API surface through which many operators can be reached. This is the category's second economic good: consolidated access.

The singular purpose is routing. The two economic goods the category delivers, shared integration cost and consolidated access, follow from that purpose.²

What the channel manager solves well

These two economic goods are real. Neither has been replaced by any other category since channel managers have been in widespread use. No subsequent layer of the stack

has made point-to-point integration cheaper, and no subsequent layer has given resellers a better route into the fragmented supply base at scale.³ The channel manager's durability is not an accident of incumbency; it is a consequence of solving a specific problem well enough that the solution has not been dislodged.

It is important to state this clearly before the rest of the essay becomes about what the category does not do, because the rest of the essay will be about what the category does not do. Both descriptions can be true at the same time. A category can solve its original problem well and still be asked, by the market's evolution, to solve problems it was not designed for.

What the channel manager was not designed to solve

Four structural problems sit beyond the category's original design envelope. None of them are failures of execution. They are features of scope.

Real-time availability consistency across deep chains. Every system in the booking chain has its own latency budget. The channel manager sits in the middle of that chain, and each hop in either direction adds to the budget. When the chain is two hops deep, the compounding is manageable. When the chain is three or four hops deep, as it often is in practice, the time any individual boundary has to complete its work shrinks toward zero. This is not a failure of the channel manager; it is a structural property of routing. A router routes, and cannot eliminate the cost of the routes it carries. Genuinely sub-second availability consistency across

a deep chain is not achievable by adding more routing. It would require the chain to get shorter, not longer.⁴

End-to-end data integrity across boundaries. Each end of a channel-manager connection has its own schema. The supplier's reservation system describes its inventory in one product vocabulary: age bands, pricing categories, unit types, meeting-point definitions. The reseller expects its own vocabulary on the way in. The channel manager negotiates a translation between them at the point of integration. That translation is a choice. When the reseller's schema does not map cleanly onto the supplier's, the channel manager picks a mapping. The mapping is an engineering decision made once and carried in every subsequent booking that passes through it. It is not possible to route heterogeneous schemas without picking compromises, and picking compromises is the work.⁵ Integrity, in the sense of no drift between what the operator thinks it is selling and what the reseller thinks it is buying, is bounded by the quality of the mappings and by the rate at which either schema changes beneath them.

Distribution economics, at structural scale. This limitation deserves more care than the others. Channel managers do influence distribution economics at the margin. Some negotiate preferential commission terms on behalf of their book of operators, offering resellers volume in exchange for rate. Some offer standardised payment and refund frameworks that their larger customers use as a baseline. These are real contributions, and they are not trivial. But the underlying commercial relationships, the commission cut itself, payment timing, refund policies, voucher-redemption

logistics, chargeback handling, remain negotiated out of band, directly between suppliers and resellers, or between their parent organisations. The channel manager carries the booking through. It does not carry the commercial framework. In other words, the category shapes economics at the edges without reshaping them structurally. The distinction matters, because the industry's distribution-economics problems are not edge problems; they are structural.

Neutral settlement. Essay 5 argued that the Billing and Settlement Plan on the international side and the Airlines Reporting Corporation on the US side were preconditions for the aviation clearinghouse, not afterthoughts.⁶ Experiences has no equivalent, and the channel manager has not become one. Money does not flow through the channel manager. Settlement is peer-to-peer between reseller and supplier, or between reseller and an intermediary separate from the connectivity layer entirely. This is not an oversight of the category; it was never part of the category's design. A routing tool routes data. A settlement layer routes money. They are different pieces of infrastructure, with different legal, regulatory, and commercial requirements. Building one does not imply building the other. The consequence reaches Essay 5 directly. If a neutral inventory layer without a neutral settlement layer is only half of what made the aviation pattern work, the closest thing the experiences industry has produced to a neutral layer does not address the other half at all.

Each of these four is a structural feature, not a flaw. Together they define the shape of the category's original envelope.

Why the category gets pushed past its scope

A category designed to do one thing well sometimes ends up doing many things imperfectly, because two pressures pull on it simultaneously and there are no alternative layers to which it can delegate.⁷

The external pressure is the absence of adjacent layers. In adjacent industries, a newly emerging distribution problem might be absorbed by a different piece of infrastructure: a standards body, a settlement network, a regulator, an institutional intermediary. The experiences industry has few of these. When a new distribution-connectivity problem appears, the only layer close enough to absorb it is, in most cases, the channel manager. Content-freshness management on behalf of operators who do not automate content, schema reconciliation between resellers with conflicting product models, pricing synchronisation under dynamic availability, workflow support for cancellations and refunds, each of these has been handed to the channel-manager category over time, mostly by default.

The internal pressure is competitive dynamics within the category itself. A channel manager that implements more reseller-specific features, more pricing logic, more content-management sophistication, more workflow assistance, tends to win commercial deals against one that does less. The economic incentive is to accumulate scope, not to stay lean.

A category pushed in this direction will absorb adjacent problems, even those it was not designed for, as long as customers reward the absorption with spend.

Both pressures are rational, and together they explain why the channel manager today carries a great deal of weight that the category in its original form did not. That weight is real work, done by real engineering teams, resolving real operational problems the rest of the stack does not resolve. Calling this a feature of the category's evolution is analytically accurate. Calling it the category's long-term equilibrium may not be.

The structural ceiling

A category can be pushed some distance past its original design envelope before the strain becomes visible from outside. That distance has been considerable in the channel-manager case. The category has absorbed schema heterogeneity, content-management responsibilities, pricing-synchronisation complexity, and workflow functions its original design did not contemplate. Each absorption was individually reasonable. The cumulative effect is that the category now sits under commercial and operational pressures it was not built to hold.

The signs of structural strain are becoming visible now. Chains of intermediation have grown deeper, not shallower, because new resellers often reach the supply base through another layer of aggregation rather than directly. The latency budget has not grown more generous in response. Translation burden has grown with the diversity of product

vocabularies, because the industry has not converged on one and shows no sign of doing so.⁸ Larger resellers, responding to the compounding cost of deep chains, have invested in proprietary supply stacks that let them bypass parts of the channel-manager layer and reach operators more directly, through their own connectivity specifications, their own integrations, or their own acquisitions of reservation-system vendors upstream.⁹ A category whose largest customers are building infrastructure to get around parts of it is a category showing structural strain.

None of this is a failure. The channel manager's core economic goods, shared integration cost and consolidated access, continue to be real, and continue to be delivered. The strain is at the edges of the scope the category has absorbed, not at the centre of what it was built for.

Close

The channel manager is the clearest partial answer the experiences industry has produced to the distribution-connectivity problem. Partial is the honest word. The category solves two real economic goods and solves them durably. It was not designed to solve real-time availability consistency across deep chains, end-to-end data integrity under heterogeneous vocabularies, distribution economics at structural scale, or neutral settlement. It has been pushed to carry some of those additional problems, with meaningful success at the edges and visible strain at the centre as scope has accumulated.

Whether a layer emerges that sits underneath or alongside the channel-manager category, carrying weight it was not built for, is the structural question the rest of this arc takes up. The next essay turns to the arrival of AI-mediated discovery and agentic distribution, and asks what that wave, now reshaping adjacent travel categories, actually changes about the architecture described so far. Less, in some ways, than the current hype cycle suggests. More, in ways the industry has not sat with, than it is ready for.

The channel manager works. The question is what it cannot carry, and what that implies for an industry whose infrastructure needs are growing faster than any one category was built to absorb. Whatever layer emerges next must build on what the channel manager already solved, not pretend it failed.

ENDNOTES

1. Stern, L. W., El-Ansary, A. I., and Coughlan, A. T., *Marketing Channels*, Pearson / Prentice Hall, 7th edition, 2006. The canonical academic textbook on distribution channels; anchors the precise distinction between routing data and carrying the commercial framework.
2. Coase, R. H., "The Nature of the Firm," *Economica*, 4(16), 1937, pp. 386-405; the modern restatement is Williamson, O. E., *The Economic Institutions of Capitalism: Firms, Markets, Relational Contracting*, Free Press, 1985. Foundational transaction-cost economics; anchors the "shared integration cost" economic good as a make-vs-buy-vs-share decision.
3. Arival, *Guide to Channel Management: Strategy, Connectivity & Technology*, November 2023, on durability of the category and operator-side and reseller-side adoption patterns. The multi-sided platform shape of the channel manager is also consistent with Parker, Van Alstyne and Choudary, *Platform Revolution*, W. W. Norton, 2016.
4. The latency-compounding argument across deep chains is a corollary of the consistency-versus-availability trade-off formalised in Brewer, E. A., "Towards Robust Distributed Systems," PODC keynote, 2000, and Gilbert, S., and Lynch, N., "Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-tolerant Web Services," *ACM SIGACT News*, 33(2), 2002, pp. 51-59. Sub-second consistency across a multi-hop chain is bounded by these properties.
5. The translation work described here is an applied case of boundary-object theory: Star, S. L., and Ruhleder, K., "Steps Toward an Ecology of Infrastructure: Design and Access for Large Information Spaces," *Information Systems Research*, 7(1), 1996, pp. 111-134, on the necessity of compromise mappings when heterogeneous schemas meet at a shared infrastructure layer.
6. See Essay 5 on the Billing and Settlement Plan and the Airlines Reporting Corporation; the underlying documentation is in IATA, BSP

documentation, and Airlines Reporting Corporation, ARC system documentation.

7. The pattern by which a category absorbs scope under path-dependence rather than design is the canonical case treated in David, P. A., "Clio and the Economics of QWERTY," *American Economic Review*, 75(2), 1985, pp. 332-337. Categories that solve a specific original problem accumulate adjacent functions through small, locally rational decisions; the cumulative shape is path-dependent rather than designed.
8. Arival, *Technology & Connectivity*, January 2023, on operator-side and reseller-side schema diversity and the absence of a converged product vocabulary.
9. The proprietary-supply-stack pattern referenced here is documented in the wave of reservation-technology acquisitions and capital-markets activity covered in Essay 2 note 4 (Phocuswire, *The Coming Wave of Consolidation in Tours and Activities, 2024-2025*; Expedia / Tiqets December 2025; American Express Global Business Travel / CWT September 2025; Klook NYSE filing November 2025), together with the broader category-level investments in supplier connectivity reported in Phocuswire trade coverage of the same period.

AI Doesn't Remove Infrastructure, It Reveals It

What looked like connectivity functioning at scale was humans absorbing the gap. That subsidy is being withdrawn.

The industry's AI conversation is almost entirely about the consumer interface. Whether the interface is a generative discovery experience, a personalised search result, an agent booking a holiday on a traveller's behalf, or a chat window answering a question about a tour, the debates in this sector converge on one layer: the thing the customer touches. "AI will replace OTAs." "AI will transform marketing." "AI will disrupt booking." The three claims differ in detail. They share a shape. Each treats the consumer-facing layer as the place where something significant happens, and treats everything beneath that layer as out of frame.

The frame is too narrow. The structural implication of AI is not that the consumer interface is replaced by a smarter one. It is that the consumer interface itself, as a category, commoditises. When the interface layer becomes infinite and cheap, the thing that defines market position stops being the interface and becomes whatever the interface sits on.¹ In every previous technology transition, the move has been the same. Value migrates from the layer that is becoming infinite to the layer that remains scarce.² The layer that was hard to build before the transition becomes easy to build after it, and the constraint relocates.

The pattern, across prior transitions

Search did this to publishing. When indexing and retrieval became a cheap utility in the late 1990s, the scarce resource stopped being the content and started being the ability to match a query to the right page at scale. Publishers who had been the layer discovered they were now above the constraint, and much of the value migrated down to the search layer that had replaced their discovery function.³

Social did a similar thing to attention. When the cost of distributing content collapsed inside the networks that became dominant through the 2000s, the scarce resource became attention allocation rather than content production. Publishers, the same category, found themselves displaced again, this time by the feed. The generating layer commoditised; the attention layer concentrated.⁴

Cloud did this to operations. When provisioning, scaling, and managing infrastructure became a commoditised utility

in the early 2010s, the scarce resource stopped being operational competence and became scale. The old constraint (running a data centre) was not removed. It was absorbed into a different layer, operated by different companies, under different economics.⁵ In each of these transitions, the same move is visible. One layer commoditises; the layer underneath becomes load-bearing in a way it was not before. The constraint is not eliminated. It relocates.

What sits underneath, in experiences

The AI-era version of this rule is already in motion elsewhere in travel. Aviation has GDS clearing with a standardised schedule and ticketing vocabulary. Hotels have a central-reservation-system layer with more standardised interface conventions than experiences has developed, and a reasonably well-understood inventory model.⁶ When AI commoditises the consumer interface in those categories, the layer underneath is reasonably well-formed. It will flex, but it is there. Experiences is the category where this is not true. What sits underneath the interface in this industry is the connectivity stack described in Essays 4 through 6. A booking crosses five to seven organisational boundaries. Each boundary carries its own translation between heterogeneous product vocabularies. The chain is latency-bound, translation-heavy, and missing a neutral settlement layer. The channel-manager category, the clearest partial answer the industry has produced, has been pushed to carry weight it was not built for.

If AI is the mechanism by which the consumer interface is commoditising across travel, then the place where that commoditisation meets a bottleneck in this industry is the connectivity stack.⁷ Not because AI is doing something new to the stack, but because the stack was already under strain, and the strain was previously absorbed by something that will no longer be present in the same form.

What AI does not do

It is worth stating this part plainly before the next.

The boundaries. AI does not remove the six boundaries of Essay 4. An AI system checking availability still has to cross the distributor, connectivity, reservation-system, operator, and voucher boundaries, in some order, with some latency, across the same three or more organisations. Replacing a human user with a machine does not collapse the topology underneath. It changes who is at the edge of it.

The translation layer. AI does not remove it. Schemas do not converge because a faster reader is parsing them. The mapping between a reseller's product vocabulary and an operator's product vocabulary is a structural feature of fragmentation, not a feature of the reader. A faster reader with the same mapping reads the same compromises.

The settlement gap. AI does not remove this either. The absence of a neutral settlement layer, named in Essay 5 as a precondition of the aviation clearinghouse pattern, does not become less of a gap because an agent is placing the booking. Money still moves through the same peer-to-peer arrangements, negotiated out of band between the same

participants. A machine at the edge does not change the plumbing behind it.

Whatever the headline version of the AI debate predicts the consumer interface will look like, none of the infrastructure beneath it is removed by the prediction.

What AI does do

What AI does is sharpen the cost of fragmentation. The connectivity stack has, until recently, been a sufficient answer to the industry's connectivity problem because the edge of that stack was a human. Humans absorb unstructured input, resolve ambiguity across systems, work around schema mismatches, wait longer than they should, and re-read a page when the data arrived slightly wrong. A great deal of what has looked like industry connectivity functioning at scale is this absorption.⁸ The gap between what the stack delivers and what the interaction requires has been held together, quietly, at the customer's own expense.

A machine at the edge does not absorb ambiguity. Where a human reader encountering a product with inconsistent age bands across two schemas will pause, compare, and make a judgment call, a structured-data reader will either pick one and commit, or fail. Schema drift becomes decision drift. Product-vocabulary mismatch becomes unbookable inventory. Latency that was tolerated across a twenty-second checkout window is not tolerated inside a machine-to-machine call that expects to complete in a fraction of a second. The translation layer, the compromises it has made,

and the inconsistencies it has been quietly accommodating, all become visible in a way they were not before.

This is most obvious at the agentic end of the AI spectrum, which is why it deserves to be named directly. An agent booking on a traveller's behalf is not merely a faster customer. It is a structurally different kind of edge. It has none of the patience a human booker has. It cannot resolve ambiguity by asking the customer what they meant. It cannot stall the transaction while a human reconciles a policy mismatch out of band. It reads structured data, makes a decision, and books. Whatever the connectivity stack has been quietly hiding behind human accommodation is precisely what stops the agent from completing the transaction reliably.⁹ The argument does not depend on agents becoming the majority booking path soon. It depends on them becoming commercially material enough that supply which is incoherent, incomplete or unreachable to machines becomes structurally disadvantaged. This is the point at which interface commoditisation meets the infrastructure of this industry. The interface layer is being rebuilt, in parallel, by many actors, across discovery, search, and agentic surfaces. The infrastructure underneath is the layer that was already strained and has now been asked to carry more load at lower tolerance.

Implication

Every piece of infrastructure that sits between the interface layer and the operator's daily manifest becomes a load-bearing wall in a way it was not before. Those walls were not

built with machine-to-machine load in mind. The translation layer in particular was built with human accommodation as a quiet dependency. Remove the accommodation and the compromises show.

The next essay takes up the agentic case in its own right, because the agent-mediated booking is the cleanest single place to watch the structural argument above operate in the specific. The essay after that asks where value actually settles when the interface layer commoditises at scale. This essay has done only the framework work: named the pattern, and named the place in this industry where the pattern meets resistance.

Close

The headline version of the AI debate in experiences is about replacement. It asks which consumer-facing participant is displaced and which one takes its place. The structural version is about revelation. It asks what the industry's infrastructure has been quietly holding together, and what happens when the thing it was holding together for, a human at the edge, is no longer what arrives. AI does not remove the infrastructure this industry depends on. It reveals which parts of it were never strong enough to hold the weight they were being asked to carry. Machine-legible, current, trusted supply becomes more important, not less. The rest of this phase is about what that recognition implies.

ENDNOTES

1. The collapse of marginal cost per token across frontier and open-weight large language models, and the corresponding commoditisation of the consumer-interface layer, is documented in Stanford Institute for Human-Centered Artificial Intelligence, *AI Index Report 2025*, on inference-cost curves and the rate at which interface-equivalent capabilities have moved from scarce to abundant.
2. Christensen, C. M., and Raynor, M. E., *The Innovator's Solution: Creating and Sustaining Successful Growth*, Harvard Business School Press, 2003, on the law of conservation of attractive profits: when an adjacent layer commoditises, value migrates to the layer that remains scarce. The same argument is articulated in operator class in Spolsky, J., "Strategy Letter V: The Economics of Open Source," *Joel on Software*, June 2002. Christensen and Spolsky together provide the academic and the practitioner statement of the structural move named in this paragraph.
3. The search transition is treated historically in Battelle, J., *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*, Portfolio, 2005, and quantitatively in Brynjolfsson, E., Hu, Y., and Smith, M. D., "Consumer Surplus in the Digital Economy: Estimating the Value of Increased Product Variety at Online Booksellers," *Management Science*, 49(11), 2003, pp. 1580–1596. Together they describe how value migrated from the publishing layer into the matching layer when retrieval became cheap.
4. The attention argument originates in Simon, H. A., "Designing Organizations for an Information-Rich World," in M. Greenberger (ed.), *Computers, Communications, and the Public Interest*, Johns Hopkins Press, 1971, pp. 37–72. The dominant-feed-era account is Wu, T., *The Attention Merchants: The Epic Scramble to Get Inside Our Heads*, Knopf, 2016.
5. Carr, N. G., "IT Doesn't Matter," *Harvard Business Review*, 81(5), May 2003, pp. 41–49, and Carr, N. G., *The Big Switch: Rewiring the World*,

from Edison to Google, W. W. Norton, 2008. As a category of capability becomes commoditised utility, competitive position migrates to the layer above or beside it.

6. The hotel-distribution comparison is treated systematically in Buhalis, D., *eTourism: Information Technology for Strategic Tourism Management*, Pearson, 2003 (cited in Essay 1), on the development of central-reservation-system infrastructure, interface conventions, and inventory standards across the hotel sector. The contrast with experiences, where no equivalent layer has formed, is the structural point this paragraph rests on.
7. See Essay 4 on the boundary count and translation burden of a single booking, Essay 5 on the absence of a neutral settlement layer relative to the aviation clearinghouse precedent, and Essay 6 on the structural strain accumulated by the channel-manager category as it absorbs problems beyond its original design envelope.
8. The argument that human users in fragmented information systems quietly perform the reconciliation work the systems do not perform themselves is the central thesis of Suchman, L. A., *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge University Press, 1987. The reconciliation argument is reinforced by Bowker, G. C., and Star, S. L., *Sorting Things Out: Classification and Its Consequences*, MIT Press, 1999, on the labour of bridging classification mismatches at infrastructure boundaries.
9. Arival, *Guide to AI in Experiences*, 2024. Field observations on operator and distributor adoption patterns during the early phase of agentic and AI-mediated discovery surfaces. The framing of AI as a layer that exposes rather than removes infrastructure is consistent with this report's findings on machine-to-machine reliability requirements.

ESSAY EIGHT

The Agent-Era Supply Problem

An agent does not compensate for incoherence. The supply problem the industry already had, revealed.

The previous essay ended on a simple observation. A great deal of what looks like industry connectivity functioning at scale is, on closer inspection, the customer quietly absorbing the gaps between what the infrastructure actually delivers and what the interaction requires. The traveller re-enters a hotel address. Accepts that availability on the booking page was wrong and picks a different time. Clarifies by email that the pick-up point moved. Confirms the voucher is legitimate by calling the supplier. That absorption has never appeared on any balance sheet, but it has been one of the industry's most important operating subsidies.¹

Agentic AI is the first interaction model that refuses to pay this subsidy. An agent does not re-enter the address. It does not accept that the time was wrong. It does not call. It

queries a surface, receives a response, and either acts on the response or does not. Whatever the human was quietly compensating for is now exposed, because the new party at the other end of the transaction is not absorbing anything.

This essay asks the obvious follow-on question. If agentic booking becomes commercially material, what does the experiences industry actually have to offer it? The answer is uncomfortable, and it is not primarily a story about AI. It is a story about supply.

What an agent actually needs

Strip away the interface narrative and an agent needs four things from a supply side in order to transact reliably on behalf of a user.²

Machine-readable product data. Not a marketing page, not a PDF, not an image with embedded text. A structured description of what the product is, what it includes, how long it takes, what restrictions apply, and how it is priced, updated when the supplier changes any of those attributes.

Real-time availability coherence. The answer to ‘is this bookable at 09:00 tomorrow for two adults’ has to be authoritative at the moment the question is asked. Not eventually consistent across a cache that refreshes overnight. Not approximately correct most of the time. Authoritative now, or an explicit signal that it cannot be determined.

Deterministic transaction completion. When the agent commits to a booking, the booking either succeeds and is acknowledged as a unique, unambiguous reservation the

supplier is aware of, or it fails and the agent knows it failed. What cannot happen is the ambiguity the industry has long tolerated, where a confirmation is sent, the supplier's calendar shows something else, and the mismatch is resolved on arrival by the person at the desk.

Commercial reachability. The agent has to be able to reach the supplier at all, by means that do not require a human on either end. A supplier whose only route to market is a listing keyed by hand into a distributor extranet is, for an agent acting on behalf of a user, functionally unreachable.

Four requirements. None of them are exotic. A hotel-chain interface has offered all four for the better part of two decades. A flight-booking interface has offered all four since the GDS layer was externalised to consumer-facing booking.³ The question is whether the experiences industry, as an industry, offers them.

What the experiences industry actually offers

A product page can be visible online, persuasive to a human reader, well-ranked in search results, and still not satisfy any of the four requirements above. Visibility and executability are different problems. The honest answer is that it offers each of the four inconsistently at the top of the supply curve and not at all through most of the long tail.

Product data. Fragmented across catalogues that are, for a substantial share of suppliers, maintained by hand inside distributor extranets rather than pulled from a supplier reservation system.⁴ The same supplier can appear with different inclusions, different cancellation terms, and

different titles on different distribution surfaces, because each surface's catalogue was keyed separately.

Availability. Correct authoritatively only at the supplier's own reservation system. Everywhere else, it is a copy that is more or less fresh, held in channel-manager queues and distributor caches that were designed to refresh fast enough for human browsing, not to satisfy an agent acting on a deterministic query at an arbitrary moment.⁵

Transaction completion. The place the gap is most visible. A booking made by a human through a distributor has a complete path to the supplier's calendar only in the subset of cases where the supplier is integrated in real time. In the rest, the confirmation is sent, a secondary process moves it toward the supplier's calendar later, and the traveller's standing at the supplier's front desk is, in effect, a well-formed query about whether those two systems agreed.

Commercial reachability. A four-way asymmetry. The operators that sit inside direct operator-to-distributor integrations are reachable. The operators that sit inside a channel manager are reachable. The operators whose only connection to the distribution graph is an account someone logs into once a week to update prices are, for an automated counterparty, not reachable at all.

Four structural mismatches

Placing the four requirements next to the four realities yields four structural mismatches, and they are not the sort that a clever interface layer can paper over.

Catalogue-first where the agent needs query-first. Catalogues are assembled ahead of time, in a distributor's staging systems, to support a browsable surface. An agent does not browse. It asks specific questions against specific attributes and expects the answer to reflect the supplier's current state, not last week's staging copy.

Eventually consistent where the agent needs real-time authoritative. The industry's connectivity layer was built to propagate calendar changes fast enough that human-speed browsing rarely sees a stale number. It was not built so that a machine, querying at any moment for a specific combination of date, time, party composition and product, receives an answer guaranteed to agree with the supplier's own system at that instant.

Loose where the agent needs deterministic. The industry has tolerated a confirmation protocol where the booking is considered 'good enough' and residual ambiguities are resolved at the moment of arrival. That tolerance has been sustained by humans on both sides of the interaction who can resolve mismatches in real time. It does not survive the removal of the human.

Network-of-networks where the agent needs a coherent counterparty. The industry's distribution graph is a sparse mesh held together by channel managers, manual extranets, and bilateral agreements of varying machine-readability. For a human shopper, this mesh is invisible because a distributor's front-end hides it behind a search box. For an agent, the mesh is the problem. It has to reach every

supplier it intends to transact with, and it has to reach them in the same shape.

The uncomfortable diagnosis

The temptation, reading this, is to conclude that agentic AI creates a new set of problems for the experiences industry. It does not. The four mismatches described above are not new. They are the conditions the industry has operated under for as long as it has been digital. What agentic AI does is remove the party that has been compensating for them, which is the human at the interface. Essay 7 named this directly. Human tolerance has been the hidden subsidy making the distribution graph appear to work. The agent is the first participant that refuses to pay the subsidy.

Which means the supply problem the industry is about to discover in the agent era is not a new supply problem. It is the one it already had. The tolerance of the human at the interface has been obscuring it, essay by essay of industry commentary, press release by press release, throughout the modern history of the category.

Who feels this first

The first-order effect falls on the long tail.

In the experiences industry, the long tail is not a minor phenomenon. It is the majority of the supply base. The overwhelming share of operators globally are independent, run by small teams, and occupy a position in the distribution graph that depends on intermediation.⁶ Many of them are reachable to distributors only through an extranet that a

human updates intermittently. This is not a failing of operators. The industry has not rewarded operators consistently for producing complete, structured, machine-actionable supply data, and many have had no rational reason to maintain it. For an agent acting on behalf of a user at commercial scale, those operators are, absent a structural change, functionally invisible.

The point deserves to be stated sharply, not as a neutral observation. A supplier unreachable by an agent is unreachable by a growing share of demand. Whether this becomes a long-run structural problem depends on whether the preconditions for agent reach can be met at the long tail, not only at the top. If they can, the industry converges toward a supply surface that is legible to an agent without flattening into a handful of megasuppliers. If they cannot, the only operators reachable by an agent will be those sitting inside channels already designed for machine-to-machine transaction, which is a much smaller population.

What happens by default

Absent a neutral infrastructure layer capable of presenting a coherent, machine-readable supply surface across the long tail, the agent era's default outcome is a concentration effect. The reasoning is structural. An agent, at scale, will transact most reliably with whichever counterparty can satisfy the four requirements end-to-end. Where no neutral layer offers that, the position falls to the participants who have built their own. Those participants are the ones for whom demand and supply sit inside a single vertically

integrated platform, because that is the only configuration in which product data, availability, transaction completion and reachability can be made coherent by fiat.⁷

The infrastructure position defaults, in other words, to whoever can present the cleanest supply surface to the agent, and for structural reasons that is the participant whose commercial interests are not neutral to the outcome. This is a first-principles observation about what happens if nothing changes. It is not a prediction about any specific participant. The point is that the agent era has a default, the default is concentration, and the question of whether the industry ends up there is the question of whether a neutral supply layer exists to offer an alternative.⁸

Close

The experiences industry has, in its modern history, optimised for attention. Interfaces, brand affinity, marketing funnels, conversion. The agent era does not care about any of these. It cares about a structured supply surface that is real-time, deterministic, and commercially reachable across the full length of the supply curve.

That surface does not currently exist at industry-wide standardisation. The tolerance that has masked its absence is being withdrawn. Who ends up holding the infrastructure position is, therefore, not really a question about AI. It is a question about supply, and about which participants are positioned to present supply coherently to a counterparty that will not compensate for incoherence on their behalf.

The next essay takes up the larger pattern this sits inside. When the interface commoditises, value does not disappear. It migrates. The migration has happened before, in search, in social, in cloud infrastructure, in retail distribution.⁹ Where value goes when interfaces die is, in every previous case, the same place. The experiences industry is about to find out whether it has built the thing that place happens to be.

ENDNOTES

1. The argument that the customer in a fragmented information system performs uncompensated reconciliation work that the system does not perform itself is the central thesis of Suchman, L. A., *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge University Press, 1987 (cited in Essay 7), reinforced by Bowker, G. C., and Star, S. L., *Sorting Things Out: Classification and Its Consequences*, MIT Press, 1999 (cited in Essay 7), on the labour of bridging classification mismatches at infrastructure boundaries. The reframing as 'operating subsidy' applies these academic insights to the commercial structure of the experiences industry specifically.
2. The four-requirements analysis of an autonomous agent transacting on behalf of a principal is grounded in the multi-agent systems literature: Wooldridge, M., *An Introduction to MultiAgent Systems*, Wiley, 2nd edition 2009, on what autonomous agents require from their environment in order to act reliably, and Shoham, Y., and Leyton-Brown, K., *Multiagent Systems: Algorithmic, Game-Theoretic, and Logical Foundations*, Cambridge University Press, 2009, on the formal foundations of agent coordination, communication protocols, and the determinism requirements that distinguish a machine counterparty from a human user.
3. The hotel-distribution comparison is treated systematically in Buhalis, D., *eTourism: Information Technology for Strategic Tourism Management*, Pearson, 2003 (cited in Essays 1 and 7), on the development of central-reservation-system infrastructure across the hotel sector. The flight-booking comparison is treated in Essay 5, drawing on Copeland, D. G., and McKenney, J. L., "Airline Reservations Systems: Lessons from History," *MIS Quarterly*, 12(3), 1988, pp. 353-370, on the externalisation of the SABRE programme to travel agents from the mid-1970s and to consumer-facing booking through the 1990s.
4. Arival, *Technology & Connectivity: The state of booking system adoption and satisfaction, distribution connectivity, and tech priorities for tours, activities & attractions*, January 2023 (cited in Essays 3 and

- 4), on the prevalence of distributor-extranet content management and the structural fragmentation of catalogue surfaces across the supplier base.
5. Essay 4 traces the consistency-versus-availability trade-off in the booking chain through Brewer, E. A., "Towards Robust Distributed Systems," PODC keynote, 2000, and Gilbert, S., and Lynch, N., "Brewer's Conjecture and the Feasibility of Consistent, Available, Partition-tolerant Web Services," ACM SIGACT News, 33(2), 2002, pp. 51-59. Essay 6 traces the structural limits of the channel-manager category as the layer the industry currently asks to carry the consistency burden across boundaries.
 6. Arival, 2022 Industry Outlook for Tours, Activities and Attractions Operators, 2022 (cited in Essay 1), on the scale and shape of the global operator base, including the long-tail composition that drives the structural argument made here.
 7. The vertical-integration argument is the central claim of Essay 4, which traces how distributor and reservation-system integration inside a single corporate perimeter collapses cross-organisation interfaces into internal function calls. The recent reservation-technology M&A wave covered in Essay 2 note 4 (Phocuswire, "The Coming Wave of Consolidation in Tours and Activities," 2024-2025; American Express Global Business Travel / CWT, September 2025; Expedia / Tigtets, December 2025; Klook NYSE filing, November 2025) is the empirical record of this configuration extending across the category in real time.
 8. The structural argument that an integrated platform displaces single-side specialists when neutral coordination layers are absent is treated formally in Eisenmann, T., Parker, G., and Van Alstyne, M. W., "Platform Envelopment," Strategic Management Journal, 32(12), 2011, pp. 1270-1285. The paper provides the canonical academic anchor for the concentration outcome described here: when standards and neutral coordination fail, ownership consolidation across the demand-supply pair becomes the equilibrium configuration that satisfies the requirements neutrally.
 9. See Essay 7 on the structural pattern across prior technology transitions (search, social, cloud), grounded in Christensen, C. M., and

Raynor, M. E., *The Innovator's Solution*, Harvard Business School Press, 2003, on the law of conservation of attractive profits, with corresponding precedent citations: Battelle, J., *The Search*, Portfolio, 2005, and Brynjolfsson, E., Hu, Y., and Smith, M. D., *Management Science*, 49(11), 2003, on the search transition; Simon, H. A., 1971, and Wu, T., *The Attention Merchants*, Knopf, 2016, on the attention transition; Carr, N. G., "IT Doesn't Matter," *Harvard Business Review*, 2003, and *The Big Switch*, W. W. Norton, 2008, on the cloud transition.

ESSAY NINE

Where Value Migrates When Interfaces Die

Search, social, cloud, retail. The same structural move each time. Experiences is next.

The previous essay ended on a promise. When the interface commoditises, value does not disappear. It migrates. The migration has happened in search, in social, in cloud infrastructure, in retail distribution, and it has gone to the same kind of place in each. Before arguing what that means for the experiences industry, it is worth laying out the pattern in the open, because the pattern is the argument.

Four cases

Search. By the mid-2000s, the front-end of search had converged toward a generic shape: a single box, a ranked list of results. The interface ceased to be a site of meaningful differentiation. Value migrated to the index underneath, to the crawl coverage, the ranking signal, and the auction

mechanism for monetising intent, none of which a new front-end could cheaply replicate.¹

Social. By the late 2010s, social-platform interfaces had converged toward a common shape: vertical infinite scroll, algorithmic ranking, short-form video as the dominant unit. Value migrated to the attention graph underneath, to the infrastructure for identity resolution, audience modelling, and ad placement that no new interface could assemble at comparable scale. New interfaces kept being built. The position underneath did not move.²

Cloud. By the early 2020s, the consumption surface had largely converged toward a standard shape: container orchestration, infrastructure-as-code tooling, a common vocabulary for object storage, queues, and serverless functions. The interface became portable. Value migrated to what the interface could not make portable, which was the physical compute, networking, and storage at scale underneath. A new provider can replicate the APIs. Replicating the physical footprint is a different kind of project.³

Retail. By the mid-2010s, the e-commerce front-end had converged toward a pattern any competent team could assemble: search, product page, cart, checkout, review. What did not converge was the layer underneath that made the front-end's promises deliverable. Warehousing density, last-mile logistics, demand forecasting, returns processing. Value migrated to that layer. The dominant competitive position in e-commerce was won at the fulfilment layer, not at the

storefront, and that position has been structurally durable in a way no front-end advantage has been.⁴

Four transitions, four different industries, the same structural move each time.

The rule

The rule these cases illustrate is simple, though it is more frequently observed than stated plainly.

An interface sits between demand and supply. When interfaces become cheap to build and user switching costs are low, the margin available at the interface compresses. Whatever sits immediately underneath the interface, and is expensive or slow for a new entrant to assemble, is where the compressed margin reappears. Value moves to the layer that holds demand or supply coherent at a scale new interfaces cannot easily reassemble for themselves.⁵

This is a structural observation rather than a strategic one. Nobody plans for it. It happens because interfaces that need a particular layer in order to function end up paying for that layer, one way or another, and the layer that multiple interfaces need accrues the returns that any individual interface cannot hold. The interface becomes a commodity. The underlying layer becomes an infrastructure position.

Two things are worth noting about the rule as it has played out historically. The first is that the migration is rarely obvious in real time. The players in each of the four transitions above continued to treat the interface as the site of competition well past the point at which the competitive surface had moved underneath. The second is that the

migration is not reversible once the underlying layer reaches sufficient scale. A cloud competitor cannot out-scale the leading hyperscaler by releasing better APIs. A search competitor cannot out-index the dominant index by releasing a better front-end. The position, once held, is durable.

Why this applies to experiences

The experiences industry has, in its modern history, treated the interface as the primary site of competition. Booking flows, conversion funnels, loyalty schemes, brand campaigns, search-engine marketing, product-detail page optimisation. An enormous share of operational and strategic effort across distributors and suppliers has gone into the interface. If the argument of Essay 8 is correct, that interface surface is about to commoditise. Agents do not care about conversion design. Agents do not respond to brand equity. Agents do not absorb confusion on behalf of the user, which means the interfaces that competed on their ability to manage human confusion cannot compete on that ability when the counterparty is a machine. What agents care about is whether the supply underneath can be queried, priced, and transacted against reliably and at scale.⁶

This is the same structural move the four prior transitions described. The interface competes itself toward a generic shape. Whatever sits immediately underneath, and holds supply coherent at scale, accrues the value that the interface can no longer hold. The experiences version of this pattern is not that booking pages disappear. It is that the page becomes less defensible than the supply state beneath it.

Where the pattern says the value goes

In the experiences case, the layer underneath is the one Essay 8 described. It is the layer that can present, across the supply base, a machine-readable product catalogue that is current rather than staged, an availability surface that is authoritative rather than eventually consistent, a transaction protocol that completes deterministically, and a reachability graph that does not depend on someone keying updates into an extranet once a week.

Whoever holds that layer, across enough of the supply base to matter, holds the position value migrates into as the interface commoditises. This is not a strategic claim. It is the same claim the pattern has already made in four other industries.

What is different about experiences is that the layer in question does not currently exist at industry-wide scale. Essay 6 traced why the channel-manager category is structurally partial: it solves for shared integration cost and consolidated reseller access, not for real-time authoritative availability, not for end-to-end data integrity, not for neutral settlement.⁷ The layer the pattern says value migrates into is therefore either still to be built, or is about to be built in a shape that locks the industry into a particular path.

Two paths

There are, structurally, two shapes that layer can take.

The neutral layer. A layer that suppliers and distributors both connect to on commercial terms that are independent of

whether they compete with each other. The layer holds the supply coherence. It does not own the supply and it does not own the demand. Its position in the graph is analogous to the position the Billing and Settlement Plan occupies in aviation settlement, or that cloud compute occupies underneath software-as-a-service, or that card-scheme infrastructure occupies underneath consumer payments. The value accrues to the layer, and the layer is available to every participant on symmetric terms.⁸

The vertically integrated supply surface. A configuration in which coherence is achieved inside a single commercial perimeter rather than across the industry. The layer exists, and is capable of presenting the supply surface an agent needs, but it is owned by a participant whose commercial interest is not neutral to the outcomes of other participants connecting to it. This is structurally analogous to how e-commerce marketplaces resolved for the long tail of independent sellers. The layer exists. Using it is conditional on terms set by a participant with its own interest in where demand flows.⁹

Both paths are consistent with the migration pattern. Both have happened in other industries. Neither is a prediction.

The asymmetry

These paths are not symmetric outcomes for the industry. They produce meaningfully different shapes.

In the neutral-layer outcome, the fragmented supply base retains access to agent-driven demand without being absorbed into any one participant's commercial perimeter.

An independent operator in a small European city stays an independent operator, reachable by every agent, on terms negotiated with a neutral layer rather than with a competitor for the end customer. Distributors stay distributors. The competitive surface between them moves back toward the things distributors can actually compete on, which are curation, trust, bundled experience, and demand generation within specific traveller segments. In the vertically integrated outcome, access to agent-driven demand is conditional on sitting inside whichever platform holds the coherent supply surface. Suppliers whose inventory is not visible inside that platform are unreachable to the agent. The platform's commercial terms become the industry's commercial terms. The long tail of independent operators either consolidates, gets absorbed into private labels, or loses reach. It is possible to believe one outcome is more likely than the other. It is also possible to hold no strong view. What is not really possible, if the pattern from the four prior transitions holds, is to assume the migration does not happen at all. The migration is what always happens. The question is only into what shape of layer it resolves.

Close

The experiences industry is inside a transition of exactly the sort that has played out several times already in adjacent industries. The commoditisation of the interface is underway. The migration of value to the layer underneath is the structural consequence. The only variable is the shape of the layer the value migrates into, and that variable is being

decided now, in the specific technical and commercial choices the industry is making about how supply is going to be presented to its next counterparty.

The choices are being made distributedly, by participants who do not always recognise that the cumulative result is a structural settlement. A reservation system that does not invest in real-time availability is contributing to the settlement. A distributor that builds its own supply integrations rather than connecting to a shared layer is contributing to it. A neutral-layer initiative that fails to reach critical mass is contributing to it by default. None of these decisions appears, locally, to be a choice about the industry's shape. Cumulatively, that is what they are.

The next essay argues that the point at which the cumulative choices crystallise into a settlement that is hard to undo is close at hand. The industry is at its inflection point.

ENDNOTES

1. The search transition is treated in Battelle, J., *The Search: How Google and Its Rivals Rewrote the Rules of Business and Transformed Our Culture*, Portfolio, 2005 (cited in Essay 7), and quantitatively in Brynjolfsson, E., Hu, Y., and Smith, M. D., "Consumer Surplus in the Digital Economy," *Management Science*, 49(11), 2003 (cited in Essays 1 and 7), on the migration of value from publishing into the matching layer when retrieval became cheap.
2. The social transition is treated in Simon, H. A., "Designing Organizations for an Information-Rich World," 1971 (cited in Essay 7), on attention as the scarce resource, and in Wu, T., *The Attention Merchants*, Knopf, 2016 (cited in Essay 7), on the dominant-feed-era migration of value from content production to attention allocation.
3. The cloud transition is treated in Carr, N. G., "IT Doesn't Matter," *Harvard Business Review*, 81(5), 2003, and Carr, N. G., *The Big Switch*, W. W. Norton, 2008 (both cited in Essay 7), on the commoditisation of the cloud-consumption surface and the migration of competitive position to the underlying physical infrastructure.
4. The retail transition is treated in Stone, B., *The Everything Store: Jeff Bezos and the Age of Amazon*, Little, Brown, 2013, on the strategic role of fulfilment-layer investment in producing structurally durable competitive position in e-commerce, and Anderson, C., *The Long Tail*, Hyperion, 2006 (cited in Essay 1), on the underlying long-tail mechanic that fulfilment-layer scale is positioned to capture.
5. The structural rule named here is the law of conservation of attractive profits, treated in Christensen, C. M., and Raynor, M. E., *The Innovator's Solution*, Harvard Business School Press, 2003 (cited in Essay 7), and articulated in operator class in Spolsky, J., "Strategy Letter V," *Joel on Software*, 2002 (cited in Essay 7).
6. The argument that an autonomous-agent counterparty does not absorb the human-accommodation work the consumer-facing interface has been performing is grounded in the four-requirements analysis of Essay

- 8, drawing on Wooldridge, M., *An Introduction to MultiAgent Systems*, Wiley, 2nd edition 2009, and Shoham, Y., and Leyton-Brown, K., *Multiagent Systems*, Cambridge University Press, 2009 (both cited in Essay 8).
7. See Essay 6 for the full argument on the channel-manager category as a partial answer to the industry's distribution-connectivity problem, including the four structural problems it was not designed to solve. The classification of those problems draws on Brewer, E. A., 2000, and Gilbert, S., and Lynch, N., 2002 (cited in Essay 4) for availability consistency; Star, S. L., and Ruhleder, K., 1996 (cited in Essay 3) for end-to-end data integrity across boundaries; and IATA Billing and Settlement Plan and ARC documentation (cited in Essay 5) for neutral settlement.
8. The neutral-layer analogies referenced here draw on three institutional precedents: Essay 5 on the Billing and Settlement Plan and the Airlines Reporting Corporation as neutral aviation settlement infrastructure; Essay 7 on the cloud-compute layer underneath software-as-a-service; and, on the card-scheme infrastructure underneath consumer payments, Evans, D. S., and Schmalensee, R., *Paying with Plastic: The Digital Revolution in Buying and Borrowing*, MIT Press, 2nd edition 2005, on the structural position of card-scheme rails as neutral infrastructure available to every participant on symmetric terms.
9. The vertical-integration path described here is treated formally in Eisenmann, T., Parker, G., and Van Alstyne, M. W., "Platform Envelopment," *Strategic Management Journal*, 32(12), 2011, pp. 1270-1285 (cited in Essay 8). The structural mechanism: in the absence of a neutral coordination layer, ownership consolidation across the demand-supply pair becomes the equilibrium configuration that satisfies the requirements neutrally inside a single commercial perimeter.

ESSAY TEN

The Experiences Industry Is Hitting Its Inflection Point

Five structural conditions that held the industry's shape are weakening simultaneously.

Industries do not turn on a single day. What looks in retrospect like a moment of change is almost always, in the evidence, a convergence. Several structural conditions that had been stable for years weaken at the same time. Several others, previously absent, begin to accumulate. Somewhere in that overlap the industry crosses a threshold, after which the prior shape cannot be held in place and a new shape becomes difficult to avoid. That threshold is what an inflection point actually is.

The argument of this essay is that the experiences industry is inside one of those crossings now. The argument is not that something dramatic happens in a particular quarter, nor that any specific event is decisive. The argument

is that five structural conditions that held the industry's shape in place throughout the digital distribution era are weakening simultaneously, that the market signals visible in public-facing sources confirm the transition is underway, and that the choices made at the supply-surface layer during this window have outsized and durable consequences, for the reasons Essay 9 described.

The conditions that held the prior shape

The experiences industry has had a remarkably stable shape since online distribution scaled. It is worth naming the structural conditions that held that shape, because each of them is now under pressure, and the simultaneity is the point.

Human-tolerant interfaces. The entire distribution graph assumed a human at the edge who would absorb incoherence (Essay 7). A misprinted pickup location, a stale availability cache, an ambiguous voucher number. All of these were quietly reconciled by the traveller standing at the desk. The interface did not need to be complete or deterministic because the human made it complete on arrival.

Fragmented demand matched to fragmented supply. The industry's fragmentation has not felt dysfunctional because both sides of it are fragmented. No single distributor held enough demand to force structural coherence on supply; no single operator category held enough supply to force structural coherence on demand. Fragmentation was stable because the asymmetry between the two sides was small.

Interface as the primary competitive surface. Distributors and suppliers alike have, in the modern history of the industry, treated the interface as where competition happens. Booking flows, conversion, brand. This was rational as long as humans were the counterparty and the interface was the surface they interacted with.

Channel managers as a good-enough routing layer. The category that emerged as a distinct piece of infrastructure solved for shared integration cost and consolidated reseller access (Essay 6). It did not solve for authoritative real-time availability, end-to-end data integrity, or neutral settlement, but the margin of error those gaps produced was absorbed by the human-tolerant interface upstream. The routing layer was sufficient because the failures downstream were human-resolvable.

Connectivity treated as a cost centre. Because none of the above conditions pressured the connectivity layer, participants at all levels of the industry treated it as overhead. Something to minimise. Something to outsource. Something that did not repay strategic investment.

Five conditions. One equilibrium.

Why each condition is weakening now

The conditions are weakening not in sequence, but together.

Human-tolerance is being withdrawn. The subsidy the industry has run on is priced at zero in the agent era. An agent does not reconcile. It queries, receives, and either acts or fails.¹

The buyer side is consolidating faster than the supplier side. The demand-side concentration visible in the major online travel aggregators has moved the industry away from the symmetric fragmentation that made the prior equilibrium stable.² A smaller number of buyers holds a larger share of the flow. That asymmetry alone, absent everything else, would be enough to pressure the distribution graph.

The interface is commoditising under agent pressure. For the reasons Essay 9 set out, accumulated investment in conversion design and brand does not transfer to a counterparty that does not convert and does not respond to brand.

The channel-manager ceiling is becoming material. Channel managers were sufficient when routing could be good-enough and failures could be resolved downstream. With agents at the other end, routing has to be authoritative and failures cannot be resolved downstream. The routing layer is now being asked to do something it was not designed to do.³

Connectivity is no longer a cost centre. If Essay 9 is right, it is the layer value migrates into when the interface commoditises.⁴ The strategic cost of treating it as overhead compounds in every quarter the migration continues.

Five conditions. Five simultaneous pressures. Not from any single direction.

What the signals confirm

A structural argument about weakening conditions is worth only as much as the signals in the market confirm it. Five are visible in public-facing sources.

Standardisation activity at the connectivity layer. Arival's Technology and Connectivity 2023 report documents the emergence of connectivity standardisation efforts within tours and activities, a category that had no equivalent work earlier in the digital era. The presence of an industry-level standards conversation is, in itself, a signal that the connectivity layer has become strategic.⁵

Shift in how operators describe their technology priorities. The same Arival research tracks a notable move of connectivity and data quality up the operator priority list. Technology discussion among operators has shifted from point-of-sale and website, which dominated prior-cycle surveys, toward connectivity and supply-feed quality. This is the operator-side expression of the pressure described above.

A dedicated industry literature on agent-mediated demand. Arival's Guide to AI in Experiences 2024 is the first industry-wide synthesis of how agent-mediated interaction changes what the supply side is required to present.⁶ Earlier no such literature existed in the category. The fact that the industry's primary research body has produced a guide is itself the signal; it is not a commentary on any single recommendation within it.

Observable moves at the buyer side toward structured supply. The largest distributors in the category have publicly communicated, over recent quarters, an investment focus on structured supply feeds, machine-readable catalogues, and direct real-time connectivity with operators, rather than on front-end optimisation alone. No specific company is named here; the observation is about a pattern visible across the set of major distributors, which is what makes it a signal about the industry rather than a claim about any one participant.⁷

Capital flows into infrastructure rather than interface. The direction of investment in the sector has visibly shifted from interface-layer businesses (new booking apps, new marketplace front-ends) toward infrastructure-layer businesses (connectivity, data, settlement). This observation is qualitative, but it is consistent with what the prior essays predict: value migrates to the layer underneath when the interface commoditises, and capital generally follows the migration.⁸

None of the five signals is decisive on its own. Together, they triangulate the same structural moment.

Why this is not the mobile cycle or the consolidation cycle

The experiences industry has had prior moments that felt, at the time, like inflections. An earlier mobile cycle, during which the interface shifted from desktop to smartphone. A subsequent OTA consolidation cycle, during which a small number of distributors acquired or squeezed out smaller ones.⁹ In both, industry commentary described a

transformation. In both, the underlying shape of the industry survived more or less intact.

The reason is structural. Both prior cycles operated at the interface layer. The mobile cycle changed the device through which the interface was consumed; it did not change what the interface had to do. The consolidation cycle changed the ownership structure of distributors; it did not change the underlying distribution graph. The five prior-equilibrium conditions described above held through both.

The current moment is categorically different because it operates at the supply-surface layer. The pressure is not on how the interface is shaped, or who owns the interface, but on whether the interface can continue to exist as the primary competitive surface at all. The five prior-equilibrium conditions are under pressure simultaneously, and four of them sit underneath the interface, at the layer the prior cycles did not disturb. A reader who assumes this cycle resembles the prior ones is making a category error, not a disagreement of degree.

The window, and why it is not symmetric

The argument of Essay 9 was that infrastructure positions, once held, are durable.¹⁰ The value that migrated to the dominant search index, to the dominant fulfilment layer, to the leading hyperscalers, did not redistribute later because the layer underneath had been built at a scale a new entrant could not assemble. The same durability applies to whatever supply-surface layer ends up holding the coherence function in experiences. The choices made in the next several years

about how supply is presented to agent-mediated demand are choices that compound. Early decisions about architecture, commercial terms, and participation requirements establish the shape, and the shape persists.

This is a structural observation about how infrastructure layers form, not a tactical prescription. The asymmetry of the window is what matters: a participant arriving late to the layer does not arrive at the same layer. The coherence function is either a neutral shared resource, a set of vertically integrated private resources, or some mixture of the two, and which of those shapes solidifies is being settled by the moves visible in the market now.

What this means for different participants

Each class of participant is facing a different version of the same structural question.

Large distributors. Already building or buying supply-surface coherence inside their own commercial perimeter. Their choice, expressed through public roadmaps and investment patterns, is consistent with the vertically integrated path from Essay 9. Whether they also participate in a neutral layer is an open question.

Smaller distributors. Facing a different choice. They cannot replicate the integrated supply surface at scale, and they may find the agent era rewards either specialisation (becoming a curation and trust layer on top of someone else's supply surface) or participation in a neutral infrastructure layer that allows them to remain competitively relevant without owning the supply coherence themselves.

Suppliers at the top of the supply curve. Already have connectivity in a shape that is agent-ready or close to it. Their question is on what commercial terms their supply surface is visible to which agents.

Suppliers at the long tail. Facing the sharpest structural question, for the reasons Essay 8 described. Reachability by agents is not a given, and the route to reachability runs either through an intermediary layer or through a vertically integrated platform whose commercial terms are not neutral.

Infrastructure participants. Operating inside the window during which the infrastructure position in this category is, for the first time, forming at industry-wide scale. The opportunity is large. The window is not indefinite, for the durability reasons already described.

Close

What changes during an inflection point is not the industry's visible surface. The booking flows remain booking flows. The marketing campaigns still run. The interface-layer conversation continues as though it were the central conversation. What changes is underneath. The five conditions that held the prior shape give way at once. The connectivity layer moves from overhead to strategic. The supply surface becomes the site of competition. The settlement is not yet decided; it is being shaped, distributedly, by the technical and commercial choices participants are making now. And the choices made at that layer solidify into an infrastructure position whose durability

means the industry's shape for the decade after the window is, in large part, set inside the window itself.

The next essay attempts the straight-faced exercise of describing what the experiences industry looks like in 2030, on each of the two paths that the inflection could resolve into. The exercise is useful not because the future is knowable but because the structural properties of each outcome are knowable, and reading them side by side makes the stakes of the current window visible in a way that a more cautious essay could not.

ENDNOTES

1. The customer-side subsidy that the agent era prices at zero is the central inheritance from Essay 7, which grounds the argument in Suchman, L. A., *Plans and Situated Actions: The Problem of Human-Machine Communication*, Cambridge University Press, 1987, and Bowker, G. C., and Star, S. L., *Sorting Things Out: Classification and Its Consequences*, MIT Press, 1999. Essay 8 extends the argument to the four-requirements analysis of an autonomous-agent counterparty, drawing on Wooldridge, M., *An Introduction to MultiAgent Systems*, Wiley, 2nd edition 2009, and Shoham, Y., and Leyton-Brown, K., *Multiagent Systems*, Cambridge University Press, 2009.
2. Arival, *Technology & Connectivity: The state of booking system adoption and satisfaction, distribution connectivity, and tech priorities for tours, activities & attractions*, January 2023 (cited in Essays 3, 4, 7, 8 and 9), on operator-side connectivity and the buyer-supplier asymmetry. The empirical record of demand-side concentration is treated in Essay 2 note 4 (Phocuswire, “The Coming Wave of Consolidation in Tours and Activities,” 2024–2025; American Express Global Business Travel / CWT, September 2025; Expedia / Tiqets, December 2025; Klook NYSE filing, November 2025).
3. See Essay 6 on the structural limits of the channel-manager category as a routing layer, including why the four problems it was not designed to solve become material under machine-to-machine load.
4. See Essay 9 on the conservation-of-attractive-profits pattern: Christensen, C. M., and Raynor, M. E., *The Innovator’s Solution*, Harvard Business School Press, 2003, and Spolsky, J., “Strategy Letter V,” Joel on Software, 2002. Value migrates to the layer that holds demand or supply coherent at a scale new interfaces cannot easily reassemble for themselves.
5. Arival, *Technology & Connectivity*, January 2023 (cited above), specifically on the emergence of an industry-level standardisation

conversation and the operator-side priority shift toward connectivity and supply-feed quality.

6. Arival, *Guide to AI in Experiences*, 2024 (cited in Essays 7, 8 and 11), on the synthesis of how agent-mediated interaction changes the supply-side requirements; the existence of a category-wide guide is itself the signal cited here.
7. The buyer-side investment pattern referenced here, in which major distributors have publicly communicated investment focus on structured supply feeds and direct real-time connectivity rather than on front-end optimisation, is consistent with the reservation-technology consolidation wave documented at Essay 2 note 4 (Phocuswire, “The Coming Wave of Consolidation in Tours and Activities,” 2024–2025) and at Essay 6 note 9. The observation is described as a pattern across the category rather than as a claim about any single participant.
8. The capital-flow shift toward infrastructure-layer businesses is the empirical pattern documented in the trade-press coverage cross-referenced at Essay 2 note 4 and Essay 6 note 9, including Phocuswire’s coverage of the 2024–2025 reservation-technology consolidation wave.
9. The mobile and consolidation cycles referenced here are treated in the longer historical record of tourism-distribution technology in Buhalis, D., *eTourism: Information Technology for Strategic Tourism Management*, Pearson, 2003 (cited in Essays 1, 7, 8 and 9), on the device transition and channel concentration; the more recent OTA consolidation cycle is documented across Phocuswire trade coverage referenced in Essay 2 note 4.
10. The durability argument referenced here is the central claim of Essay 9, grounded in Christensen, C. M., and Raynor, M. E., *The Innovator’s Solution*, Harvard Business School Press, 2003, and treated formally in Eisenmann, T., Parker, G., and Van Alstyne, M. W., “Platform Envelopment,” *Strategic Management Journal*, 32(12), 2011, pp. 1270–1285 (cited in Essays 8 and 9), on why infrastructure positions, once held at scale, do not redistribute.

ESSAY ELEVEN

What the Experiences Industry Looks Like in 2030

Two paths, read side by side. Structural reasoning forward, not forecasting.

Four years is not very long. It is also, in this particular industry at this particular moment, enough for the inflection described in Essay 10 to resolve into a structurally different equilibrium. This essay takes a straightforward look at what that equilibrium looks like in 2030 on each of the two paths Essay 9 laid out, reads them side by side, and then describes the most realistic middle case.

The exercise is not prediction. It is structural reasoning forward from premises the earlier essays established. Prediction would require knowing a great deal that is not known. Structural reasoning requires only knowing the shape of the forces and allowing them to compound.

Ground rules

Certain things the prior arc treats as structurally determined are held constant across both scenarios.

Agent-mediated interaction reaches commercial scale. This does not mean every booking is made by an agent. It means the share is large enough that being unreachable to agents materially reduces a supplier's access to demand. The Arival Guide to AI in Experiences 2024 treats this trajectory as directionally established, though it does not project a 2030 share.¹

The interface commoditises under agent pressure. Value migrates to the supply-surface layer.² This is the migration from Essay 9; it is not conditional in either scenario. What is conditional is the shape of the layer the value migrates into.

Category demand continues to grow. The experiences category sits at roughly 270 billion dollars in annual gross booking value at the time of writing, with online penetration still in the low thirties by most industry estimates, and is the largest remaining undigitised segment of travel.³ Digital share continues to rise through the window. The absolute size of the category in 2030 is larger than it is now.

What varies between the two scenarios is which of the two paths from Essay 9 resolves. Scenario B is the default-gravity path: vertically integrated supply surfaces held by participants with existing demand. Scenario A is the path that requires active construction: a neutral infrastructure layer.

Scenario B: the default-gravity path in 2030

In this outcome the supply-surface layer is not a neutral shared resource. It is several private ones, each owned by a participant whose commercial interests are not neutral to the outcomes of other participants connecting to it.

Connectivity architecture. A small number of large distributors have built, acquired, or consolidated the connectivity and supply-coherence layer into their own commercial perimeters. Product data, availability, transaction determinism, and operator reach are maintained to an agent-ready standard inside each perimeter, and not maintained to that standard anywhere else.

Operator position. Reach to agent-driven demand is conditional on sitting inside one of the large commercial perimeters. Operators at the top of the supply curve, those with resources and technical capability, participate on the commercial terms the perimeter owners set. Operators in the long tail face a choice that is not really a choice: accept the terms offered, consolidate, convert to private-label supply, or accept reduced reach.

Distributor position. The distribution layer has narrowed. A handful of large platforms hold structural positions. Smaller distributors either specialise into curation atop someone else's supply surface, acquire into one of the large platforms, or lose relevance. The middle of the distribution market hollows.

Consumer-facing surface. Travellers interact with the category through agent surfaces that call a small number of

supply endpoints, each of which returns coherent results from inside one perimeter but does not guarantee coherent results across perimeters. Inside a perimeter, the experience is good: bookings complete reliably, supply is current, the agent transacts without friction. The cost is paid at the seams between perimeters. Cross-perimeter comparison is worse than current open-web comparison, not better, because the incentive to be cross-comparable is weaker for a perimeter owner whose advantage depends on being the inside of the perimeter.

Economic shape. A substantial share of the category's economic surplus migrates to the infrastructure-holding platforms, because as Essay 9 argued, the layer that value migrates into is durable and rent-extractive. Operators retain gross revenue; margin compresses because take rates are set by a smaller number of counterparties with stronger pricing positions. Infrastructure participants that do not own demand either integrate upward into one of the large perimeters or become thin commodity suppliers to them.

Historical analogue. This is the structural shape that e-commerce reached for the long tail of independent retail sellers in the preceding cycle.⁴

Scenario A: the constructed-neutral path in 2030

In this outcome a neutral infrastructure layer has formed during the window, and holds the supply-surface coherence function on commercial terms that are symmetric across distributors and suppliers.

Connectivity architecture. The connectivity layer is held by a neutral actor or a small set of neutral actors. Product data is maintained to an agent-ready standard across the supply base because the layer's commercial model aligns its incentives with coherence at scale rather than with any participant's demand capture. The architecture is structurally analogous to the Billing and Settlement Plan in aviation settlement or to card-scheme infrastructure in consumer payments.⁵ Neither analogy is non-commercial: both require governance, funding, and adoption mechanisms substantial enough to sustain the layer across decades, and the experiences-industry version would face the same requirements.

Operator position. Reach to agent-driven demand runs through the neutral layer. The long-tail operator in a small European city is reachable on the same structural terms as the top-ten-city operator, because the layer is designed to present coherent supply across the full length of the supply curve rather than to privilege the top.

Distributor position. The distribution layer is broader, not narrower. Distributors compete on the things the interface era taught them to compete on (curation, trust, bundled experience, traveller-segment depth) without needing to own the supply coherence function. The middle of the distribution market thickens.

Consumer-facing surface. Travellers interact with the category through agent surfaces that query a layer that has structural coverage across suppliers. Cross-surface comparison is better than current open-web comparison,

because the coherence function is held by a layer whose commercial interest is being queried by every participant.

Economic shape. Economic surplus distributes more evenly across the value chain. The infrastructure layer captures a durable but bounded share (the price of holding coherence at scale). Distributors retain margin on the curation-and-demand-generation surface. Operators keep a larger share of gross booking value because the take-rate is set in a more competitive market for agent reach.

Historical analogue. This is closest in structural shape to aviation distribution after the establishment of the Billing and Settlement Plan and IATA accreditation.⁶ It is not identical. The operators are more fragmented than airlines. The neutral layer carries more coordination load than the Billing and Settlement Plan alone.

Reading the two outcomes side by side

The two scenarios differ structurally on five axes. They are best read in parallel.

Who owns supply coherence. Scenario B: several private participants, each inside their own perimeter, non-neutral to each other and to everyone outside. Scenario A: a neutral layer, commercially independent of demand capture.

Commercial terms of agent reach. Scenario B: set by platform owners whose interest is not neutral. Scenario A: set in a market where the coherence layer is a counterparty to everyone rather than a competitor to some of them.

Position of the long tail. Scenario B: attritional pressure toward consolidation, acquisition, or private-label

absorption. Scenario A: structural retention of access on symmetric terms.

Distributor competition. Scenario B: compressed to the inside of each perimeter, with the middle of the distribution market hollowed. Scenario A: broader, played on curation and bundled experience atop a shared supply surface.

Distribution of economic surplus. Scenario B: concentrated toward the platforms holding the infrastructure position. Scenario A: distributed more evenly, with a durable but bounded share held by the neutral layer.

These are structural properties. They do not depend on which specific companies end up where. They depend only on which of the two shapes the layer takes.

The hybrid middle case

The most likely real-world outcome is neither pure scenario. Both paths partially resolve. Large platforms build supply surfaces inside their own perimeters for the portion of supply they can reach on their own commercial terms. A neutral layer forms for the portion they cannot. The interesting structural question is not which label wins but which proportion resolves which way.

Three variables matter in the hybrid case.

The reach of vertically integrated platforms into the long tail. If the integrated platforms reach most of it on acceptable terms to the operators, Scenario B is structurally close to the outcome even when a neutral layer exists alongside. If they reach only the easier portion, the neutral

layer covers the rest, and the overall shape looks more like Scenario A.

The timing of neutral-layer formation. Infrastructure positions, once held, are durable.⁷ A neutral layer that forms before the vertically integrated platforms have built supply-surface coverage at scale ends up as the structural default. One that forms after does not displace what is already in place. It serves the residual.

The governance of the neutral layer. A layer that begins neutral but accumulates dominant equity from a single strategic owner, or that is captured by a subset of its participants, becomes structurally closer to the vertically integrated outcome regardless of its early positioning. Governance design is not a cosmetic feature.

What is true in 2030 regardless of path

Some things hold across both scenarios and across the hybrid middle case. These are the features of 2030 that are not conditional.

Agents are a material share of bookings.

The interface is not the primary competitive surface.

Supply coherence is strategic.

Connectivity investment has higher marginal ROI than interface investment.

The industry's competitive vocabulary has moved, across operators and distributors alike, from conversion to reachability. The language of funnel optimisation, which has dominated distribution strategy in the modern history of the

industry, has receded into a subset of the work rather than the centre of it.⁸

These are the features worth noticing, because they describe the shape of the industry that any reader of this essay will be operating inside, regardless of which version of the supply-layer outcome resolves.

Close

The question the current window puts to the industry is not which of the two paths is more likely. The gravity of Scenario B is visible. The inflection-point window and the agency of participants during the window are the variables that determine how much of Scenario A's structural features end up in the outcome.

The real question is a different one. It is the question an industry puts to itself when it can see the two scenarios clearly: which of these is the industry its participants would actually choose to operate inside, given that the choice is available to be influenced only during a short window and has compounding consequences afterwards. That is the question the next essay addresses, because it is also the question the industry has been, structurally, not asking itself.

ENDNOTES

1. Arival, Guide to AI in Experiences, 2024 (cited in Essays 7, 8 and 10). The agent-era trajectory premise of the 2030 scenarios draws on field observations in this report; the report does not project a 2030 share, and no specific share is assumed in either scenario here.
2. See Essay 9 on the conservation-of-attractive-profits pattern that drives interface commoditisation and value migration to the layer underneath: Christensen, C. M., and Raynor, M. E., *The Innovator's Solution*, Harvard Business School Press, 2003, and Spolsky, J., "Strategy Letter V," *Joel on Software*, 2002. The migration is treated as structurally determined; only the shape of the destination layer is conditional.
3. The 270-billion-dollar gross-booking-value figure and the low-thirties online-penetration estimate are documented in Arival, *Global Operator Landscape*, 4th edition, January 2026 (cited in Essays 1 and 2), with longer-run digital-share trajectory drawn from Arival, *2022 Industry Outlook for Tours, Activities and Attractions Operators*, 2022 (cited in Essay 1). The "largest remaining undigitised segment of travel" framing follows from comparing these category figures with Buhalis, D., *eTourism: Information Technology for Strategic Tourism Management*, Pearson, 2003 (cited in Essays 1, 7, 8, 9 and 10), on the comparatively higher digital penetration in adjacent travel categories.
4. The structural shape of e-commerce concentration around the long tail is treated in Stone, B., *The Everything Store: Jeff Bezos and the Age of Amazon*, Little, Brown, 2013 (cited in Essay 9), and Anderson, C., *The Long Tail*, Hyperion, 2006 (cited in Essays 1 and 9), on the fulfilment-layer mechanism by which independent sellers either consolidated, became private-label suppliers, or lost reach as the dominant platform's supply-surface coherence advantage compounded.
5. The neutral-layer analogies referenced here are treated in Essay 5 on the Billing and Settlement Plan and the Airlines Reporting Corporation as neutral aviation settlement infrastructure, and in Essay 9 note 8 on card-scheme infrastructure underneath consumer payments, drawing

- on Evans, D. S., and Schmalensee, R., *Paying with Plastic: The Digital Revolution in Buying and Borrowing*, MIT Press, 2nd edition 2005, on the structural position of card-scheme rails as neutral infrastructure available to every participant on symmetric terms.
6. The aviation distribution analogue is treated in Essay 5, drawing on Doganis, R., *The Airline Business*, Routledge, 2nd edition 2006, on the institutional scaffold of aviation including IATA standards, route authority, alliance interlining, and the regulatory framework within which the GDS layer operates, and on Copeland, D. G., and McKenney, J. L., "Airline Reservations Systems: Lessons from History," *MIS Quarterly*, 12(3), 1988, pp. 353–370, on the SABRE programme and the externalisation of airline reservation systems to travel agents.
 7. The durability argument is the central claim of Essay 9, grounded in Christensen, C. M., and Raynor, M. E., *The Innovator's Solution*, Harvard Business School Press, 2003, and treated formally in Eisenmann, T., Parker, G., and Van Alstyne, M. W., "Platform Envelopment," *Strategic Management Journal*, 32(12), 2011, pp. 1270–1285 (cited in Essays 8, 9 and 10), on why infrastructure positions, once held at scale, do not redistribute. The implication for the hybrid middle case is that timing of formation has compounding consequences.
 8. The invariants enumerated here are consequences of the prior arc rather than free-standing claims. Agents-as-material-share inherits from Arival, *Guide to AI in Experiences*, 2024 (note 1 above) and Essay 8. Interface-not-the-primary-competitive-surface inherits from Essay 9. Supply-coherence-strategic inherits from Essays 6, 8, 9 and 10. The shift in competitive vocabulary from conversion to reachability is the operator-side and distributor-side expression of the same structural move.

The Question the Industry Is Not Asking Itself

The question is being answered either way. In the industry's voice, or in its silence.

The eleven prior essays tried to lay out, as carefully as the material allowed, the structural forces that describe the experiences industry at this moment. The persistent under-digitisation of a category at the scale of two hundred and seventy billion dollars in annual gross booking value, fragmentation as a structural state rather than a transitional one, connectivity topology, the anatomy of a booking, the GDS analogy and what it does not transfer, the partial answer the channel-manager category provides, the way AI reveals infrastructure rather than removing it, the agent-era supply problem, the migration pattern visible across prior industry transitions, the inflection window the industry is inside now, and the two scenarios the window resolves into.

Across all of it, the industry's public conversation has been about a different set of questions. Conversion optimisation. Interface design. OTA consolidation and market share. Marketing performance. Growth rates. Traveller sentiment. The interface layer. These are the questions being asked, inside operator rooms, inside distributor strategy decks, inside conference programmes, inside trade press.

There is a question the industry is not asking itself. The silence is not an oversight. It is structural, for reasons this essay will name. What this essay will not do is answer the question, because the answer is what the industry's participants will determine, either by engaging with it during the window or by accepting the default that forms in their absence.

The question

When the value migrates to the supply-surface layer, who holds the infrastructure position, on what commercial terms, and under what governance?¹

This is one question that separates cleanly into four, each of which has to be answered for the question as a whole to be answered.

Coherence. Who owns the layer that holds machine-readable product, real-time availability, deterministic transaction, and coverage across the supply base coherent at industry-wide scale.

Commercial terms. On what terms do suppliers reach agent-driven demand through that layer, and on what terms do distributors participate. Specifically, whether the terms are symmetric across participants or set by a participant with an interest in their own demand capture.

Long-tail access. Whether the independent operator outside the top of the supply curve retains structural access to agents, or whether reach becomes conditional on consolidation into larger commercial perimeters.

Governance. Whether the layer's governance is designed to keep it neutral to its participants, or whether the structure allows drift into capture by a subset of them over time.

These four are the substance. The rest of this essay is about why they have not been asked as a coordinated question in the industry's public conversation, and what happens if they continue not to be.

Why the question has gone unasked

Four structural reasons, each of which reinforces the others.

Connectivity has been treated as a cost centre. Because strategic attention follows profit-and-loss visibility, and connectivity's visibility has been on the cost side rather than the revenue side, senior operating and investment attention in the industry has not been drawn to it in the way it has been drawn to interface, marketing, and distribution economics. The participants asking industry-wide questions have therefore not been treating connectivity as a site where industry-wide questions arise.

The participants best positioned to answer the question publicly have no incentive to ask it in public. The large

distributors building supply-surface coverage inside their own commercial perimeters are structurally positioned to answer the coherence question, because they are the ones doing it. Their answer is visible in their roadmaps. Their incentive to frame it as an industry-wide question is weak, because the answer is self-serving and the absence of a forum for coordinated discussion is not adverse to their interests.

The institutional apparatus through which the industry asks coordinated questions has focused on narrower ones. Arival, which is the category's primary research and convening body, has done excellent work on operator practice, traveller behaviour, distribution economics, and the state of connectivity at the point of adoption. It has not, to date, convened a coherent industry-wide conversation about supply-surface architecture at the infrastructure level, and to be fair to it, the question has not been one the category has been actively asking. Standards-setting efforts that have emerged within the category represent the seed of such a conversation, but the scope of those efforts is narrower than the four-part question above, and the participation in them is a subset of the industry rather than the whole of it.

Inflection-window dynamics pull cognitive attention away from infrastructure questions at exactly the moment those questions matter most. Interface commoditisation, the arrival of agents at commercial scale, and the visible pressure on operator and distributor economics generate a large volume of short-horizon operational decisions. These decisions are real and demand attention. They also crowd

out the longer-horizon infrastructure question that the same forces make most consequential.

None of these reasons is a failure of any particular participant. They are the structural conditions under which the question has not been asked. Naming them is not blame. It is the diagnosis the question needs before it can be asked properly.

Silence is not neutral

The particular thing about an inflection window is that the absence of a question is not a neutral state. It is an active concession to whichever outcome forms under default gravity. Essay 9 described why. Infrastructure positions, once held, are durable; they form fastest in the absence of coordinated opposition; and they are set by the participants moving first rather than by the industry collectively deciding what shape the layer should take. Essay 10 described the timing: the window inside which the shape is being determined is short, and the choices made inside it compound.²

Taken together, this means the question being unasked is itself the mechanism by which the default answer resolves. A silent industry is not an industry that is reserving judgement. It is an industry whose participants are, collectively, not exercising the agency available to them during the only window in which agency is available. The answer that resolves in that silence is the answer the industry would have had regardless of anyone's preferences. Which for structural reasons is the answer that compounds to the vertically integrated side of Essay 11's hybrid case.

What asking the question actually looks like

Asking the question is not a commitment to any particular answer. It is a diagnostic exercise, conducted first by each participant in isolation and then, if the outputs converge, as a coordinated industry conversation.

The individual diagnostic is straightforward. Read Essay 11's two scenarios. Identify where you sit in each. Identify which of the two you would prefer to operate inside, setting aside, for the purpose of the exercise, any specific views about how to get there. Operators, distributors, infrastructure participants, and trade bodies can all do this in an afternoon.³

The coordinated diagnostic follows from the individual one. If preferences across the industry's participants converge on one scenario, the question of what to do about the coherence layer becomes a coordinated industry question rather than an individual participant's strategic bet. If preferences diverge, the industry at least knows where it disagrees, which is a materially better position than its current one of not having surfaced the question at all.

The reason this matters is that neither pure scenario resolves without coordinated work. Scenario A, the neutral infrastructure layer, does not emerge by accident. It is built, and the economics of building it at industry-wide scale require participation from operators, distributors, and infrastructure participants in a form that is not currently organised. Scenario B, the vertically integrated path, resolves by default gravity and does not require coordination. This asymmetry is the entire point of treating

the question as one the industry has to ask itself rather than wait to have answered for it.

What happens if the question continues to go unasked

If the question continues not to be asked through the close of the window, the infrastructure position is held by whoever builds coherent supply-surface coverage fastest. For structural reasons this is the participants with existing demand at scale, because supply coherence is easiest to achieve inside a perimeter that already owns the demand side of the transaction. The 2030 outcome lands on the vertically integrated side of Essay 11's hybrid case, not because any operator preferred it, not because any independent distributor chose it, not because any trade body advocated it, but because no one convened the conversation about what the alternative would actually require and how it would be organised.⁴ The default is the default precisely because it requires nothing to form. The alternative is the alternative precisely because it requires something to form.

This is not a partisan statement about which of the two outcomes is better. Readers of this arc who prefer Scenario B on commercial or philosophical grounds have every right to that preference and should argue for it in the public conversation. Readers who prefer Scenario A, likewise. The argument of this essay is not about which answer is right. It is about the fact that the question itself has not been put, and that the absence of the question is producing an answer by default rather than by deliberation.

Close

This essay does not answer the question. It does not answer it because the answer is not the author's to give. The question is about industry-wide infrastructure, and an industry-wide question has to be answered by the industry, through the forums it has for coordinated conversations, or through new forums if the existing ones are too narrowly scoped for the question as stated.

What the essay has tried to do, and what the eleven essays preceding it have tried to do, is lay out the structural conditions under which the question becomes visible and precisely stated. Not because stating it produces an answer, but because the question has been unasked for long enough that its absence is now itself one of the most consequential facts about the industry's near future.

The question will be answered. The only thing in doubt is whether the answer is the one the industry's participants would have chosen with their eyes open.

When the value migrates to the supply-surface layer, who holds the infrastructure position, on what commercial terms, and under what governance? The industry has perhaps twenty-four months, at the outside, to decide whether this is a question it intends to ask itself, or one it intends to have answered in its silence.⁵

ENDNOTES

1. The four-part question stated here synthesises arguments from across the prior arc. Coherence draws on Essay 8 (agent-era supply requirements: machine-readable product, real-time availability, deterministic transaction, commercial reachability), grounded in Wooldridge, M., *An Introduction to MultiAgent Systems*, Wiley, 2nd edition 2009, and Shoham, Y., and Leyton-Brown, K., *Multiagent Systems*, Cambridge University Press, 2009. Commercial terms and long-tail access draw on Essay 9 (the migration pattern and the two paths) and Essay 11 (2030 scenarios and structural axes), grounded in Christensen, C. M., and Raynor, M. E., *The Innovator's Solution*, Harvard Business School Press, 2003, and Eisenmann, T., Parker, G., and Van Alstyne, M. W., "Platform Envelopment," *Strategic Management Journal*, 32(12), 2011, pp. 1270-1285. Governance draws on Essay 11's hybrid-case third variable on neutral-layer governance design.
2. The durability and timing of infrastructure-position formation are treated formally in Essay 9 note 9 (Eisenmann, Parker, Van Alstyne 2011, on platform envelopment and the irreversibility of integrated-position acquisition once held at scale) and developed empirically in Essay 10 (the five-conditions and five-pressures argument on the inflection window).
3. The individual and coordinated diagnostic referenced here is grounded in Essay 11's "Reading the two outcomes side by side" comparison, which decomposes the two scenarios into five structural axes (supply-coherence ownership, commercial terms of agent reach, position of the long tail, distributor competition, distribution of economic surplus). Each axis is independently evaluable by any participant in the industry against their own commercial position.
4. The default-gravity argument is the central claim of Essay 11's hybrid-middle-case section: in the absence of coordinated work, the integrated path resolves first because supply coherence is easiest to achieve inside a perimeter that already owns the demand side. The argument is

grounded in Eisenmann, Parker, Van Alstyne 2011 (cited above) on platform envelopment as the structural mechanism, and in Coase, R. H., "The Nature of the Firm," *Economica*, 4(16), 1937, pp. 386-405, and Williamson, O. E., *The Economic Institutions of Capitalism*, Free Press, 1985 (both cited in Essay 6 note 2), on the transaction-cost basis for ownership consolidation under coordination failure.

5. The window-timing claim referenced here is a structural projection from Essay 10's five-conditions and five-pressures argument and Essay 11's hybrid-case timing variable (the second of three): a neutral layer that forms before integrated platforms have built supply-surface coverage at scale ends up as the structural default; one that forms after does not displace what is already in place. The "twenty-four months" figure is offered as an upper-bound order of magnitude rather than a precise estimate; the structural point is that the window is short and closing rather than that any particular date is decisive.

About the Author and Disclosures

Authorship, citation relationships, and the stance the book takes.

About the author

Marc Wieland was the founder and CEO of Palisis. He serves today as Group Executive Chairman of the Experience Technology Group Holding (ETG Holding AG), based in Switzerland. He writes on the structural dynamics of the tours, activities and experiences industry. The arguments in this book are derived from a first-principles reading of the market and reflect no specific organisation's strategy. Correspondence: [linkedin.com/in/marcwieland](https://www.linkedin.com/in/marcwieland).

Citation relationships

This book draws on Arival research at multiple points across Essays 1, 2, 3, 4, 5, 6, 7, 8, 10 and 11. Arival is the category's principal independent research and convening body, and its

field reports underpin much of the empirical claim-making in those essays.

Palisis is listed among Arival's fielding partners. Most recently, Palisis is one of the 44 fielding partners of the *Global Operator Landscape 4th Edition* (January 2026), cited principally in Essays 1 and 2. Palisis was previously among the partners who helped distribute Arival's *Technology & Connectivity* report (January 2023, cited in Essays 3, 4, 7, 8 and 9) and the *Guide to Channel Management: Strategy, Connectivity & Technology* (November 2023, cited in Essay 6). The relationship is a distribution role, not an editorial one. Arival's editorial positioning, research sponsorship and analytical findings are independent of the fielding-partner relationship. The disclosure is made in the interest of citation integrity.

In addition, Palisis has contributed to Arival's *Guide to AI in Experiences* (2024), cited in Essays 7, 8, 10 and 11, in an advisory capacity. The contribution is to category-level framing, not to specific findings or recommendations.

Palisis operates TourCMS, a connectivity platform in the channel-manager category analysed structurally in Essay 6. The category is treated in structural terms throughout the book, without reference to any specific product, including TourCMS. TourCMS appears by name in two places in the cited research, and is referenced in Essays 3 and 4 only as a category-level instance: it is included in the 13-firm provider matrix of the *Guide to Channel Management* (Essay 3), and on the public Top Connectivity Partners list maintained by a

major distributor (Essay 4). In both cases the reference is to the category, not to the product.

Stance

No participant in the experiences industry is named as the subject of an active argument anywhere in this book. Where firms are referenced, the reference is to historical, institutional or bibliographic fact. The arguments are structural and apply to the category, not to any specific actor within it.