

The Evolution of Holding Companies

Control, client value and three paths for marketing work

By **Evgeny Popov** · No Fluff Advisory · Research edition, 7 October 2026 update · research cutoff 5 October 2026 · published 6 October 2026, updated 7 October 2026

01 The groups have bought the parts of the loop between a brand's goal and its next decision. The dated record runs from the Digitas agreement of December 2006 to the IPG closing of November 2025. The audited record holds no matched client result for any of these deals.	02 Buying scale has not gone away. In the only comparable pair of readings held here, the ANA reported principal-media use at 58% of 114 client-side members in 2026, up from 47% in its prior wave. That is a member incidence, not a share of spend.	03 The record carries real counterexamples: a €1.44 billion Publicis charge in 2016 citing integration delays, a ¥396.1 billion Dentsu goodwill impairment for 2025, a staged WPP sale of Kantar control, and a digital-first challenger whose 2025 net revenue fell 10.8%.
04 The last large link is still pending. Publicis announced its LiveRamp agreement on 17 May 2026. The latest dated issuer status, a 16 September 2026 bond release, still said approvals were pending. No completion was found by the 5 October 2026 cutoff. This paper never calls the asset owned.	05 Value captured and value created are different quantities. Spreads, subscriptions, data services and switching costs may let a seller's revenue grow faster than its delivery labour; in this review's register they are also the mechanisms with no client-side test on record. The layers the groups are adding can improve client outcomes, reinforce principal-media economics or raise switching costs. Their value depends on the task, after full cost.	

210 atomic claims in the public ledger: 160 active, 10 withdrawn, 40 editorial; 148 supported as stated	51 sources in the public register, each with a stored audit passage: 49 passages inspected in the prior audit, 2 repositories pinned by commit, 49 rechecked by search only, none re-opened in this edition	104 numerical and comparability checks
255 rows in the corrections ledger	21 figures, each with a text equivalent and an editable SVG	

The findings and sources are unchanged since the update of 6 October 2026. Since then the 20 figures were rebuilt and checked twice against the live page, a 21st figure in section 20 sets the present-day loop beside a possible future, and the retailer example in section 26 now pauses new spending after a failed evidence gate. No source was re-opened; the research cutoff is still 5 October 2026. Every dated change is listed in the update history after the method section. Update history (6 dated notes)

ABSTRACT

What this paper covers

Over twenty years the large marketing groups bought the parts of the loop between a brand's goal and its next decision: digital skill, enterprise systems, identity and customer data and commerce operations. The original edition's records, not re-opened here, add creators, production and AI tools. This edition reads that record through dated deal, impairment and results documents. The groups are adding data, software and coordination to buying scale and specialist services. Those layers can improve client outcomes, reinforce principal-media economics or raise switching costs, and the record proves the bid, not the client's gain. Their value depends on the task: whether they improve creative quality, decisions, delivery or incremental results after the full cost of the arrangement. Buying scale persists beside the new tools. Brands keep accountability for goals and evidence while choosing the degree of operational control and portability they can justify. What a brand keeps is decided by the rights it can test, stop and take away: set the goal, act, stop the spend, challenge the proof, keep the learning. Part I gives the audited record for six groups and the alternatives. Parts II and III set the tests and three conditional paths.

HOW TO READ THIS

31 sections. Every paragraph carries its classification (fact, attributed, inference, assumption, scenario, analysis or illustrative) and every figure its status (observed, illustrative or conditional).

SECTION 1

The argument after review

The groups bought the parts of the loop. The record proves the bid, not the client's

gain.

IN SHORT

What a brand keeps turns on the rights it can test, stop and take away. It does not turn on who owns the most tools. The layers the groups add can improve outcomes, reinforce principal economics or raise switching costs; their value depends on the task after full cost.

A holding company is a parent firm that owns or controls several marketing businesses. This paper calls it a group. For twenty years the large groups have bought new kinds of work. They bought digital skill, enterprise systems, identity and customer data and commerce operations; the original edition's records, not re-opened here, add creators, production and AI tools. The deal records are real and dated. Publicis agreed to buy Digitas in December 2006. It closed Sapient in February 2015 and Epsilon in July 2019 ^{[1][2][3]}. Omnicom closed Flywheel on 2 January 2024 and IPG on 26 November 2025 ^{[4][5]}. Stagwell and MDC combined on 2 August 2021 ^[6]. ^{[1][2][3][4][5][6]} . **FACT** A008 A009 A025 A036

FIGURE 1 ILLUSTRATIVE


Five control questions

Control is five rights a brand can test, stop and take away. A contract, a log and a tested exit grant them. No system design does.

Control is five rights a brand can test, stop and take away

Illustrative

A contract, a log and a tested exit grant them; no system design does. A set of five, not a cycle: no order, no cause. Each card: the question, the answer in a word, who is accountable (the section 27 proposal) and what to ask to see.



Roles are proposed in section 27, not observed. No source in the record shows a brand holding all five rights (section 20).

Five control questions ILLUSTRATIVE Control has five parts, a set of rights, not a cycle: who sets the goal (owner); who may act (grant); who can stop the spend (stop right); who can challenge the proof (separate proof); who keeps the learning when the deal ends (tested exit). A contract, a log and a tested exit grant them, not a system design. Roles shown are proposed; a title grants no legal authority, and no role is a separate hire. Illustrative. Units: none; a set of rights, not stages or timing. Scope: the paper's own design. Text equivalent: appendix "Figure text equivalents".

This paper reads those purchases as parts of one loop: the loop that runs from a brand's goal to its next decision. The record proves the bid. It does not prove the client's gain. Filings show assets bought, charges taken and group margins earned. They almost never show a matched client result: profit after fees, internal cost and long-term brand effects. The one client-reported result in this review is a recognition rate, not a sales figure. It comes from a vendor release with no control group [7]. That is the gap this paper keeps open on purpose. [7] · INFERENCE A028

A029 A048

Buying scale has not gone away. Principal media is media an agency buys on its own account and resells to the client. In March 2026 the ANA reported that 58% of 114 qualified client-side members had used principal media in the past year, up from 47% in its prior wave. Fieldwork ran from October 2025 to January 2026 ^[8]. That is a member sample, not a share of spend. It is still the only comparable pair of readings this review holds. It points up, not down.^[8] · **ATTRIBUTED** A061 A062 A063

The thesis is therefore narrow and testable. The groups are adding data, software and coordination to buying scale and specialist services. Those layers can do three different things: improve client outcomes, reinforce principal-media economics, or raise switching costs. The record proves the bid, not the client's gain. The value of each layer depends on the task: whether it improves creative quality, decisions, delivery or incremental business results after the full cost of the arrangement. Two effects sit behind that. Shared tools can make a group more useful to a client. The same tools can make the group's own media and data products easier to sell. The new way of doing the work can prop up the old way of earning money. Neither effect proves the other. A client must test both. What a brand keeps is decided by the rights it can test, stop and take away. Those rights explain who captures value; on their own they do not explain what creates it. A brand keeps accountability for its goals and its evidence while choosing the degree of operational control and portability it can justify. · **ANALYSIS** A001

Control is the useful word. It has five parts. Who sets the goal? Who may act? Who can stop the spend? Who can challenge the proof? Who keeps the learning when the deal ends? No system design grants those rights by name. A contract, a log and a tested exit do. Everything that follows is built around that claim. · **ANALYSIS**

Three paths can coexist. A brand may run the core itself. An agency may run it for the brand. A platform may run much of the work inside its own system. The right route can differ by market and by task. No route wins by name. An agency can earn a premium, a closed platform can create value, and a brand can delegate operation without giving up accountability; each is a case to test, not a rule. Part III sets out the three paths as conditions, not forecasts. It names the one decision a CMO can take in each. · **ANALYSIS**

One more thing stays in view throughout. Cheaper production does not automate judgment: what to say, what to stop, what is distinct and what the brand should mean. Cheaper versions are not a stronger idea. That judgment is the part of the loop a brand must keep, or knowingly buy with a right to overrule. · **ANALYSIS** A047

SECTION 2

What this edition covers

Cutoff 5 October 2026. History from January 2006. Every count says what it is

counted against.

IN SHORT

A public-source study of six groups and selected alternatives. No source was re-opened in this edition's own review, and that limit is part of the result. An external review of 6 October 2026 is named where this update relies on one of its checks; those checks are attributed, not adopted.

This is a public-source research edition. It is written for CMOs and senior leaders in media, advertising and technology. The research window runs from 1 January 2006 to 5 October 2026. It covers Publicis, WPP and WPP Media, Omnicom and the IPG lineage, Dentsu, Havas and Stagwell. It also looks at selected smaller firms, independent agencies and media platforms. It is a set of cases that matter. It is not a census of every subsidiary, affiliate or local brand. · **FACT** A004 A005

The evidence is strongest for large global groups and for US and European brand buyers. The surveys quoted are member samples: 114 ANA client-side members in 2026, 162 ANA respondents in 2023, 45 companies in the WFA and Observatory study, and 69 multinationals in the WFA 2026 study ^{[8][9][10][11]}.

None is a market share or a global rate. This paper gives no full legal and financial comparison for China, India, Africa, Latin America or most Middle East markets. Conclusions about rights, proof and exit travel across markets. Prices, fee forms and legal buying roles do not. ^{[8][9][10][11]} · **ANALYSIS** A063 A065 A068

A069

Coverage, with denominators

The audit behind this edition is dated 5 October 2026. It read named passages in 49 sources. It judged 72 selected atomic claims from the original paper and 20 more about Horizon, 92 in all. It recomputed the 32 inherited calculations; 31 reproduce, and one, a statutory-margin check, cannot be recomputed from the stored records and is not printed. The calculation ledger of this edition holds 104 numerical and comparability checks: those 32 recomputations, 48 numbers checked from the revised edition against their stored inputs, 15 task-model edge cases, 7 peer-growth observations and 2 precision rules. Only the 32 are recomputations in the strict sense. It listed all 362 paragraphs of the original edition. The original register of 232 claim records and 329 financial observations was not all audited again. Those records stay in a marked private archive. The counts the build computes for this page are 210 claim records, 51 source records and 255 correction entries. · **FACT**

In this edition's own review no source was re-opened. Network policy blocked every source host. So each of the 49 sources was rechecked only at search level: did its URL, title and date still appear in results dated on or before the cutoff? A search result is a discovery aid, not support. The stored audit passages were

carried forward unread, and the source register labels them as such. Two protocol repositories were cloned through git and pinned to a commit. An external critical review dated 6 October 2026 reports that it opened 13 of these sources and reproduced all 47 arithmetic rows. Where this update relies on one of its checks, the text names the review. Those checks are attributed to it; they are not this paper’s own verification, and the re-opening of the 49 sources remains open. · **FACT**

The review used three passes: evidence, challenge and editing. One assistant made all three. They were not three separate reviews. There were no client interviews, no private fee audits, no contract reads and no live buying tests. · **FACT**

The evidence rule

Facts in this edition rest on one of two tiers. Tier 1 is a claim the audit found supported or partly supported, within its stated limits; the block names the claim id and the source. Tier 2 is a dated statement from one of the 49 register sources, written as what the source said on its date. Anything found only by search, or only in the original edition's archive, is not used as fact here. Where such a record would have mattered, the text says what this review did not show. It names the check that would settle it. · **FACT**

The five evidence gates used throughout

Column classifications: · **ANALYSIS**

GATE ANALYSIS	WHAT IT SHOWS ANALYSIS	WHAT IT DOES NOT SHOW ANALYSIS
Announcement	A dated claim about what a firm can do, or about a deal.	That the thing was built, used or closed.
Documented implementation	Named systems, tasks, inputs, outputs and limits.	That any client uses it.
Named live use	A named deployment with role, market, date and approval scope.	That it works at scale or pays.
Meaningful scale	Active clients or spend, repeat use, duration and failure rate.	That the client gained.
Independently demonstrated value	A credible test with full cost, quality and added effect.	Supports value within the tested setting; durability, transferability and performance at larger scale still need evidence.

FIGURE 2 OBSERVED


From a claim to evidence

Five gates separate what a claim shows from what it does not. Most of the record stops at the first two. Empty evidence is not a zero score and not proof of no capability.

From a claim to evidence: five gates

Observed

Each gate says what a record shows and what it does not. The record in this paper mostly stops at the first two. Thin grey arrows: reading order, what would move a claim to the next gate. Not every gate is needed for every useful study. Empty evidence is not a zero score and not proof of no capability.



Examples sit at the gate the paper assigns them. Select a gate to read them with their claim ids and sources.

From a claim to evidence OBSERVED The five evidence gates (section 2): announcement; documented implementation; named live use; meaningful scale; independently demonstrated value. Each card states what a record shows and what it does not; the arrows are reading order, not a prerequisite chain. A second view places the section 17 items on the paper's six protocol stages, linked only to the record and version verified. Empty evidence is not a zero score. Observed. Units: none; gates and stages, not scores. Scope: the paper's own evidence rule (section 2) and the section 17 table: documentation-level claims; no source re-opened; no test run.^{[8] [12] [13] [14] [7] [15] [16] [17] [18] [19] [20] [21] [22] [23] [24] [25] [26]} A013 A028 A029 A045 A048 A055 A056 A057 A060 A061 A062 A063 H005 H006 H007 H008 H009 H010 V3-009 V3-043 V3-050 V3-089 Text equivalent: appendix "Figure text equivalents".

How to read the labels. A fact cites a source. A reported statement is what a company or body said on a date. A supported reported claim means the speaker made the claim, not that the promised result followed. Analysis is this paper's reasoning. An assumption or scenario states a condition, not a forecast. An illustrative number exists for arithmetic only. Missing public proof means not shown here. It does not mean the skill is absent. Where this paper says a brand must, should or never does something, that is a design choice of the paper, carried

in a block labelled analysis, scenario, assumption or illustrative; it is not a finding about prevailing practice. · **ANALYSIS** A045

SECTION 3

Use the right unit

A parent, an operating group, an agency brand, a product, a non-equity partnership, an equity partnership and a joint venture are seven different things. Keep them apart.

IN SHORT

Most errors in this field come from mixing units. A rename is not an acquisition. A joint venture is not a merger.

Most confusion about groups comes from the unit of analysis. A parent owns businesses. An operating group sits inside a parent. An agency brand is a route to clients and talent; it need not be a legal company. A product is software or a service with an owner and a licence. A non-equity partnership is a contract. An equity partnership is a shared stake. A joint venture is a jointly owned business whose parents keep their own. Each has its own dates, rights and economics. The table fixes the working terms used here. · **ANALYSIS**

Who owns what, since when

Six groups and one independent, on the dated record only. Owned, accessed, shared and proposed are kept apart. A count of deals says nothing about how good a skill is.

Who owns what: Publicis

Type · relation · status and dates · verification · capability columns (analysis, not a score; none mapped is not zero).

Publicis Groupe

Parent · Parent, the owner of the rows below · columns: none mapped

Chief Client Officers across disciplines (plan of 2 December 2015)

Operating group · Structure or plan · In place Jan 2016 · record 2 Dec 2015 · A plan record; not today's roster · columns: none mapped

Digitas

Business · Owned · Purchase announced 20 Dec 2006 · Owned on the record; not verified by this review: closing not checked · columns: Creative / production (announced or proposed), Activation (announced or proposed)

Sapient (Publicis.Sapient)

Business · Owned · Completed 6 Feb 2015 · Ownership verified by this review: completion record held · columns: Strategy / consulting

Epsilon

Business · Owned · Completed 1 Jul 2019 · record 2 Jul 2019 · Ownership verified by this review: completion record held · columns: Data / identity

Lotame

Business · Owned · Completed Apr 2025 · record H1 2026 report, note 3.2 · Owned on the company's own note; to be re-read · columns: Data / identity

LiveRamp

Business · Proposed purchase · Agreed 17 May 2026 · pending, status 16 Sep 2026 · Not owned: pending at the cutoff; never shown as owned · columns: Data / identity (announced or proposed)

Retail media products (June 2023)

Product · Built or launched · Launched 16 Jun 2023 · Issuer's own description; no deal record · columns: Activation, Commerce

Who owns what: WPP

Type · relation · status and dates · verification · capability columns (analysis, not a score; none mapped is not zero).

WPP

Parent · Parent, the owner of the rows below · columns: none mapped

GroupM, renamed WPP Media (28 May 2025)

Operating group · Part of the group · Renamed 28 May 2025 · Issuer's own description; no deal record · columns: Activation

Kantar (60% sold December 2019; 40% retained at that date)

Business · Stake kept · Part sold 5 Dec 2019 · The stake at that date; not checked since · columns: Measurement (announced or proposed)

InfoSum

Business · Announced · Purchase announced 3 Apr 2025 · Not drawn as owned: completion not in the record held · columns: Data / identity (announced or proposed)

Mindshare

Agency brand · Part of the group · Named in the record 28 May 2025 · Issuer's own description; no deal record · columns: none mapped

Wavemaker

Agency brand · Part of the group · Named in the record 28 May 2025 · Issuer's own description; no deal record · columns: none mapped

EssenceMediacom

Agency brand · Part of the group · Named in the record 28 May 2025 · Issuer's own description; no deal record · columns: none mapped

Xaxis (trading desk, launched June 2011)

Trading desk · Built or launched · Launched Jun 2011 · Issuer's own description; no deal record · columns: Data / identity, Activation

Who owns what: Omnicom and IPG

Type · relation · status and dates · verification · capability columns (analysis, not a score; none mapped is not zero).

Omnicom

Parent · Parent, the owner of the rows below · columns: none mapped

Omnicom Media

Operating group · Part of the group · In place on the record 28 Aug 2026 · Issuer's own description; no deal record · columns: Activation

Interpublic (IPG), a lineage inside Omnicom from 26 November 2025

Business · Owned · Completed 26 Nov 2025 · Ownership verified by this review: completion record held · columns: Data / identity, Activation

Flywheel

Business · Owned · Completed 2 Jan 2024 · Ownership verified by this review: completion record held · columns: Commerce

Acxiom (IPG lineage)

Business · Owned · Named in the record Jul 2026 · Owned on the record; not verified by this review: the record is unread · columns: none mapped

Hearts United (Hearts & Science with Mediahub, 28 August 2026)

Agency brand · Part of the group · Renamed 28 Aug 2026 · Issuer's own description; no deal record · columns: none mapped

OMD, PHD, UM, Initiative (siblings named in the Hearts United release)

Agency brand · Part of the group · Named in the record 28 Aug 2026 · Issuer's own description; no deal record · columns: none mapped

Omni (the group's shared platform)

Product · Part of the group · Named in the record 2 Jan 2024 · Owned on the record; not verified by this review: the record is unread · columns: none mapped

Who owns what: Dentsu

Type · relation · status and dates · verification · capability columns (analysis, not a score; none mapped is not zero).

Dentsu

Parent · Parent · Impairment booked FY2025 · columns: none mapped

Aegis (Dentsu Aegis Network from 26 March 2013)

Business · Owned · Completed 26 Mar 2013 · Ownership verified by this review: completion record held · columns: Activation

Merkle (100% from 15 April 2020)

Business · Owned · Completed 15 Apr 2020 · Ownership verified by this review: completion record held · columns: Data / identity

Who owns what: Havas

Type · relation · status and dates · verification · capability columns (analysis, not a score; none mapped is not zero).

Havas (listed on Euronext Amsterdam 16 December 2024)

Parent · Parent · Listed 16 Dec 2024 · columns: none mapped

Havas Media Network

Operating group · Part of the group · In place on the record 29 Sep 2025 · Issuer's own description; no deal record · columns: none mapped

Horizon Global (joint venture with Horizon Media, 50/50)

Joint venture · not a merger · Announced 29 Sep 2025 · Issuer's own description; no deal record · columns: Activation (announced or proposed)

Who owns what: Stagwell

Type · relation · status and dates · verification · capability columns (analysis, not a score; none mapped is not zero).

Stagwell (combined with MDC 2 August 2021)

Parent · Parent · Recorded Q3 2025 · columns: Creative / production, Activation, Measurement

Enterprise-software revenue label (\$16 million committed, H1 2026)

Reported figure · not an entity · Reported in a filing of 30 Jul 2026 · Not an entity: no product entity in the record held · columns: none mapped

S4 Capital (counterweight, not part of Stagwell)

Parent · Counterweight · Not part of the group · results Mar 2026 · not part of the group · columns: none mapped

Who owns what: Horizon Media (the independent agency-group case)

Type · relation · status and dates · verification · capability columns (analysis, not a score; none mapped is not zero).

Horizon Media Holdings (private parent)

Parent · Parent, the owner of the rows below · columns: none mapped

Horizon Media (the agency issuing the releases)

Business · Part of the group · Named in the record 18 Dec 2025 · Issuer's own description; no deal record · columns: none mapped

HorizonOS (launched 18 December 2025)

Product · Built or launched · Launched 18 Dec 2025 · Issuer's own description; no deal record · columns: Strategy / consulting, Activation

Blu (announced 18 June 2026)

Product · Announced · Announced 18 Jun 2026 · Issuer's own description; no deal record · columns: Data / identity (announced or proposed), Activation (announced or proposed)

HorizonOS Labs (partner programme)

Programme · access, not ownership · Cohort announced 1 Oct 2026 · Issuer's own description; no deal record · columns: none mapped

SharkNinja (named client; audience collaboration with retailers)

Named client · Named client use · Use reported 18 Jun 2026 · Issuer's own description; no deal record · columns: none mapped

Horizon Global (joint venture with Havas Media Network, 50/50)

Joint venture · not a merger · Announced 29 Sep 2025 · Issuer's own description; no deal record · columns: none mapped

Who owns what, since when **OBSERVED** Six groups and the independent agency-group case (Horizon Media), each in two short views. Ownership shows each entity's name, type, relation to the parent and status; the dates, whether this review verified ownership, and the passages behind each row are in its detail and the tables. Capability mapping is analysis, not a score. LiveRamp stays pending. Status: observed. Units: none. Scope: six groups and the independent agency-group case on the Tier 1 and Tier 2 record held; a set of cases, not a census. [6] [27] [28] [29] [30] [31] [32] [33] [34] [13] [14] [35] [4] [5] [7] [36] [1] [2] [37] [3] [17] [18] [38] [19] [39] [40] [26] [41] [42] A006 A007 A008 A009 A010 A011 A012 A013 A014 A015 A016 A022 A023 A024 A025 A026 A027 A028 A029 A030 A031 A032 A036 A037 A038 A039 A040 A048 H001 H002 H003 H004 H005 H006 H008 H009 H010 H011 H012 H017 H018 H019 H020 V3-009 V3-043 V3-048 V3-050 V3-051 Text equivalent: appendix "Figure text equivalents".

Working terms

Column classifications: **ANALYSIS**

TERM ANALYSIS	MEANING IN THIS PAPER ANALYSIS
Holding company (group)	A parent that owns or controls several marketing businesses.
Operating group or media network	A structure inside the parent that runs several agencies or teams.
Agency brand	A client-facing name. It may not be a separate legal entity.
Trading desk	A team or system that does specialist media buying and optimisation. It may select technology without owning it.
Product or platform	Software or a service with named functions. Owning the interface does not mean owning the data, inventory or models below it.
Commercial partnership (non-equity)	A contract for access or supply. No ownership changes hands; the rights are the contract's terms.
Equity partnership	One party takes a stake in the other, or both take stakes in a shared vehicle. Ownership changes; control depends on the terms and the size of the stake.
Joint venture	A jointly owned business with its own scope, set up by its parents, who continue outside it. Horizon Global is described by its parents as a commercial joint venture and as a 50/50 partnership; this paper treats it as a jointly owned venture.
Principal buying	A supplier buys media on its own account and resells it. The contract sets price, risk and rights.
Agency buying	A supplier acts for the client under a mandate.
In-housing	The brand does a defined task inside. The label says nothing about share of spend or of work.
Orchestration	Joining tasks, people, data, tools, rights and feedback across a process. The word says nothing about fees or neutrality.
Software agent	A system that performs a bounded sequence of tasks. It is not the client's legal agent because of its name.
Attribution	A rule that assigns credit. It need not prove added sales.
Incrementality	The change caused by an action, against a credible view of what would have happened without it.

Three rules follow. First, a rename is not an acquisition. When WPP Media replaced the GroupM name in May 2025, Mindshare, Wavemaker and EssenceMediacom stayed as client brands. No business changed hands ^[40]. Second, a retired brand is not proof that a function, a fee or a margin has ended. Omnicom's 2019 move of specialist staff into OMD, PHD and Hearts & Science kept central skills in place ^[35]. Third, a joint venture is not a merger. Horizon Global, announced on 29 September 2025, is a commercial joint venture of Havas and Horizon Media. The parents did not combine ^[33] ^[40]^[35]^[33] . ANALYSIS A006 A022 A023 A032

Accounting labels need the same care. Gross or net revenue follows an accounting judgment about who controls the service before it reaches the client. It is not a clean proxy for the media bought for resale. The gap between the two is not a disclosed principal margin. Publicis reports €17.399 billion of revenue and €14.547 billion of net revenue for 2025. The €2.852 billion difference is pass-through cost, money billed on to clients for media and other third-party costs. It is not a profit line ^[43] ^[43] . ANALYSIS A017 A018

Dated entity checks: the structure on the date of the source, not a current map^[40]^[39]^[36]^[33]^[18]

Column classifications: · FACT · ATTRIBUTED

DATE OF SOURCE FACT	PARENT / UNIT FACT	NAMED BRANDS OR ASSET FACT	WHAT THE SOURCE SUPPORTS [SOURCE] ATTRIBUTED
28 May 2025	WPP / WPP Media	Mindshare; Wavemaker; EssenceMediacom	Client brands kept at launch; GroupM name replaced. Not a full current legal entity map. ^[40]
2 December 2015 (plan)	Publicis / Chief Client Officers	Client officers across disciplines and countries; effective January 2016	A historical design, not today's roster. The original edition also cited a 10 March 2016 Publicis Media structure; that source was not among the 49 audited and is not used here. ^[39]
28 August 2026	Omnicom / Omnicom Media	Hearts United (Hearts & Science with Mediahub); UM, PHD, OMD and Initiative named as siblings	A joining of two brands, not a new acquisition. UM was not the merger partner. ^[36]
29 September 2025	Havas and Horizon Media	Horizon Global	A commercial joint venture. Parents did not merge. Combined billings are not Havas revenue. ^[33]
16 September 2026	Publicis → LiveRamp	Proposed purchase	Latest dated issuer status: approvals pending, close expected before the end of 2026. No completion verified by the cutoff. Never draw an owned link. ^[18]

Twenty years without a false straight line

The scope of work grew around the buying engine. So did the number of reversals.

IN SHORT

A dated event table replaces the smooth story. Every row is a closed, pending, renamed or impaired event with its own source.

The history reads as layers. From 2006 the groups bought digital breadth and built specialist buying teams. From the middle of the 2010s they bid for customer and enterprise work: consulting, engineering, CRM (customer relationship management systems) and identity data. From the late 2010s they merged brands and put a client officer across disciplines. The 2020s added commerce on the audited record; creators, production tools and AI rest on records not re-opened here. Each layer sat on the last. None replaced it. The dates below are anchors, not stages that every group passed on the same day. · ANALYSIS

Two rules govern the table. Only events that a Tier 1 or Tier 2 record supports appear. And the status column keeps announced, closed, pending, renamed and impaired events apart. A plan, a purchase and a write-down are different facts. The original edition's timeline held about sixty events. The rows not checked again here stay in the private archive and are listed in section 31. They are not smoothed into a series. · ANALYSIS

Dated events, 2006–2026, from Tier 1 and Tier 2 records only [1] [42] [29] [2] [39] [37] [3] [35] [41] [30] [6] [19] [4] [32] [26] [38] [40] [33] [5] [34] [31] [17] [13] [36] [18] [14]

Column classifications: · FACT · ATTRIBUTED · ANALYSIS

DATE FACT	GROUP FACT	EVENT ATTRIBUTED	ENTITY TYPE ANALYSIS	STATUS FACT	SOURCE FACT
20 December 2006	Publicis	Agreement to buy Digitas; proposed value \$1.3 billion; stated motives digital expertise and distribution	Agency (digital)	Announced (closing date not checked again here)	[1]
27–28 June 2011	WPP / GroupM	Xaxis launched, joining data and technology with GroupM trading scale	Trading desk	Launch	[42]
July 2012	Dentsu	Agreement to buy Aegis announced	Parent-level deal	Announced	[29]

DATE FACT	GROUP FACT	EVENT ATTRIBUTED	ENTITY TYPE ANALYSIS	STATUS FACT	SOURCE FACT
26 March 2013	Dentsu	Aegis purchase completed; Dentsu Aegis Network established	Parent-level deal	Closed	[29]
6 February 2015	Publicis	Sapient purchase completed; approximate consideration \$3.7 billion as stated in the completion release	Consulting and engineering firm	Closed	[2]
2 December 2015	Publicis	Transformation plan: Chief Client Officers across disciplines and countries, effective January 2016	Operating structure	Reorganisation	[39]
9 February 2017 (FY2016 results)	Publicis	€1,440 million impairment, including €1,392 million for Publicis.Sapient; legacy digital assets and integration delays cited	Accounting charge	Impairment	[37]
2 July 2019	Publicis	Epsilon completion announced; consolidated from 1 July 2019	Data and CRM business	Closed	[3]
2019 (publication date not re-shown)	Omnicom	Specialist programmatic, search and social staff moved into OMD, PHD and Hearts & Science; central skills retained	Operating structure	Reorganisation	[35]
5 December 2019	WPP	Sale of 60% of Kantar closed for about 90% of the business; remainder expected in 2020; 40% retained	Research unit	Disposal (staged)	[41]
15 April 2020	Dentsu	100% ownership of Merkle confirmed; payment for the remaining stake due in the third quarter of 2021	Agency (CRM and data)	Closed (payment later)	[30]
2 August 2021	Stagwell / MDC	Stagwell and MDC combination completed	Parent-level combination	Closed	[6]
16 June 2023	Publicis	Retail media products launched; bidding described as triggered by stock and price signals	Product	Launch	[19]

DATE FACT	GROUP FACT	EVENT ATTRIBUTED	ENTITY TYPE ANALYSIS	STATUS FACT	SOURCE FACT
2 January 2024	Omnicom	Flywheel purchase closed; planned link between commerce data and Omni stated	Commerce business	Closed	[4]
16 December 2024	Havas	Listing on Euronext Amsterdam after the spin-off from Vivendi	Parent	Listing	[32]
3 April 2025	WPP	InfoSum purchase announced for the data offer	Data-collaboration technology	Announced	[26]
April 2025	Publicis	Full purchase of Lotame, per note 3.2 of the half-year 2026 financial report	Data business	Closed (per company note; to be re-read)	[38]
28 May 2025	WPP	WPP Media replaces the GroupM name; Mindshare, Wavemaker and EssenceMediacom remain	Media network	Rename	[40]
29 September 2025	Havas and Horizon Media	Horizon Global commercial joint venture announced	Joint venture	Announced	[33]
26 November 2025	Omnicom	Purchase of Interpublic (IPG) completed	Parent-level deal	Closed	[5]
18 December 2025	Horizon Media	HorizonOS launch: specialist staff, client pilots and a partner marketplace described	Product and programme	Launch	[34]
FY2025 (report published 2026)	Dentsu	Goodwill impairment of ¥396,074 million; ¥299,656 million Americas and ¥96,418 million EMEA	Accounting charge	Impairment	[31]
17 May 2026	Publicis	Agreement to buy LiveRamp announced; equity value \$2.546 billion; closing subject to conditions	Proposed purchase	Pending	[17]
18 June 2026	Horizon Media	Blu agentic orchestration layer announced; APIs, MCP connections and agentic buying capabilities described	Product	Announced	[13]

DATE FACT	GROUP FACT	EVENT ATTRIBUTED	ENTITY TYPE ANALYSIS	STATUS FACT	SOURCE FACT
28 August 2026	Omnicom	Hearts United launched, combining Hearts & Science with Mediahub	Agency brands joined	Combination (rename)	[36]
16 September 2026	Publicis	Bond release: LiveRamp purchase still subject to approvals; close expected before the end of 2026	Transaction status	Pending	[18]
1 October 2026	Horizon Media	HorizonOS Labs third cohort: 11 partners chosen from more than 600 submissions	Programme	Announced	[14]

FIGURE 4 OBSERVED


Dated events, 2006–2026

Twenty years of dated moves by six groups and one independent. Each is a record with its own status. There is no straight line and no ranking.

Dated events by group, 2006 to 5 October 2026

26 events from Tier 1 and Tier 2 records; a selection, not a census. Nothing ranks the groups.

Level ● Parent deal ■ Group ◆ Brand ▲ Product ⬡ Partnership | Record
Status ● Closed, launched or recorded ○ Announced only ⬢ Pending at the cutoff



Dated events, 2006–2026 **OBSERVED** Twenty-six dated events from Tier 1 and Tier 2 records on one linear axis, one lane per group plus Horizon Media as an independent. Shape gives entity level; fill gives status: filled closed, launched or recorded; hollow announced only; dashed pending (LiveRamp, agreed 17 May 2026, approvals pending at the 16 September 2026 bond release). A second view shows 2024 to the cutoff on its own axis. Nothing ranks the groups. Status: observed. Units: Dates; currency amounts as reported by each company. Scope: Events supported by Tier 1 or Tier 2 records only; the original edition's wider timeline stays in the private archive.^{[6] [28] [29] [30] [31] [32] [33] [34] [13] [14] [35] [4] [5] [36] [1] [2] [37] [3] [17] [18] [38] [19] [39] [40] [26] [41] [42]} A006 A007 A008 A009 A010 A011 A012 A013 A022 A023 A024 A025 A026 A027 A030 A031 A032 A036 A038 A039 A040 H002 H003 H004 H005 H008 H009 H010 H018 V3-043 V3-048 V3-050 V3-051

Text equivalent: appendix “Figure text equivalents”.

Read the table for what it does not show. It holds no price for most deals, because the audited records state none. It holds no abandoned deal. The original edition had two: the Publicis–Omnicom merger plan, ended in 2014, and the Dentsu international sale, ended in February 2026. Both are carried forward unread and listed in section 31. And it holds no client result at all. Twenty years of buying left a long record of what the groups wanted and a short record of what clients gained. · **ANALYSIS**

The trading desk did not vanish

The trading desk is the clearest lesson in units. A trading desk is a team or system that buys media through automated auctions, which the industry calls programmatic buying. Xaxis launched in June 2011. It joined data and technology with GroupM's trading scale ^[42]. By 2019 Omnicom's media group had moved its specialist programmatic, search and social staff into OMD, PHD and Hearts & Science. It kept central product, engineering and data duties ^[35]. A name disappeared. The staff moved closer to the client teams. Nothing in either record says that a fee, a risk position or a principal margin ended with the name.^[42]
[35] · **FACT** A006

The original edition traced seven desk lineages across the groups: Xaxis, Accuen, Cadreon and Matterkind, Amnet and Amplifi, Affiperf and Varick. That lineage is carried forward unread, listed in section 31, and not repeated here as fact. The conclusion it supports stands on the two records above: trace the function and the contract, not the name. · **ANALYSIS**

Selection matters too. The largest surviving groups get most of the attention. Their histories alone cannot prove that scale is the best model. A fair test must also keep sold assets, impaired assets, failed plans and viable smaller firms in view. The group chapters that follow keep those counterweights beside each chain. Section 11 adds a digital-first challenger whose net revenue fell. · **ANALYSIS**

How to read a deal

A deal buys an option to change the business. Integration and client value still have to be shown.

IN SHORT

Six motives, five assessment steps and three value pools turn each deal into a chain. Each link can be checked on its own.

This section carries no factual burden. It is the method the group chapters apply. The original edition set it out. The revised edition dropped it, though nothing in it was contradicted. It returns here because a dated deal record is useless without a way to read it. · **ANALYSIS**

FIGURE 5 ILLUSTRATIVE

How to read a deal

Read a deal link by link. The record proves the bid. Each link, the countercase and the rival reading must be filled from dated records, or left blank.

How to read a deal

Template: no company data

Six links in reading order, then two checks and three value pools. Fill each from a dated record, or leave it blank.

No case chosen: the template only. Numbers give the reading order; nothing here is a flow.

1 Prior gap What was missing before the deal? Fills it: The buyer's own statement of purpose, dated. fill from a dated record	2 Asset or skill What was bought, when, and for how much where stated? Fills it: Agreement and completion records, kept separate. fill from a dated record	3 Intended connection How was the asset meant to join other work? Fills it: The integration plan or operating design, dated. fill from a dated record
4 Use on record Did a named system, team or product actually connect? Fills it: A product description, a named deployment or a reorganisation record. fill from a dated record	5 Client effect Did a client do better after all costs? Fills it: A matched test with full cost and added effect. fill from a dated record	6 Group economics What did the group's own accounts show? Fills it: Revenue, net revenue, margin, impairment, with definitions attached. fill from a dated record

Then check: countercase What went wrong, was sold or was reversed? Fills it: Impairments, disposals, separations, lost scope. fill from a dated record	And: rival reading What else explains the same record? Fills it: Cross-selling, cost removal, investor narrative, cycle. fill from a dated record
--	---

Three value pools, kept apart None shows the others; a client win says little about margin.

Acquired team may gain: distribution fill from a dated record	Clients may gain: a better outcome, a lower total cost or less coordination work fill from a dated record	Shareholders may gain: cost removal fill from a dated record
--	--	---

Strategy readings are this paper's inference, not a deal fact. A blank link stays blank; a status word is not a score. "Not assessed in this review" means no record was opened or held for the link; "Not shown in the reviewed record" means the paper looked and found none. Neither means the capability does not exist.

How to read a deal ILLUSTRATIVE Section 5's method as a numbered checklist: six links in reading order (prior gap, asset or skill, intended connection, use on record, client effect, group economics), then the countercase and the rival reading, with three value pools kept apart. A worked case fills each link from dated records or leaves it blank. A logic, not a ranking. Status: illustrative (template). Units: none. Scope: the paper's own design.^{[6][27][28][44][29][30][31][35][4][5][7][36][45][1][2][37][3][17][18][38][19][43][46][40][47][26][41]} A006 A008 A009 A010 A011 A012 A013 A014 A015 A016 A017 A018 A019 A020 A021 A022 A023 A024 A025 A026 A027 A028 A029 A030 A031 A036 A037 A038 A039 A040 A041 A042 A048 A050 A072 V3-009 V3-043 V3-110 V3-117 Text equivalent: appendix "Figure text equivalents".

Six recurring motives

Deals in this field show six recurring motives. Buy a scarce skill. Buy a route to a client or a market. Buy a repeatable tool. Buy access to permitted data. Buy influence over a high-value decision. Or buy scale that lowers cost or improves terms. Most large deals mix several. A technology label can hide a large services business. A services label can hide valuable software, data processes and customer access. The first job is to name which motives a deal served, in the buyer's own words and then in this paper's. · **ANALYSIS**

Five assessment steps

1. Start with the buyer's missing link. What could the group not do, or not do at scale, before the deal?
2. State what the target really supplied: people, clients, software, data, rights or scale. Separate the label from the contents.
3. State which rights changed. Owning a company is not owning its partners' data, its clients' permissions or its vendors' models.
4. State how the asset was joined to other work, and which result can be seen. Separate management's stated reason from this paper's inference.
5. Compare a build or partner route where one is plausible. Do not claim that the unobserved alternative would have happened.

Three value pools that can diverge

Shareholders may gain from cost removal. The acquired team may gain distribution. Clients may gain a better outcome, a lower total cost or less coordination work. None of these shows the others. A client win may support the combined offer while showing little about margin or measured effect. A goodwill write-down, a charge that cuts the book value of a past purchase, may challenge the deal thesis without proving that every acquired skill is useless. So the group chapters keep group economics and client effect in separate rows. They leave a row blank rather than fill it by inference. · **ANALYSIS** A050

The chain template

The chain each group chapter follows

Column classifications: · ANALYSIS

LINK ANALYSIS	QUESTION IT ANSWERS ANALYSIS	EVIDENCE THAT WOULD FILL IT ANALYSIS
Prior gap	What was missing before the deal?	The buyer's own statement of purpose, dated.
Asset or skill	What was bought, when, and for how much where stated?	Agreement and completion records, kept separate.
Intended link	How was the asset meant to join other work?	The integration plan or operating design, dated.
Integration or use on record	Did a named system, team or product actually connect?	A product description, a named deployment or a reorganisation record.
Client effect	Did a client do better after all costs?	A matched test with full cost and added effect.
Group economics	What did the group's own accounts show?	Revenue, net revenue, margin, impairment, with definitions attached.
Counter case	What went wrong, was sold or was reversed?	Impairments, disposals, separations, lost scope.
Rival reading	What else explains the same record?	Cross-selling, cost removal, investor narrative, cycle.

Two rules close the method. One plausible deal logic is to enter more of the client's work: understand demand, know the customer, make the work, reach the customer, support the sale, then learn. That is a logic, not a ranking. No test in this review ranks one group's deal strategy above another, and this paper does not. The firmest cross-case inference is modest: groups want more recurring access to the client's decisions. A service contract can be renewed each year. A data process, a commerce system or an embedded workflow is harder to replace. That can create real value and switching cost at the same time. The right buyer question is whether the benefit of staying justifies the cost of leaving. · ANALYSIS V3-116

SECTION 6

Publicis: links from data to client work

The strategic case is coherent. The causal proof is narrower. The last link is still

pending.

IN SHORT

Sapient, Epsilon and Lotame are completion records; Digitas is a December 2006 agreement whose closing was not checked here. The 2016 impairment is the countercase. LiveRamp was still conditional at the last dated check.

FROM FIGURE 3 · WHO OWNS WHAT, SINCE WHEN

Publicis on the dated record

Publicis: every entity on the dated record with its type, relation to the parent, transaction status, effective date, evidence date, verification, dated events, capability columns (analysis), deal card and the passages it rests on

ENTITY	TYPE	RELATION TO THE PARENT	TRANSACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNERSHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Publicis Groupe	Parent	Parent	Parent, the owner of the rows below	none in the record held	none drawn	the parent itself	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s03/P1-0022.
Chief Client Officers across disciplines (plan of 2 December 2015)	Operating group	Structure or plan	In place Jan 2016 · record 2 Dec 2015	Jan 2016	2 Dec 2015	A plan record; not today's roster	Reorganised 2 Dec 2015	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s03/P1-0022.
Digitas	Business	Owned — Agreement 20 December 2006 (an announcement, not a closing); the closing was not checked here. The paper's deal order treats Digitas as bought; this review holds only the agreement record.	Purchase announced 20 Dec 2006	none in the record held	20 Dec 2006	Owned on the record; not verified by this review: closing not checked	Agreement announced 20 Dec 2006	Creative / production (announced or proposed), Activation (announced or proposed)	Digital expertise and distribution, in the buyer's words; the paper reads it as digital and direct-marketing skill at scale → gate reached: announcement Basis passages: s06/P1-0043, s06/P1-0042, s04/P1-0025, s31/P4-0015.

ENTITY	TYPE	RELATION TO THE PARENT	TRANS ACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNER SHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Sapient (Publicis.Sapient)	Business	Owned	Completed 6 Feb 2015	6 Feb 2015	6 Feb 2015	Ownership verified by this review: completion record held	Completed 6 Feb 2015; Impairment booked 9 Feb 2017	Strategy / consulting	Consulting and engineering → gate reached: announcement Basis passages: s06/P1-0043, s04/P1-0025.
Epsilon	Business	Owned	Completed 1 Jul 2019 · record 2 Jul 2019	1 Jul 2019 (consolidated from 1 July 2019; completion announced 2 July 2019)	2 Jul 2019	Ownership verified by this review: completion record held	Completed 2 Jul 2019	Data / identity	Identity and customer data (a data and CRM business) → gate reached: announcement Basis passages: s06/P1-0043, s04/P1-0025.
Lotame	Business	Owned — Full purchase dated April 2025 by note 3.2 of the half-year 2026 financial report; closed on the company's note, to be re-read. Lotame and LiveRamp are different deals.	Completed Apr 2025 · record H1 2026 report, note 3.2	Apr 2025	H1 2026 report, note 3.2	Owned on the company's own note; to be re-read	Completed Apr 2025	Data / identity	A data business → gate reached: announcement Basis passages: s06/P1-0043, s04/P1-0025.

ENTITY	TYPE	RELATION TO THE PARENT	TRANS ACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNER SHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
LiveRamp	Business	Proposed purchase — Agreed 17 May 2026; subject to approvals on 16 September 2026, the latest dated issuer status; pending at the cutoff; never drawn as owned	Agreed 17 May 2026 · pending, status 16 Sep 2026	none in the record held	16 Sep 2026	Not owned: pending at the cutoff; never shown as owned	Agreement announced 17 May 2026; Still pending, status checked 16 Sep 2026	Data / identity (announced or proposed)	A data-collaboration platform; proposed, not completed → gate reached: announcement Basis passages: s06/P1-0051, s03/P1-0022.
Retail media products (June 2023)	Product	Built or launched	Launched 16 Jun 2023	16 Jun 2023	16 Jun 2023	Issuer's own description; no deal record	Launched 16 Jun 2023	Activation , Commerce	Not a deal: a product launch → gate reached: documented implementation Basis passages: s06/P1-0045.

A capability mark means the paper's category words for that asset fall in this column. It is not a measure of capability quality, and a deal count is not a capability score. **Open the Publicis tab of Figure 3.**

Prior gap. Publicis set out to buy digital and direct-marketing skill at scale. Then it bought consulting and engineering. Then identity and customer data. Then commerce. Client coordination was a stated goal long before generative AI. The December 2015 plan put Chief Client Officers across disciplines and countries, effective January 2016 ^[39]. Media buying remains a named service. The sequence reads as a bid to own more of the links between a goal, customer data, content, execution and feedback. That is a reading of the deal order. It is not a claim that the links work as one system. ^[39] · **INFERENCE** A007

Asset. The Digitas agreement of 20 December 2006 proposed a \$1.3 billion deal. It named digital expertise and distribution as motives. The date is an announcement, not a closing ^[1]. Sapient closed on 6 February 2015; the closing release put the price at about \$3.7 billion ^[2]. Epsilon's closing was announced on 2 July 2019, with the business consolidated from 1 July ^[3]. The half-year 2026 financial report, in note 3.2, dates the full purchase of Lotame to April 2025 ^[38].

Lotame and LiveRamp are different deals. The first is closed on the company's note. The second is not.^{[1][2][3][38]} · **FACT** A008 A009 A020

Intended link. The July 2019 Epsilon release announced completion and consolidation^[3]. The idea is simple. One customer identity lets work from several teams reuse the same view of the customer. That could help from one campaign to the next. It also raises the questions this paper keeps asking. Who can license the identity? Who can export matched audiences? Who can check where the data came from? Who can change the measurement vendor? Who keeps the learning after exit?^[3] · **ATTRIBUTED**

Use on record. In June 2023 Publicis described retail media products in which bids are triggered by stock and price signals, with a content product that recommends page changes^[19]. That is a named input, decision and action. It is stronger than a logo chart. It is still a seller's description of an available skill, not an outside result. On the five gates it sits at the second: documented implementation.^[19] · **ATTRIBUTED**

Client effect. Not shown in any source reviewed. No matched client comparison with full cost was found. · **ANALYSIS** A048

Group economics. The 2025 Universal Registration Document shows €14.547 billion of net revenue and €17.399 billion of revenue. Net revenue rose from €13.965 billion in 2024; that is reported growth of 4.2%, recomputed from the two figures^[43]. The half-year 2026 report shows net revenue of €7.229 billion and an operating-margin subtotal of €1.261 billion. That is 17.4% after €7 million of LiveRamp transaction costs and 17.5% before them; the company headline uses the 17.5%^{[38][46]}. The subtotal is a company measure, not statutory operating income.^{[43][38][46]} · **FACT** A017 A018 A019 A020 A021

Countercase. The 2016 results booked a €1.44 billion non-cash charge. Of that, €1.392 billion fell on Publicis.Sapient. The release spoke of legacy digital assets and delays in joining them up^[37]. The charge is not a write-off of the Sapient purchase alone; the unit held more than that deal.^[37] · **FACT** A010 A011 A012

Rival reading. The 2016 charge is evidence against an easy platform story. Buying more can raise the value of joining the parts while making the joining harder. Results consistent with a strategy do not isolate its cause. Simpler cross-selling to the same clients through one client officer fits the same record. So does retention: a shared identity asset raises the cost of leaving whether or not it raises the client's return. · **ANALYSIS**

LiveRamp: a pending link, by date

The LiveRamp record is five events, and the paper keeps them apart. Publicis announced the agreement on 17 May 2026, a Sunday. The release stated an equity value of \$2.546 billion, net cash of \$379 million and an enterprise value of \$2.167 billion. It said closing was subject to conditions^[17]. The announcement also states an offer price of \$38.50 per share; the external review of 6 October 2026

confirmed the price, the equity value and the enterprise value in that announcement, which this edition did not re-open. A LiveRamp shareholder vote on 17 August 2026 appears in the search record of LiveRamp's filings. On 16 September 2026 a Publicis bond release still said the purchase was subject to the usual approvals and expected to close before the end of 2026 ^[18]. That is the latest dated issuer status checked. No completion was found by the cutoff of 5 October 2026. ^{[17][18]} . **FACT** A013 A014 V3-009

The proposed asset is not like an agency. A data-collaboration platform joins data held by many firms. Its value may depend on rivals of its owner still using it. Publicis has promised neutral, open access ^[17]. That is a promise to test. It is not outside proof of neutrality, and it is not evidence of misuse. The test is this: equal terms of access, walls around sensitive data, service quality, and whether rivals keep connecting. A scarce shared utility can lose value if its owner is not trusted to be neutral. Nothing in this paper draws an owned link. ^[17] . **ANALYSIS** A015

A016

WHAT THE RECORD PROVES, AND DOES NOT

The record proves five things. Publicis made a dated run of purchases, from digital skill to identity and commerce. It stated a client-coordination design in 2015. It described a stock-to-bid workflow. It booked one large impairment that names integration delay. And it held a conditional agreement for LiveRamp as of 16 September 2026. The record does not prove four things. It does not show that the links work as one system for any client. It does not show that any client earned more after fees. It does not give the size or margin of principal media resale. And it does not show that LiveRamp is owned. Records this update did not re-open are listed once in section 31 with their effect on this chapter.

SECTION 7

WPP: less friction, with buying scale still in view

A shared platform can change the work without changing every reason to sell.

IN SHORT

Xaxis, Kantar, InfoSum and WPP Media are dated records. H1 2026 shows the net measure down 4.7% like for like. The largest counterexample the original edition reported, the 2025 AKQA and Grey separation, is carried forward unread and listed in section 31.

FROM FIGURE 3 · WHO OWNS WHAT, SINCE WHEN

WPP on the dated record

WPP: every entity on the dated record with its type, relation to the parent, transaction status, effective date, evidence date, verification, dated events, capability columns (analysis), deal card and the passages it rests on

ENTITY	TYPE	RELATION TO THE PARENT	TRANSACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNERSHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
WPP	Parent	Parent	Parent, the owner of the rows below	none in the record held	none drawn	the parent itself	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s03/P1-0022.
GroupM, renamed WPP Media (28 May 2025)	Operating group	Part of the group	Renamed 28 May 2025	28 May 2025	28 May 2025	Issuer's own description; no deal record	Renamed 28 May 2025	Activation	Not a deal: a rename. No business changed hands → gate reached: announcement Basis passages: s03/P1-0020, s03/P1-0022.
Kantar (60% sold December 2019; 40% retained at that date)	Business	Stake kept — 60% sold on 5 December 2019 for about 90% of the business; 40% kept at that date, not a checked current stake	Part sold 5 Dec 2019	5 Dec 2019	5 Dec 2019	The stake at that date; not checked since	Sold in part 5 Dec 2019	Measurement (announced or proposed)	A disposal, not a purchase: 60% of Kantar sold → gate reached: announcement Basis passages: s07/P1-0060.
InfoSum	Business	Announced — Purchase announced 3 April 2025; terms not disclosed; completion not in the record held	Purchase announced 3 Apr 2025	none in the record held	3 Apr 2025	Not drawn as owned: completion not in the record held	Agreement announced 3 Apr 2025	Data / identity (announced or proposed)	Data-collaboration technology; terms not disclosed → gate reached: announcement Basis passages: s07/P1-0055, s31/P4-0015.

ENTITY	TYPE	RELATION TO THE PARENT	TRANS ACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNER SHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Mindshare	Agency brand	Part of the group — A client brand named as kept when WPP Media replaced the GroupM name; no deal record applies	Named in the record 28 May 2025	none in the record held	28 May 2025	Issuer's own description; no deal record	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s03/P1-0022, s03/P1-0020.
Wavemaker	Agency brand	Part of the group — A client brand named as kept when WPP Media replaced the GroupM name; no deal record applies	Named in the record 28 May 2025	none in the record held	28 May 2025	Issuer's own description; no deal record	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s03/P1-0022, s03/P1-0020.
EssenceMedia	Agency brand	Part of the group — A client brand named as kept when WPP Media replaced the GroupM name; no deal record applies	Named in the record 28 May 2025	none in the record held	28 May 2025	Issuer's own description; no deal record	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s03/P1-0022, s03/P1-0020.

ENTITY	TYPE	RELATION TO THE PARENT	TRANS ACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNER SHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Xaxis (trading desk, launched June 2011)	Trading desk	Built or launched — Built from existing teams, not bought	Launched Jun 2011	Jun 2011	Jun 2011	Issuer's own description; no deal record	Launched Jun 2011	Data / identity, Activation	Not a deal: a trading desk built from existing teams → gate reached: announcement Basis passages: s07/P1-0055, s03/P1-0019.

A capability mark means the paper's category words for that asset fall in this column. It is not a measure of capability quality, and a deal count is not a capability score. **Gap:** WPP: 24/7 Real Media, TNS, AKQA, Satalia, Goat, the AKQA and Grey separation, the three-layer stack and usage counts: carried forward unread. **Open the WPP tab of Figure 3.**

Prior gap. WPP's gaps came in this order: search and ad technology, audience buying at scale, research scale, and then one data and workflow layer shared across its agencies. Each gap was met by a different kind of move. A purchase. A launch built from inside. A staged sale. A rename. • **ANALYSIS**

Asset. Xaxis launched in June 2011. It joined the group's own data tools with GroupM's trading scale; it was built from existing teams, not bought ^[42]. WPP announced the InfoSum purchase on 3 April 2025 as an investment in its data offer. The release did not disclose the terms ^[26]. ^[42]^[26] • **FACT**

Intended link. On 28 May 2025 WPP Media replaced the GroupM name. Mindshare, Wavemaker and EssenceMediacom stayed as client-facing brands ^[40]. That is an operating change. It does not show one contract, one team or one data right for every client. It is the May 2025 launch structure, not a current legal entity map. ^[40] • **FACT** A022 A023

Use on record. The InfoSum purchase and the WPP Media launch are the recorded steps this review holds. The original edition described more: a three-layer stack, counts of staff using the platform, a production site, a self-serve route and an agent hub. Each is a sign of use reported by the company, and none is a client result. • **ANALYSIS**

Client effect. Not shown in any source reviewed. • **ANALYSIS** A048

Group economics. The 2026 interim results, published on 6 August 2026, show revenue of £6,373 million and revenue less pass-through costs of £4,745 million. The net measure fell 4.7% like for like, that is, on a constant scope and currency as the company defines them. Headline operating profit was £398 million, which

is 8.4% of the net measure. Reported operating profit was £261 million, which is 4.1% of revenue ^[47]. Those two percentages differ in both numerator and denominator. They cannot be read as one profit test. ^[47] · **FACT**

Counterexample. The Kantar sale is the recorded counterexample to endless asset building. On 5 December 2019 WPP said the sale of 60% of Kantar had closed for about 90% of the business. The remaining steps were expected in 2020, and WPP kept 40% ^[41]. The 40% is the December 2019 position, not a checked current stake. The original edition also recorded a 2020 combination of AKQA and Grey that was separated in 2025, with £641 million of goodwill impairments in the 2025 annual report. If that record holds, it is the largest reversal of an integration for any group in this paper, and it is the first WPP item in section 31. ^[41] · **FACT**

A024

Rival reading. A new name and a shared platform had not yet produced group-wide growth by the cutoff. A cost reset after weak results fits the same record as well as a strategy of selective control does. WPP attributed WPP Media's H1 2026 decline partly to prior-year client losses ^[47]; the external review of 6 October 2026 confirmed that attribution in the interim release, which this edition did not re-open. It is management's explanation of its own division, not an independently established cause. Shared tools can lower WPP's own cost while leaving client cost, quality and control unchanged. Both readings can be true at once. A new structure must earn its keep through delivery, not through less overhead.

^[47] · **ANALYSIS** V3-117

What a buyer can test now. Ask to see one brief pass through the real teams. Which steps are shared? Which still need manual repair? Who may reject a group product? Can the client keep an outside proof team? The claim weakens if lower group costs leave client cost, quality and control unchanged. · **ANALYSIS**

WHAT THE RECORD PROVES, AND DOES NOT

The record proves five things. WPP built an audience-buying desk in 2011. It sold research control in stages from 2019. It announced a data-collaboration purchase in April 2025. It renamed its media network and kept three media brands in May 2025. And its net revenue measure fell 4.7% like for like in the first half of 2026. The record does not prove two things. It does not show that the shared platform lowered any client's total cost. It does not give the group's principal share or spreads. Records this update did not re-open, including the 2025 impairment and brand separation, are listed once in section 31 with their effect on this chapter.

SECTION 8

Omnicom and IPG: a wider

perimeter

Scale, commerce and identity now share a parent. Joining them is still a test.

IN SHORT

Flywheel and IPG are closed. Hearts United joined Hearts & Science with Mediahub, not UM. One vendor case reports a recognition rate, not sales.

FROM FIGURE 3 · WHO OWNS WHAT, SINCE WHEN

Omnicom and IPG on the dated record

Omnicom and IPG: every entity on the dated record with its type, relation to the parent, transaction status, effective date, evidence date, verification, dated events, capability columns (analysis), deal card and the passages it rests on

ENTITY	TYPE	RELATION TO THE PARENT	TRANSACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNERSHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Omnicom	Parent	Parent	Parent, the owner of the rows below	none in the record held	none drawn	the parent itself	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s03/P1-0022.
Omnicom Media	Operating group	Part of the group	In place on the record 28 Aug 2026	none in the record held	28 Aug 2026	Issuer's own description; no deal record	Reorganised 2019	Activation	Not a deal: a reorganisation → gate reached: announcement Basis passages: s03/P1-0022, s04/P1-0029.
Interpublic (IPG), a lineage inside Omnicom from 26 November 2025	Business	Owned	Completed 26 Nov 2025	26 Nov 2025	26 Nov 2025	Ownership verified by this review: completion record held	Completed 26 Nov 2025	Data / identity, Activation	A parent-level combination: scale; IPG's lineage ran from first-party data work to applications and activation → gate reached: named live use Basis passages: s08/P1-0064, s08/P1-0065.
Flywheel	Business	Owned	Completed 2 Jan 2024	2 Jan 2024	2 Jan 2024	Ownership verified by this review: completion record held	Completed 2 Jan 2024	Commerce	A commerce business → gate reached: announcement Basis passages: s08/P1-0065.

ENTITY	TYPE	RELATION TO THE PARENT	TRANS ACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNER SHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Acxiom (IPG lineage)	Business	Owned — Owned on the record (an Acxiom release in Omnicom's newsroom, July 2026, in the IPG lineage); IPG's 2018 Acxiom purchase was carried forward unread, so ownership is not verified by this review. Acxiom and LiveRamp are different companies on different paths	Named in the record Jul 2026	none in the record held	Jul 2026	Owned on the record; not verified by this review: the record is unread	Named use reported Jul 2026	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s08/P1-0065, s08/P1-0067, s31/P4-0015.
Hearts United (Hearts & Science with Mediahub, 28 August 2026)	Agency brand	Part of the group — A joining of two brands, not a purchase; UM was not the merger partner	Renamed 28 Aug 2026	28 Aug 2026	28 Aug 2026	Issuer's own description; no deal record	Renamed 28 Aug 2026	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s08/P1-0067, s03/P1-0022.

ENTITY	TYPE	RELATION TO THE PARENT	TRANS ACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNER SHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
OMD, PHD, UM, Initiative (siblings named in the Hearts United release)	Agency brand	Part of the group — Siblings named in the Hearts United release; no deal record applies	Named in the record 28 Aug 2026	none in the record held	28 Aug 2026	Issuer's own description; no deal record	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s03/P1-0022, s08/P1-0067.
Omni (the group's shared platform)	Product	Part of the group — Named as the group's shared platform in the Flywheel closing release; its 2018 launch and the January 2026 release were carried forward unread	Named in the record 2 Jan 2024	none in the record held	2 Jan 2024	Owned on the record; not verified by this review: the record is unread	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s08/P1-0065, s08/P1-0066, s31/P4-0015.

A capability mark means the paper's category words for that asset fall in this column. It is not a measure of capability quality, and a deal count is not a capability score. **Gap:** Omnicom and IPG: Annalect, Credera, IPG's 2018 Acxiom purchase, Omni's 2018 launch, the January 2026 Omni release, Huge and R/GA, the 2024 impairment: carried forward unread. **Open the Omnicom and IPG tab of Figure 3.**

Prior gap. Omnicom filled its gaps in this order: shared analytics, then consulting access to client systems, then commerce data and operations, then scale. IPG began with first-party data work, the data a company collects from its own customers. It moved to applications and activation, and then into Omnicom. The two histories stay apart until 26 November 2025. After that date IPG is a lineage inside Omnicom. It is not a peer to add to a total. • **ANALYSIS**

Asset. Omnicom closed Flywheel on 2 January 2024. The closing release proposed a link between commerce data and Omni, the group's shared platform ^[4]. Omnicom completed its purchase of Interpublic on 26 November 2025 ^[5].

Both are closing records, not plans. One boundary from that record is worth a caution. Acxiom and LiveRamp are different companies on different paths. Nothing in the audited records joins them.^{[4][5]} · **FACT** A025

Intended link. The Flywheel closing release stated the design: link commerce data with Omni ^[4]. The original edition reported that Omni launched in 2018 around partners the group did not own and that a January 2026 release named Acxiom and Flywheel inside the platform. That shift, from partner-neutral to owned, is the chapter's central point, and it is kept here as analysis of the original edition's account.^[4] · **ANALYSIS**

Use on record. On 28 August 2026 Omnicom launched Hearts United by combining Hearts & Science with Mediahub. The release listed UM, PHD, OMD and Initiative as siblings. UM was not the merger partner ^[36]. This is a joining of two brands, not a purchase. In July 2026 an Acxiom release in Omnicom's newsroom quoted a Guitar Center executive: recognised website visitors rose from 23% to 35% after an Acxiom link into its Salesforce stack ^[7]. That is 12 percentage points, or roughly half again in relative terms. It is a vendor-published before-and-after with no disclosed control group. It supports a named use. It does not show added sales caused by the wider group.^{[36][7]} · **FACT** A026 A027 A028 A029

Client effect. One reported recognition rate, and nothing on added sales or profit. · **ANALYSIS** A029

Group economics. The second-quarter 2026 release, dated 28 July 2026, reports Core Operations revenue of \$5,995 million and core adjusted EBITA of \$1,068.8 million. EBITA is earnings before interest, tax and amortisation. The core base leaves out sold and held-for-sale businesses. Its 2025 comparator combines Omnicom and IPG ^[45]. It is a chosen perimeter, not a statutory series. The revised edition printed a bridge from total revenue to core revenue. In it, the excluded amount was just the gap between the other two numbers. This edition does not repeat it.^[45] · **FACT**

Countercase. In 2019 Omnicom's media group moved specialist staff out of a central desk and into OMD, PHD and Hearts & Science ^[35]. That is integration running the other way. The original edition held more countercases: IPG's sale of Huge and R/GA before the close, a 2024 goodwill impairment on IPG's digital specialists, a merger with Publicis abandoned in 2014, and US regulatory conditions on the merger. They would make the chapter's countercase larger, not smaller.^[35] · **ANALYSIS** A006

Rival reading. Once a skill becomes common work, placing it near the client can make more sense than selling access to a distant team. A margin on a chosen perimeter cannot say how much came from breadth, how much from scale and how much from cost removal. Cutting duplicate cost and gaining bargaining power can lift a margin with no client gain at all. Selling weak units makes the survivors look better while the sold units take the losses. A bigger group makes a shared platform easier to sell to the same clients. · **ANALYSIS**

What a buyer can test now. Test breadth at the seams. Can commerce stock data change a media plan? Can a creative team challenge the targeting logic? Can a proof team disagree with a sales team? The group thesis gains strength only when those tests work at full cost across repeated client scopes. · **ANALYSIS**

WHAT THE RECORD PROVES, AND DOES NOT

The record proves four things. Omnicom completed two purchases on known dates. It joined two brands on 28 August 2026, and the partner was Mediahub. One vendor reported a recognition result. And one quarter shows a margin on a chosen perimeter. The record does not prove three things. It does not show that commerce data changes media plans for any client. It does not show that the merger created value a client could not buy elsewhere. It does not give the group's principal share. Records this update did not re-open, including those disposals, the impairment and the regulatory conditions, are listed once in section 31 with their effect on this chapter.

SECTION 9

Dentsu: skill is not the same as group value

Local strengths and global costs need separate tests.

IN SHORT

Aegis and Merkle are closed records. The FY2025 goodwill impairment of ¥396.1 billion is real contrary evidence. It cannot be assigned to one deal.

FROM FIGURE 3 · WHO OWNS WHAT, SINCE WHEN

Dentsu on the dated record

Dentsu: every entity on the dated record with its type, relation to the parent, transaction status, effective date, evidence date, verification, dated events, capability columns (analysis), deal card and the passages it rests on

ENTITY	TYPE	RELATION TO THE PARENT	TRANSACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNERSHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Dentsu	Parent	Parent	Parent · Impaired not booked FY2025	none in the record held	none drawn	the parent itself	Impairment booked FY2025	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s09/P1-0075.
Aegis (Dentsu Aegis Network from 26 March 2013)	Business	Owned — Agreement announced July 2012; completed 26 March 2013, when Dentsu Aegis Network was established	Completed 26 Mar 2013	26 Mar 2013	26 Mar 2013	Ownership verified by this review: completion record held	Agreement announced Jul 2012; Completed 26 Mar 2013	Activation	Reach outside Japan and digital skill in those markets → gate reached: announcement Basis passages: s09/P1-0075, s04/P1-0025.
Merkle (100% from 15 April 2020)	Business	Owned — 100% ownership confirmed 15 April 2020; payment for the remaining stake due in the third quarter of 2021; ownership, payment and reporting dates differ	Completed 15 Apr 2020	15 Apr 2020	15 Apr 2020	Ownership verified by this review: completion record held	Completed 15 Apr 2020	Data / identity	Customer data and CRM: customer systems and services → gate reached: announcement Basis passages: s09/P1-0075.

A capability mark means the paper's category words for that asset fall in this column. It is not a measure of capability quality, and a deal count is not a capability score. **Gap:** Dentsu: the 2016 Merkle rationale, the 2023 structure, the 2026 audience product and the 2026 decisions: carried forward unread. **Open the Dentsu tab of Figure 3.**

Prior gap. Dentsu's gaps were reach outside Japan, digital skill in those markets, and then customer data and CRM. Aegis answered the first two. Merkle answered the third. Both are closed records. The question the record leaves open is whether the group could run those assets at a global cost base. · **ANALYSIS**

Asset. The Aegis agreement was announced in July 2012 and the deal completed on 26 March 2013, when Dentsu Aegis Network was established ^[29]. On 15 April 2020 Dentsu confirmed 100% ownership of Merkle, with payment for the remaining stake still due in the third quarter of 2021 ^[30]. Ownership, payment and reporting dates differ, which is why the entity map needs separate fields for each. ^{[29][30]} · **FACT**

Intended link. The case was to link demand with customer experience. Aegis brought media reach. Merkle brought customer systems and services. The original edition added more: a 2016 Merkle rationale, a 2023 one-management-team structure and a 2026 audience product. · **ANALYSIS**

Use on record. No Tier 1 or Tier 2 record in this review documents a named link between Merkle's work and the group's media work. That is a gap in the review, not a finding of absence. · **ANALYSIS**

Client effect. Not shown in any source reviewed. · **ANALYSIS** A048

Group economics and countercase. The FY2025 financial report records goodwill impairment of ¥396,074 million. Of that, ¥299,656 million relates to the Americas and ¥96,418 million to EMEA; the two regional figures sum exactly to the total ^[31]. The charge is revised value against book amounts, not that year's cash. From regional totals it cannot be assigned wholly to Merkle or to any single deal. ^[31] · **FACT** A030 A031 A050

The original edition leaned on a regional reading: Japan grew in the first half of 2026 while the Americas, EMEA and APAC outside Japan fell. It also reported a February 2026 decision not to sell the international business and a change of chief executive in March 2026. If those records hold, they show a group choosing to run the acquired assets rather than sell them, which sharpens the cost-base question without settling it. · **ANALYSIS**

Rival reading. The charge is real contrary evidence to a simple deal-success story. It is not proof that all the acquired skills failed. Strong acquired skills run at an unaffordable global cost base fit the record as well as a failed strategy does. Japan's model may fund a focused recovery. Both can be true in different markets. · **ANALYSIS**

What a buyer can test now. Test the proposed local team, its continuity and its full cost. Ask which team owns the brief, the customer data model, the creative work,

the buying accounts and the proof of impact in the market concerned. The group's size is a weak substitute for that work. · **ANALYSIS**

WHAT THE RECORD PROVES, AND DOES NOT

The record proves three things. Dentsu bought international reach in 2013. It owned all of Merkle from April 2020, with payment in 2021. And it booked a ¥396.1 billion goodwill impairment for FY2025, split between the Americas and EMEA. The record does not prove two things. It does not show which deal lost value. It does not show any client's result. Records this update did not re-open, including the regional growth split and the 2026 decisions, are listed once in section 31 with their effect on this chapter.

SECTION 10

Havas: a case for access as well as ownership

Partnerships may offer reach. They can also leave hard handoffs.

IN SHORT

A listing in December 2024, a joint venture in September 2025 and a 12.9% adjusted margin for 2025 are the audited record. The control structure at the cutoff is not.

FROM FIGURE 3 · WHO OWNS WHAT, SINCE WHEN

Havas on the dated record

Havas: every entity on the dated record with its type, relation to the parent, transaction status, effective date, evidence date, verification, dated events, capability columns (analysis), deal card and the passages it rests on

ENTITY	TYPE	RELATION TO THE PARENT	TRANSACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNERSHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Havas (listed on Euronext Amsterdam 16 December 2024)	Parent	Parent — Control structure at the cutoff not in the audited record	Parent · Listed 16 Dec 2024	none in the record held	16 Dec 2024	the parent itself	Listed 16 Dec 2024	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s10/P1-0085.
Havas Media Network	Operating group	Part of the group	In place on the record 29 Sep 2025	none in the record held	29 Sep 2025	Issuer's own description; no deal record	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s10/P1-0086.
Horizon Global (joint venture with Horizon Media, 50/50)	Joint venture	Joint venture — Announced by both parents on 29 September 2025; described by Horizon as 50/50; the parents did not merge; no formation record beyond the announcements is held. Shared with the Horizon tab	Announced 29 Sep 2025	none in the record held	29 Sep 2025	Issuer's own description; no deal record	Announced 29 Sep 2025	Activation (announced or proposed)	Not a purchase: a commercial joint venture, 50/50; the parents did not merge → gate reached: announcement Basis passages: s03/P1-0019, s03/P1-0020, s10/P1-0086.

A capability mark means the paper's category words for that asset fall in this column. It is not a measure of capability quality, and a deal count is not a capability score. **Gap:** Havas: control structure at the cutoff, FullSIX, Havas CX,

Uncommon, Channel Bakers and the 2026 deals: carried forward unread; no Havas purchase is listed. **Open the Havas tab of Figure 3.**

Prior gap. Havas wanted depth in customer experience and digital work, a health specialism, creative work that stands out, and scale without owning every layer of technology. Its route mixes a few purchases with partnerships. The test is whether a mixed system can deliver better work with clear client rights. The client must also know which firm stands behind each part. • **ANALYSIS**

Asset and structure. Havas listed on Euronext Amsterdam on 16 December 2024, after its spin-off from Vivendi ^[32]. A separate listing does not by itself show who controls the company. The original edition reported family-linked holdings as of September 2026, and purchase records for FullSIX, Havas CX, Uncommon, Channel Bakers and the 2026 deals. Those records are carried forward unread; section 31 lists them with their effect on this chapter. ^[32] • **FACT**

Intended link. On 29 September 2025 Havas and Horizon Media announced Horizon Global, a commercial joint venture ^[33]. Horizon's own December 2025 release describes it as a 50/50 partnership with Havas Media Network ^[34]. The parents did not merge. The combined billings quoted in coverage are not Havas revenue. ^{[33][34]} • **FACT** A032 H018

Use on record. No Tier 1 or Tier 2 record in this review documents a named client use of the Havas operating system or of the joint venture. The original edition named released tools. • **ANALYSIS**

Client effect. Not shown in any source reviewed. • **ANALYSIS** A048

Group economics. For 2025 Havas reported net revenue of €2,783 million, adjusted EBIT of €358 million and organic growth of 3.1%; organic growth is growth excluding deals and currency, as the company defines it. The margin is 12.9%, and the ratio recomputes to 12.9% at one decimal ^[48]. That is a full-year, company-defined adjusted measure. It cannot be ranked against a peer's half-year or statutory margin. ^[48] • **FACT** A033 A034 A035

Countercase. The original edition reported two things from the listing prospectus: a large budget loss at one health client, and a €400 million programme for 2024 to 2027 of which about three-fifths was operating expense. Both matter to a claim of breadth, and both would count against it. No impairment record for Havas was audited. • **ANALYSIS**

Rival reading. A joint venture is a route to scale without ownership. It also leaves two firms behind one service. Partnerships can keep the brand dependent on outside systems and leave the joining work with people. Open does not mean low cost or easy exit. A workable shared process need not own all the data or tools below it. But the contract must name the party that fixes a broken handoff. • **ANALYSIS**

What a buyer can test now. Test the joint service under load. Name the firm that repairs a failed handoff. Demand a clear route to move the work if one partner leaves. Ask which data can move to another partner and which results can be tested by an outside party. · **ANALYSIS**

WHAT THE RECORD PROVES, AND DOES NOT

The record proves three things. Havas listed on 16 December 2024. It announced a joint venture on 29 September 2025, which Horizon describes as 50/50. And it earned a 12.9% adjusted EBIT margin on €2,783 million of 2025 net revenue. The record does not prove two things. It does not show who controls the listed company at the cutoff. It does not show any client's result from the mixed model.

SECTION 11

Stagwell and the smaller-group test

Specialist skill can stay distinct while tools are shared. A smaller size proves nothing by itself.

IN SHORT

The 2021 combination, a \$16 million committed-revenue label and a 3.3% organic bridge on a changed method are the record. S4 Capital is the counterweight.

FROM FIGURE 3 · WHO OWNS WHAT, SINCE WHEN

Stagwell on the dated record

Stagwell: every entity on the dated record with its type, relation to the parent, transaction status, effective date, evidence date, verification, dated events, capability columns (analysis), deal card and the passages it rests on

ENTITY	TYPE	RELATION TO THE PARENT	TRANSACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNER SHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
Stagwell (combined with MDC 2 August 2021)	Parent	Parent	Parent · Recorded Q3 2025	2 Aug 2021	2 Aug 2021	the parent itself	Completed 2 Aug 2021; Recorded Q3 2025	Creative / production, Activation, Measurement	Creative, research and digital firms in the roster → gate reached: announcement Basis passages: s11/P1-0095.
Enterprise-software revenue label (\$16 million committed, H1 2026)	Reported figure	Reported figure — A revenue-category label in the Q2 2026 earnings presentation (filed 30 July 2026): a commercial signal, not recognition and not a product entity. No product entity is identified in the record held; the three 2026 product releases were carried forward unread	Reported in a filing of 30 Jul 2026	none in the record held	30 Jul 2026	Not an entity: no product entity in the record held	Recorded 30 Jul 2026	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s11/P1-0097, s11/P1-0096, s31/P4-0015.

ENTITY	TYPE	RELATION TO THE PARENT	TRANS ACTION STATUS	EFFECTIVE DATE	EVIDENCE DATE	OWNER SHIP VERIFIED BY THIS REVIEW ?	DATED EVENTS	CAPABILITY COLUMNS (ANALYSIS)	DEAL CARD; BASIS PASSAGES (PARAGRAPHS IN THIS PAPER)
S4 Capital (counterweight, not part of Stagwell)	Parent	Counterweight — A separate listed digital network used as the smaller-group test	Not part of the group · results Mar 2026	none in the record held	Mar 2026	no ownership claim	no dated event drawn	none mapped (not assessed in this review; not zero)	No deal card: no deal event for this entity in the record held. Basis passages: s11/P1-0101, s11/P1-0102.

A capability mark means the paper's category words for that asset fall in this column. It is not a measure of capability quality, and a deal count is not a capability score. **Gap:** Stagwell: 72andSunny, Anomaly, the Harris assets, Brand New Galaxy, ADK GLOBAL and the three 2026 product releases: carried forward unread. **Open the Stagwell tab of Figure 3.**

Prior gap. Stagwell is a test of whether a smaller group can connect specialist firms, software and media access without erasing the agencies that brought it clients. Its gaps were scarce creative talent, research, software products, geographic reach and preferred supply. The group did not begin with one finished system. · **ANALYSIS**

Asset. The Stagwell and MDC combination was announced as complete on 2 August 2021. Creative, research and digital firms appear in the roster ^[6]. The original edition also recorded 72andSunny, Anomaly, the Harris assets, Brand New Galaxy and ADK GLOBAL. ^[6] · **FACT** A036

Intended link. The original edition described three 2026 products. In January, a workflow layer across existing tools. In June, a media operating system with more than twenty agents and human decision gates. In July, a marketplace of preferred supply. Together they are the clearest statement by any group that workflow and preferred access are meant to feed each other. Orchestration means joining tasks, data, tools and rights across one process. The limit this paper places on the whole orchestration thesis is here: a shared workflow can steer buying toward the supply the group prefers. Those three releases are carried forward unread and listed in section 31; the argument stands here without them. · **ANALYSIS**

Use on record. The second-quarter 2026 earnings presentation, filed on 30 July 2026, labels \$16 million as committed enterprise-software revenue for the first half ^[27]. That is a commercial signal. The slide alone does not show recognition in the period. It does not separate subscriptions from services. It does not show that agent systems are the group's main profit source. ^[27] · **ATTRIBUTED** A037

Client effect. The original edition held one cost claim from a named client; the audited record holds no client result. · **ANALYSIS**

Group economics. The same presentation shows first-half net revenue of \$1,216.2 million against \$1,162.3 million a year earlier. The organic bridge is \$38.9 million, which is 3.3% of the prior-year base ^[27]. The company headline rounds that to 3%; the original edition printed it with an extra decimal it never had. Stagwell changed its organic-growth method from the third quarter of 2025. Acquired businesses now enter organic growth after twelve months, and the deal adjustment uses current-period acquired revenue rather than prior-period revenue ^[28]. The 3.3% is on the new method. It must not be spliced onto any earlier series without a constant-method bridge. No such bridge exists in the audited records. ^{[27][28]} · **FACT** A038 A039 A040 A072

Rival reading. A set of specialist firms sharing tools is a services business with a software label. Software sales may be a way to sell services. A marketplace of preferred supply may make supply economics stronger, not replace them. The gap between gross and net revenue is not a disclosed principal margin. · **ANALYSIS**

The smaller-group test: S4 Capital

S4 Capital's audited 2025 results, presented in March 2026, show net revenue of £673.0 million, down 10.8% as reported and 8.4% like for like on the company's definition. Operational EBITDA (earnings before interest, tax, depreciation and amortisation) was £81.2 million and year-end net debt £86.9 million ^[44]. The decline does not isolate an AI effect. The revised edition said the margin rose and net debt fell. Those directions rest on prior-year values not in the audited passage, so this edition states the 2025 figures only. ^[44] · **FACT** A041 A042 A043 A044

What a buyer can test now. S4 is the counterweight to any claim that one digital brand protects demand. A digital model can improve cost and cash while demand weakens. Consolidation and technology coexist with client caution. So judge the actual team and the joins, not the size of the parent. Ask whether the default supply choices inside a buying system improve the client's result, and whether an outside party can challenge them. Keep three things apart: delivery the group owns, work done with affiliates, and access to software. A smaller group can offer all three under one name. · **ANALYSIS**

WHAT THE RECORD PROVES, AND DOES NOT

The record proves four things. Stagwell and MDC completed their combination on 2 August 2021. Stagwell labelled \$16 million as committed enterprise-software revenue for the first half of 2026. Its organic growth was 3.3%, on a method changed from the third quarter of 2025. And S4 Capital's 2025 net revenue fell 10.8% as reported. The record does not prove two things. It does not show that agent products are a profit source. It does not show that any client's total cost fell.

The alternatives: independents, specialists, platforms and the in-house boundary

The alternative to a group is not one kind of firm. Compare the entry point to the client, not the label.

IN SHORT

In-house teams and outside agencies coexist in every sample held here. Horizon is an independent agency with a proposed way of working. Platforms automate inside their own inventory.

The rival to a large group is not one kind of firm. The set includes consulting firms, listed digital networks, groups backed by private equity, founder-led firms, teams embedded in the client, and the media platforms themselves. Each enters the client through a different need: media buying, enterprise systems, creative production, commerce or a skill the brand wants inside. Each can then grow into the next task. This paper compares where each one gets its grip on the client, not their revenues. A consulting firm's total revenue, an agency's fee revenue, media billings and software commitments do not measure the same thing. This review holds no audited revenue for any consulting firm's marketing arm. • **ANALYSIS**

Not all independents are open, and not all groups are closed. Independence from a listed group can mean private-equity ownership with its own return targets, founder control with minority investors, or an agency that sells embedded teams and owned software. None of those makes advice more impartial or exit easier by itself. The test is the same as for a group: rights the client can exercise, a proof route the seller does not control, and a working export. • **ANALYSIS**

The in-house boundary

In-house work moves the boundary. It does not end the need for an agency in any sample held here. The ANA's 2023 member survey had 162 respondents, with fieldwork in February and March 2023. It reported 82% with an in-house agency and 92% still using outside agencies ^[9]. The WFA and Observatory International study of 45 companies, published on 20 December 2023, reported 66% with an in-house agency and all respondents still working with outside agencies ^[10]. These are incidence rates in member samples, not shares of spend. The two populations must not be pooled. ^{[9][10]} • **ATTRIBUTED** A064 A065 A066 A067 A068

The results support coexistence. They do not show that agency work or spend stayed level. A brand can bring one task inside and buy more of another outside. A small internal team can own a key choice while a partner does most of the work. Dependence falls only when the brand keeps the brief, the account ownership, the test plan and the budget authority, and can run an exit. Fixed cost rises with headcount whoever owns the decision. · **ANALYSIS**

Compare tasks, not labels. Who owns the brief, the account, the test plan and the final budget? Which skills are needed every week, and which twice a year? Count the full cost of staff, tools, training, cover, rework and the links between teams. Map dependence by task and by exit time, the time it would take to move each task to another supplier. Do not map it by the presence of an internal team. · **ANALYSIS**

Horizon, as an independent

Horizon Media is the independent agency this paper follows. Its releases describe a way of working, not a deal. Horizon Media Holdings is a private parent ^[34]. On 18 December 2025 it launched HorizonOS. The release described human specialists, client pilots and a partner marketplace. It also disclosed the 50/50 Horizon Global partnership with Havas Media Network ^[34]. On 18 June 2026 it announced an agentic layer for media decisions, with APIs, MCP connections and agentic buying. An API is a software interface; MCP is a standard for connecting AI tools; 'agentic' means software agents act within the system. The one named client, SharkNinja, was using audience APIs with retail partners. That is audience collaboration, not autonomous buying ^[13]. On 1 October 2026 it named a third cohort of 11 partners, chosen from more than 600 submissions. More than 50 partners had been through the programme. The next step was to scope pilots with clients ^[14] ^[34] ^[13] ^[14]. · **ATTRIBUTED** H001 H002 H003 H004 H005 H006 H008 H009 H010 H011 H017 H018

Three limits apply, and section 22 works through them. A marketplace is not portability. The promise of an ecosystem with no vendor lock-in is a sales claim until a contract and a migration drill back it. A count of partners is not a count of deployments. An announcement is not proven autonomy, and the SharkNinja example does not prove autonomous buying. And no fee mix, principal margin or client economics for Horizon appears anywhere in this paper, because no release reviewed states them. Horizon is a proposed way of working. It sits closest to the agency-led path. It is not proof that independents have won. · **ANALYSIS** H012 H020 H006

Platform substitutes

Platforms automate the work inside their own inventory. Google's help page describes Performance Max as a campaign type that spans Google Ads inventory and bids toward goals the advertiser sets. The page is a living document with no version history ^[15]. Amazon's page dated 21 March 2025 had a title that put Brand+ in open beta. Its body, when the audit read it, said both Performance+ and

Brand+ were available globally. By the cutoff the title in search results had changed again. The page cannot date the later stage on its own, and this paper does not claim the exact stage from it ^[16]. On the five gates, both sit at the second, documented implementation, as described by the seller. ^[15]

^[16] · **ATTRIBUTED** A057

The substitution is real and bounded. A platform removes manual buying work inside its own silo. It gives no outside view across platforms. It keeps its own model state, the learned settings behind its bidding, under its own terms. An agency can compete with that automation. It can also run it for the client. So platform use is an input to an agency as well as a rival to it. Neither a platform's report nor an agency's claim settles whether the sales were new, moved from another channel, or bought with discounts. · **ANALYSIS** A052

For the CMO the evidence supports a choice of operating model, not a choice between old and new logos. Use a group where shared talent, local access and cross-discipline work solve a real coordination problem. Use a consultancy where implementation and enterprise change dominate. Use embedded teams where closeness to the brief and repeat production matter. Use software where the brand can own the process and govern the outputs. These are decision rules, not rankings. Part III turns them into three paths. · **ANALYSIS**

SECTION 13

What the numbers can and cannot show

Keep each figure's unit, period, currency and cost base attached. No shared axis.

IN SHORT

A dated bridge table replaces the peer chart. The LiveRamp values are explained line by line. Survey figures keep their denominators.

Numbers in this field mislead in four ways. Gross revenue is compared with net revenue. A quarter is compared with a half-year. An organic rate on one definition is compared with a like-for-like rate on another. And a rounded headline gains a decimal it never had. This section prints each observation on its own row with its own definition. It says which can be recomputed from disclosed figures and which can only be quoted. · **ANALYSIS**

Growth observations, one row each

Dated growth observations from Tier 1 and Tier 2 records; no shared axis, no ranking ^[43] ^[46] ^[48] ^[44] ^[47]
^[27] ^[28]

COMPANY FACT	MEASURE AS LABELLED BY THE COMPANY ATTRIBUTED	PERIOD FACT	CURRENCY FACT	RATE AS REPORTED ATTRIBUTED	METHOD NOTE ANALYSIS	RECOMPUTA BLE HERE? ANALYSIS	SOURCE AND LOCATOR FACT
Publicis	Net revenue, reported change	1 January to 31 December 2025	EUR	4.2%	Net revenue €14,547m against €13,965m; organic growth is a separate company figure not recomputed here	Yes: (14,547 / 13,965 – 1) = 4.2% at one decimal	[43] URD income statement, printed p. 260
Publicis	Organic net revenue growth, second quarter	1 April to 30 June 2026	EUR	4.8%	Company- defined organic rate; full-year guidance on the same page is prospective	No: quoted as reported	[46] results page, 16 July 2026
Havas	Organic growth of net revenue	1 January to 31 December 2025	EUR	3.1%	Company- defined; full year	No: quoted as reported	[48] key results, 17 February 2026
S4 Capital	Net revenue, reported change	1 January to 31 December 2025	GBP	–10.8%	Reported; the prior-year input is not in the audited passage, so no recomputation is printed	No: quoted as reported	[44] results RNS, March 2026, financial summary
S4 Capital	Net revenue, like-for-like change	1 January to 31 December 2025	GBP	–8.4%	Company- defined like-for- like; scope and currency adjustments not reconstructed	No: quoted as reported	[44] results RNS, March 2026
WPP	Revenue less pass-through costs, like-for- like change	1 January to 30 June 2026	GBP	–4.7%	Company- defined like-for- like; base is revenue less pass-through costs, not net revenue	No: quoted as reported	[47] interim results, 6 August 2026, key figures

COMPANY FACT	MEASURE AS LABELLED BY THE COMPANY ATTRIBUTED	PERIOD FACT	CURRENCY FACT	RATE AS REPORTED ATTRIBUTED	METHOD NOTE ANALYSIS	RECOMPUTA BLE HERE? ANALYSIS	SOURCE AND LOCATOR FACT
Stagwell	Organic net revenue growth, first half	1 January to 30 June 2026	USD	3.3% from the detailed bridge; headline rounded to 3%	Method changed from the third quarter of 2025: acquired businesses enter after twelve months with a current-period adjustment; not comparable with earlier series	Yes: 38,936 / 1,162,316 = 3.3%	[27] slide 14 bridge; [28] MD&A p. 30

Five rows from the original peer chart are missing on purpose. They are Publicis first-half 2026 organic growth, Omnicom second-quarter 2026 core organic growth, Dentsu first-half 2026 organic growth, Havas first-half 2026 organic growth and S4 Capital first-half 2026 like-for-like change. Each rests on a document this update did not re-open; section 31 lists them with the primary document that would settle each. Even when they return, they must not share an axis. A quarter is not a half. A chosen core perimeter is not a statutory group. A changed organic method is not an unchanged one. · **ANALYSIS** A049

Margins, one decimal, with their definitions

Margin observations recomputed from rounded disclosed inputs; not a league table^{[38] [47] [45] [48] [44]}

Column classifications: · **FACT** · **ANALYSIS**

COMPANY AND PERIOD FACT	NUMERATOR / DENOMINATOR FACT	RECOMPUTED AT ONE DECIMAL FACT	WHY IT CANNOT BE RANKED AGAINST THE OTHERS ANALYSIS
Publicis, first half 2026	Operating-margin subtotal €1,261m / net revenue €7,229m	17.4% after €7m of LiveRamp costs; 17.5% before them ^[38]	A company subtotal, not statutory operating income; half-year
WPP, first half 2026	Headline operating profit £398m / revenue less pass-through £4,745m	8.4% ^[47]	Headline profit excludes items; net base
WPP, first half 2026	Reported operating profit £261m / revenue £6,373m	4.1% ^[47]	Different numerator and denominator from the row above
Omnicom, second quarter 2026	Core adjusted EBITA \$1,068.8m / core revenue \$5,995m	17.8% ^[45]	Selected core perimeter; one quarter; 2025 base combines Omnicom and IPG
Havas, full year 2025	Adjusted EBIT €358m / net revenue €2,783m	12.9% ^[48]	Full year; adjusted, company-defined

COMPANY AND PERIOD FACT	NUMERATOR / DENOMINATOR FACT	RECOMPUTED AT ONE DECIMAL FACT	WHY IT CANNOT BE RANKED AGAINST THE OTHERS ANALYSIS
S4 Capital, full year 2025	Operational EBITDA £81.2m / net revenue £673.0m	12.1% ^[44]	EBITDA, not operating profit; company-defined

No two of these six ratios share both the top line and the bottom line. The inputs are rounded to the nearest million. A fourth decimal would be false precision. One decimal is all the inputs allow. No figure in this paper is converted to another currency or scaled to a full year. · **ANALYSIS** A021

LiveRamp: which number is which

Press reports of the proposed LiveRamp deal quote \$2.5 billion and about \$2.2 billion for the same transaction. Both are right. They measure different things. The Publicis release of 17 May 2026 states an equity value of \$2.546 billion, net cash of \$379 million and an enterprise value of \$2.167 billion ^[17]. Equity value is what shareholders receive: the per-share price times the fully diluted share count. Enterprise value is what the buyer pays net of the cash it acquires: \$2.546 billion less \$0.379 billion equals \$2.167 billion, which rounds to \$2.2 billion. Publicis’s 17 May 2026 announcement states an offer price of \$38.50 per share, an equity value of \$2.546 billion and an enterprise value of \$2.167 billion; completion remained conditional in its 16 September update ^[18]. The price was confirmed in the announcement by an external review on 6 October 2026; this edition did not re-open the source, and its stored passage carries the two values but not the price. The net-cash figure is as of a date stated in the release and will differ at closing. Neither number is a completed purchase price. The two must never be added or swapped. ^{[17][18]} · **FACT** A013 A014 V3-009

Survey denominators

The ANA reported in March 2026 that 58% of respondents had used principal media in the past year, against 47% in its prior wave. The sample was 114 qualified client-side members, with fieldwork from October 2025 to January 2026 ^[8]. The gap is 11 percentage points between two member samples that were not matched. The prior wave’s sample size and the base for the 58% are not in the audited passage. No weighting is stated. The figure is a rate of use, not a share of spend. A relative change should not be computed from it. The 2023 in-house figures are rates of use too: 82% and 92% of 162 ANA respondents, and 66% of 45 companies in the WFA and Observatory study. They must not be pooled ^{[9][10]}. The WFA’s 2026 study covered 69 multinationals, 71% of whose respondents work in global marketing procurement. It reported that 63% expected more performance-based fees, and that such fees are still a small share in most disciplines ^[11]. ^{[8][9][10][11]} · **ATTRIBUTED** A061 A062 A063 A064 A065 A066 A067 A068 A069 A070 A071

What growth cannot say

Growth can come from price, wins, scope, currency or demand. Profit can rise through staff cuts. Cash can rise through client payment timing. None of these, on its own, shows that AI or one deal caused a better client result. The records in this section do not isolate a causal effect of AI or of any one deal. A strong test needs the same scope, quality and cost boundary on both sides. No such test has been published for any group in this review. · **ANALYSIS** A048

SECTION 14

Value created and value captured

A group can gain margin without a client gaining profit. A client can gain from a closed system.

IN SHORT

Nine mechanisms, who gains from each, whether the client can gain too, and the evidence that would show it.

Value created for a client and value captured by a seller are different quantities. Value is captured in fees, spreads, subscriptions, data services, switching costs or platform spend. A group can gain margin without a client gaining profit. A client can gain from a closed system. An open system can cost more to run. This section does not rank business models by virtue. It names each mechanism, who gains from it as volume rises, whether the client can gain too, and what evidence would show that the client did. · **ANALYSIS**


FIGURE 6 ILLUSTRATIVE

Value created and value captured

Nine ways value is captured as volume rises. On every row the client can gain only if a condition holds. The record holds no proof that it does.

Nine mechanisms: who gains as volume rises

From the section 14 table. Each chip sits under the party that gains as volume rises. Select one for the payer, the benefit condition, the conflict and the evidence.



Not a ranking. Chip position and size carry no amount. On every row the client can gain only if a condition holds; no client-side test for a named group account is on record.

Value created and value captured ILLUSTRATIVE Nine mechanisms from the section 14 table, drawn under the party that gains as volume rises: Brand, Agency or group, Provider or platform. Select one to see the payer, the benefit condition, the conflict and the evidence needed. Every client gain is conditional and none is measured here. Not a ranking. Status: illustrative (analysis). Units: none. Scope: the paper's own design. A060 A058 A059 V3-097 A061 A062 A037 A015 A028 A029 A057 V3-089 A069 A071 V3-094 Text equivalent: appendix "Figure text equivalents".

Mechanisms of value capture and the evidence that would show a client gain [12] [49] [50] [8] [27] [7] [17] [14] [15] [16] [11]

Column classifications: ANALYSIS ATTRIBUTED

MECHANISM ANALYSIS	WHO GAINS AS VOLUME RISES ANALYSIS	CAN THE CLIENT GAIN? ANALYSIS	EVIDENCE THAT WOULD SHOW IT ANALYSIS	EVIDENCE STATUS IN THIS REVIEW ATTRIBUTED
Labour or service fees	The agency, when billable scope or hours rise	Yes, when the work changes demand or decisions at a fee below the brand's own cost	Approved useful output per total hour; quality held; rework falling; fee against retained-team cost	Task-level productivity evidence exists outside advertising: a 2025 study of 5,172 support agents found a 15% average gain [12]. No matched agency study reviewed.
Spend-linked fees	The agency, when billable spend rises	Only if the extra spend is profitable at the margin	Marginal client profit per extra unit of spend, after returns and discounts	Not shown. Two disclosed CMO pay plans reward company revenue, margin and shareholder return, not spend [49] [50].
Principal media spreads	The reseller, before its service cost and inventory risk	Yes, if the bundled price for comparable inventory is below the agency route and the brand can assess the product	Disclosed role per transaction; a match on equal inventory, timing, reach, quality, data rights and payment terms; inventory risk actually borne	Use rose to 58% from 47% in the ANA member sample [8]. No spread is disclosed by any group reviewed.
Technology subscriptions and platform fees	The provider, from seats, usage and renewals	Yes, if the tool replaces a costlier process and exit charges are bounded	True usage; overlapping licences removed; exit cost priced; cost per approved asset	One group labels \$16 million as committed, not recognised, enterprise-software revenue [27].
Data services: identity, collaboration, audiences	The supplier, from permitted data work and access	Yes, if match quality and lawful use improve targeting or measurement and the data can move	Provenance; match rate; use rights; portability; added effect of matched audiences	One vendor case with a before-and-after recognition rate and no control group [7]. Neutrality promised for a proposed asset, not proven [17].
Switching costs: embedded workflow, learned state, integrations	The incumbent operator or vendor, through retention	Indirectly, if continuity lowers rework and the client accepted the cost knowingly	Exit drill cost, downtime and lost state documented and priced before signing	A scenario condition, not a checked right. One independent agency's lock-in promise is an unverified commercial claim [14].
Platform spend and in-platform automation	The platform, as media and paid use rise	Yes, on bounded tasks with reliable local signals	An outside cross-platform test showing added demand rather than substitution or harvesting	Capability documented by the sellers [15] [16]. No outside portfolio-level proof reviewed.
Outcome or performance fees	The partner, when the agreed metric clears the hurdle	Yes, if the metric is added profit over a matching window	Fee tracks added profit after substitution, discounting and brand effects; refunds and caps exercised	Expected to grow by 63% of 69 multinationals; a small share today in most disciplines [11].
Working capital and float	The intermediary when the client pays before suppliers are due; the client when it pays late	Either side, depending on dates	Payment dates, supplier due dates, advance funding, cancellation liabilities	No float income can be estimated from billings; no audited disclosure reviewed.

Three readings follow from the table. First, spreads, subscriptions, data services and switching costs may allow a seller's revenue to grow faster than its delivery labour; in this review's register they are also the mechanisms with no client-side test on record. This review does not compare their scalability or client returns. That is not an accusation. It is a measure of where the proof is missing. Second, the one mechanism with task-level evidence, labour saving, comes from customer support, not advertising. The study's authors note that the gain varied with skill [12]. It cannot be applied to agency cost or profit. Third, every row has a client-side test that a contract can require. None has been published for a named group account. [12] · ANALYSIS A060

The reading rule for the rest of the paper is this. For each mechanism, state who funds spend before proof arrives, who funds failed tests and who captures savings. A spending cap is not a liability ceiling. Spend awaiting proof is not a loss. A group that gains margin has not thereby cost a client anything. A client that gains from a closed system has not thereby lost control. The question for each mechanism is whether the brand can see it, test it, stop it and leave it. Decision rights explain who captures value; the quality of the work, which section 19 takes up, is what creates it. Part II puts eleven claims about these mechanisms under pressure. Part III turns the tests into three paths. · ANALYSIS

SECTION 15

Eleven claims put under pressure

For each claim: the strongest support in this register, the strongest contrary evidence, a rival explanation, the condition under which it holds, and the test that could change it.

IN SHORT

The public record settles none of the eleven claims. Claim (b) is graded as three claims: a revenue-model change, a delivery change and a client-value change. The thinnest support in this register sits under the revenue-model claim, the protocol claim and the cluster about autonomy, self-assessment and who holds the learning.

Part I set out the record. This part tests the claims people draw from it. The eleven below are the ones a CMO is likely to hear in pitches, trade press and board papers; claim (b) is graded in three parts, because a revenue model, a delivery model and client value are three different claims. For each, the table gives the strongest support in this edition's register, the strongest contrary evidence in it, a rival explanation, the condition under which the claim would hold, and the test that could change this paper's reading. Every row uses only the sources in this

edition's register. A blank is a gap in what this review inspected, not a finding of absence. · ANALYSIS

Eleven claims, five tests each. Support and contrary cells cite the register; the last three columns are this review's analysis. [15] [35] [20] [21] [8] [42] [39] [19] [27] [43] [37] [31] [47] [7] [36] [41] [9] [10] [12] [46] [38] [44] [22] [23] [14] [13] [16] [51] [17] [11] [49] [50]

Column classifications: · ANALYSIS · ATTRIBUTED

CLAIM ANALYSIS	STRONGEST SUPPORT ATTRIBUTED	STRONGEST CONTRARY ATTRIBUTED	RIVAL EXPLANATION ANALYSIS	CONDITION FOR THE CLAIM TO HOLD ANALYSIS	TEST THAT COULD CHANGE THIS READING ANALYSIS
(a) Buying scale is losing value.	Google documents automated bidding toward advertiser goals across its own inventory, which removes manual buying work [15]. Omnicom moved specialist programmatic staff out of a central desk into OMD, PHD and Hearts & Science in 2019, while central capabilities remained [35]. Two protocol efforts propose structured discovery, proposal and booking between buyer and seller agents [20] [21].	ANA reported on 19 March 2026 that 58% of 114 client-side members used principal media in the past year, against 47% in the prior wave [8]. GroupM's 2011 Xaxis launch already joined data and technology to trading scale; technology did not remove scale then [42].	Scale is being repriced, not removed: from unit price to access, payment terms, inventory risk and service. Part of it is captured as group margin rather than passed to the client.	Holds where inventory is auctioned inside one platform. Fails where inventory is negotiated: premium video, private marketplaces, retail media deals, sponsorships.	For the same inventory, period and scope, compare net terms and results across agency buying, principal buying and direct platform buying, with the commercial role disclosed per line.
(b1) Groups are changing their revenue model toward recurring technology fees.	Stagwell labelled \$16 million as committed enterprise-software revenue in its Q2 2026 presentation [27].	Income still comes from services and media. Publicis reported FY2025 net revenue of €14,547 million against revenue of €17,399 million, and the gap is pass-through cost, not a software line [43]. The Stagwell figure is committed, not recognised, and mixes products with services [27].	An investor story that supports a higher multiple; a services business with a software label; a systems-integration business that resells third-party tools.	Holds for a group only if a material, recurring technology fee line appears with disclosed retention in audited accounts.	An audited revenue mix showing software or platform fees as a material recurring line with client retention. Three more years of launches without that line would support the rival explanations.

CLAIM ANALYSIS	STRONGEST SUPPORT ATTRIBUTED	STRONGEST CONTRARY ATTRIBUTED	RIVAL EXPLANATION ANALYSIS	CONDITION FOR THE CLAIM TO HOLD ANALYSIS	TEST THAT COULD CHANGE THIS READING ANALYSIS
(b2) Groups are changing how they deliver: coordinating tools, data and decisions across services.	Publicis placed Chief Client Officers across disciplines in its December 2015 plan, so coordination predates generative AI [39]. Publicis describes bids triggered by stock and price signals, a named input-to-action link [19]. Omnicom combined two media brands into Hearts United in August 2026 [36].	Publicis booked a €1,440 million charge in FY2016, including €1,392 million at Publicis.Sapient, citing integration delays [37]. Dentsu's FY2025 goodwill impairment was ¥396,074 million [31]. WPP's H1 2026 revenue less pass-through costs fell 4.7% like for like [47]. No logged cross-service workflow on a named account is in the register.	Cross-selling the same services through one client officer; a cost reset dressed as a platform; a rename.	Holds for a group where repeated cross-service workflows, decision rights inside the client workflow and accountable hand-offs can be shown on named accounts. A retainer, project or managed-service fee can carry this; no software revenue is needed, and subscription revenue does not show it.	Trace one real brief through the proposed teams with every approval and transfer logged. Count shared steps, manual repairs and the right to reject a group product. Repeat on two accounts.
(b3) Those changes improve client value.	A Guitar Center executive reported recognised site visitors rising from 23% to 35% after an Acxiom and Salesforce integration [7].	That case is a vendor-published before-and-after with no disclosed control group [7]. No matched client result with full cost is in the register for any group.	Gains from better inputs or a favourable cycle; successes selected for publicity; a premium paid for service rather than for results.	Holds where quality, total cost or incremental business effect improves at equal scope after the full cost of the arrangement. Recurring revenue at the group shows nothing about this; a wide scope shows nothing about it either.	A matched-scope test with full cost and added effect, read by a party without a sales interest, on a named account.
(c) A wider service scope proves integrated delivery or client benefit.	A Guitar Center executive reported recognised site visitors rising from 23% to 35% after an Acxiom and Salesforce integration [7]. Omnicom combined Hearts & Science with Mediahub into Hearts United on 28 August 2026 [36].	The Guitar Center case is a vendor-published before-and-after with no disclosed control group [7]. Publicis' FY2016 charge cited delays integrating legacy digital assets [37]. Dentsu's regional impairments cannot be assigned to any single deal [31]. WPP sold 60% of Kantar in December 2019 [41].	Breadth serves cross-selling and investor positioning. Firms are bought for talent and client ties. A shared profit line is not a shared data model.	Holds at the seams: when stock data changes a media plan, when a creative team can challenge the targeting logic, when a proof team can disagree with a sales team, and when the client gains one of four distinct benefits at equal scope: lower total cost at the same quality, better quality at the same cost, higher incremental profit, or lower material risk.	Trace one real brief through the proposed teams with every approval and transfer logged. Count shared steps, manual repairs and the right to reject a group product. Repeat on two accounts.

CLAIM ANALYSIS	STRONGEST SUPPORT ATTRIBUTED	STRONGEST CONTRARY ATTRIBUTED	RIVAL EXPLANATION ANALYSIS	CONDITION FOR THE CLAIM TO HOLD ANALYSIS	TEST THAT COULD CHANGE THIS READING ANALYSIS
(d) In-housing reduces dependence or total cost.	ANA's 2023 sample: 82% of 162 respondents had an in-house agency [9]. The WFA/Observatory sample: 66% of 45 companies [10].	In the same ANA sample 92% still used outside agencies [9]; all WFA respondents did [10]. Work can move to platforms whose model state the brand cannot export.	Staff location follows cost, speed and talent. The broad survey definition counts any internal team doing agency-type work, so incidence is not share of work or spend.	Holds when the brand keeps the brief, the accounts, the test plan and budget authority, and can run an exit. A small team that owns decisions reduces dependence more than a large one that does execution.	Map dependence by task and by exit time (how long to move each task to another supplier). If total cost and exit time fall after a retained-team build, the claim holds for that brand.
(e) AI efficiency improves client economics or agency pricing power.	A 2025 study of 5,172 customer-support agents found a 15% average gain in issues resolved per hour [12]. The task model in section 18 shows a 25% labour saving with 5 units of new tool and review cost cutting cost from 80 to 70 on a fee of 100 (illustrative). Publicis reported a 17.5% headline H1 2026 margin before LiveRamp transaction costs and 17.4% including them, a company subtotal rather than statutory operating income [46] [38].	The support study is not an advertising profit estimate, and its gains varied by skill [12]. S4 Capital's FY2025 net revenue fell 10.8% while it reported operational EBITDA of £81.2 million: cost can improve while demand weakens [44]. In the task model a fee cut to 85 takes agency profit to 15, below the starting 20 (illustrative). No margin or fee result measures added client sales.	Margins move with headcount, client mix, pricing discipline, currency and the cycle. Pricing power comes from scarce judgment and relationships, not from a tool.	Holds when labour savings are cash savings or useful redeployed capacity, quality is held, and the fee form lets the client share the gain.	A matched-quality total-cost test: approved useful work per total hour, including briefing, review, corrections, legal, integration and failed output, under the old and new process, with the fee change recorded.
(f) Open protocols ensure compatibility, portability or neutrality.	IAB Tech Lab introduced AAMP 3.0 with OpenProposal on 22 September 2026 for comment to 22 October 2026 [20]. The AAMP 2.0 release describes buyer and seller reference SDKs with approval before booking [21]. The AdCP repository, pinned at commit 6adfebe2 on 5 October 2026, holds schemas, reference work and governance documents [22].	AdCP's own FAQ separates task meaning from transport and marks some governance surfaces as experimental; shared transport does not establish the same tasks, permission model or wire compatibility [23]. Horizon's statement that its ecosystem is free of vendor lock-in is unverified [14]. API or MCP access does not establish revocation, liability or shared meaning [13].	Protocols cut integration cost for everyone, including the largest groups, so they do not imply group decline or client portability. Neutrality is a property of the operator's incentives and the contract, not of the wire format.	Compatibility holds between two named implementations at a pinned version after a passed test. Portability holds when state, rights and licences move, not only messages. Neutrality is strongest when the party choosing media has no stake in the choice; having no stake is not enough by itself, because the method, the budget and the reviewer's freedom also shape the result.	A pinned-version, cross-vendor run: a failed request, a retry, a revoked grant and a duplicate booking attempt, with the surviving record identified on both sides, plus an operator exit that moves rules and test history.

CLAIM ANALYSIS	STRONGEST SUPPORT ATTRIBUTED	STRONGEST CONTRARY ATTRIBUTED	RIVAL EXPLANATION ANALYSIS	CONDITION FOR THE CLAIM TO HOLD ANALYSIS	TEST THAT COULD CHANGE THIS READING ANALYSIS
(g) More autonomy improves outcomes.	Google documents goal-led bidding in its inventory [15]. An Amazon Ads page dated 21 March 2025 describes automated buying products; its title and body now give different stages [16]. Publicis describes stock-triggered bids [19].	A 2019 comparison of 15 US Facebook experiments with observational methods found the estimates often differed, so a fast proxy can mislead the system that optimises on it [51]. \$10,000 a day for seven days is \$70,000 spent before delayed feedback arrives (illustrative; not a loss). No reviewed evidence establishes material cross-platform budget authority with profit proof from a party without a sales interest; that is a gap, not absence.	Gains come from better inputs, skilled setup, cleaner feeds and the selection of successes that get published. Speed improves pacing, not necessarily profit.	Holds on bounded tasks with reliable fast signals and a low cost of error. Fails where the goal is wrong, the proxy is gamed or the proof arrives late.	Compare delegated scope at equal risk, quality and total cost, cutting authority whenever errors outgrow benefit. A higher authority level that produced a higher incremental return at the same loss budget would support the claim for that task.
(h) A seller can fairly assess media it also sells.	Publicis promised neutral, open access for the LiveRamp asset it proposes to buy; that is a commitment to evaluate, not proof of neutrality [17]. A seller's own data can improve analysis even while its interest remains.	GroupM's 2011 Xaxis launch joined data and technology with trading scale inside the buying group [42]; the record shows that design, not any separation of advice from sale. No reviewed contract establishes a working separation for any case.	Internal firewalls reduce the interest; they do not erase it. Fairness depends on who sets the proof budget, who holds raw access and who may publish an adverse result. A client-funded reviewer can also depend on the budget owner.	Holds when the method is pre-registered, raw data are accessible to the client, and the safeguards permit an adverse result to be published and acted on.	Ask whether the proof function could report an adverse result on the group's own media product and change the allocation: it needs access, freedom and a route to leadership. Where those are absent the claim stays unestablished; a published adverse result that changed an allocation would be the strongest support.

CLAIM ANALYSIS	STRONGEST SUPPORT ATTRIBUTED	STRONGEST CONTRARY ATTRIBUTED	RIVAL EXPLANATION ANALYSIS	CONDITION FOR THE CLAIM TO HOLD ANALYSIS	TEST THAT COULD CHANGE THIS READING ANALYSIS
(i) Outcome fees align incentives.	WFA's 2026 study of 69 multinationals, 71% in global procurement roles, reported that 63% expected more performance fees [11].	WFA also says such fees remain a small share in most disciplines [11]. Attribution assigns credit and does not prove incremental growth, so an outcome fee can pay for harvested demand [51]. Disclosed CMO pay plans weight company revenue, profit and shareholder return, not media returns, so 'outcome' means different things on each side of the table [49] [50].	Outcome fees reward easy attribution, cheap short-term sales, discounts and channel shifting unless the base case excludes them. Expected growth of a fee model is an intention, not alignment.	Holds with an agreed counterfactual, a profit base after discounts and returns, a window matching the effect, refunds and caps, rules for price, stock and client delays, and audit rights. A fixed price for an output is an output price, not an outcome fee.	Does the fee track added profit after wider effects over the agreed window? If a pool paid out while incremental profit was flat, the claim fails for that contract.
(j) Owning the data means owning the models, state and learning.	Access to source records, under rights that allow export and reuse, is needed for any rebuild; ownership is one way to secure that access, not the only one. A brand that holds its briefs, approvals and test results can at least reconstruct what was decided.	The statement that the brand retains core data, permissions, decision history and allocation is a scenario condition and is unverified as a right. Horizon's lock-in statement is unverified [14]; its API and MCP access does not establish portability [13]. Platform model state stays under platform terms.	Ownership of records is shorthand for control. The usable asset is the decision history plus the rules and tests that produced it, and those are often agency or vendor property.	Holds when the four kinds of learning are separated and the contract grants the access, licence and transfer rights each task needs: export, licence, retention, deletion and exit terms (section 25).	Export the record and rebuild a small workflow with a new supplier inside an agreed period. A lawful, successful rebuild weakens the lock-in concern; a file download the next team cannot use confirms it.
(k) Better agency judgment needs a wholly new business model.	Hours-based fees fall as task time falls, so an agency that sells judgment cannot keep pricing by hours. WFA reported on 29 September 2026 that global brands are shifting to paying for outputs rather than time [11].	Coordination and advice were paid for before AI: Publicis' client-officer plan dates to 2015 [39]. WFA reports output fees growing while outcome fees stay small, which is a change of mix, not of model [11].	Judgment can be paid for through a base retainer, output prices and a bounded outcome pool, each an incremental change. The 'wholly new model' claim may serve to justify technology fees rather than to pay for judgment.	Holds where the fee basis cannot be separated from task time and the client refuses to pay for judgment it cannot see.	Hold scope constant and change the fee form. If clients pay more per unit of demonstrated judgment (fewer wasted tests, better allocation, protected brand health), the claim gains support; if fees move from hours to seats, it does not.

Where the support is thinnest

Claim (b) is graded in three parts. The revenue-model claim (b1) has the thinnest support in this register, because its evidence is a category error. Launches, colleague counts and committed software revenue show intent and skill. A

software revenue model shows up in audited accounts as a recurring fee line with retention. No group in this edition's register shows that line. Publicis' FY2025 accounts split revenue from net revenue by pass-through cost, not by software [43]. Stagwell's \$16 million is committed and mixed [27]. This paper treats (b1) as unestablished for every group. The delivery claim (b2) is different in kind. An agency can run an orchestration service on retainers, project fees or managed-service fees, so a service-fee model can pass the delivery test without passing the revenue test. The register holds described designs, not a logged cross-service workflow on a named account, so (b2) is described but unshown. The client-value claim (b3) is unshown for every group. A subscription product does not pass it by having recurring revenue, and a managed service does not pass it by having a wide scope. The rival readings, a cost reset and cross-selling, fit the same facts at least as well. [43][27] . INFERENCE A037 A017 A018

Claim (f) is weak because the documents that exist are the wrong kind. A draft for public comment [20], an SDK description [21] and a repository at a pinned commit [22] are proposals and specifications. An SDK is a software kit for building on a protocol. Compatibility is a property of two named implementations at a named version after a test. Portability is a property of state, rights and licences. Neutrality is a property of incentives. The AdCP FAQ itself says shared transport does not settle shared meaning [23]. This review ran no compatibility test. It found no public cross-vendor result. Until one exists, (f) is a hope with a schema attached. [20][21][22][23] . INFERENCE A053 A054 A055 A056

Claims (g), (h) and (j) fail together, and for one reason. Each assumes a right the brand has not been shown to hold. Autonomy improves outcomes where a fast signal is a true signal; the 2019 experiment comparison says it often is not [51]. A seller's assessment of its own media can be trusted where someone outside the sales line is free to publish an adverse result; no reviewed contract shows that freedom. Access to data becomes usable learning when the export, licence and transfer rights the task needs are in the contract. The one case in this edition that promises no lock-in has not been verified [14]. The common fix is in Parts II and III: a tested stop, a separate proof function (the team that tests results and reports to the brand, not to the spender) and an exit drill (a rehearsed export and rebuild with another supplier) before scale. [51][14] . INFERENCE A051 H012 H005

HOW TO USE THIS TABLE

Read the condition column before the support column. A claim that holds only under a condition is a claim the CMO must contract for, not assume. Sections 25 and 28 turn the conditions into terms and a 90-day plan.

SECTION 16

What “agentic” can mean:

the authority ladder

Six levels of authority. Each is defined by task, surface, spend scope, duration, approval limit and a stop path.

IN SHORT

Do not call a tool autonomous without stating its level, the money it can move, the surfaces it touches, the time it may run and who can stop it. The ladder is a teaching aid; the task profile beside it records autonomy, integration breadth and financial authority as separate facts. One product can sit at several levels for different tasks. This review ran no live test of any level.

'Agentic' is sold as one thing. It is six different grants. A tool that drafts a plan and a tool that moves budget across platforms without fresh approval carry different risks. They need different logs and fail in different ways. The ladder below names the grant at each level. The same product can sit at level 1 for copy, level 3 for pacing (keeping spend on schedule) and level 5 for a hand-off between two systems. The level belongs to the task, not to the vendor. The ladder is a teaching aid. It orders grants by the discretion they give an agent, and it should not be read as saying that more coordination means more autonomy; the task profile below records the dimensions the ladder mixes as separate facts. • **ANALYSIS**

FIGURE 7 ILLUSTRATIVE

The task profile

Six questions a grant answers for each task, kept separate. A tool can be low autonomy with broad integration. Another can be high discretion with narrow integration.

A task profile, not a ladder

Six questions a grant answers, recorded separately. Stated conditions, not scores: nothing here ranks a task.

	Human-approved cross-system workflow Level 2: human-approved workflow	Single-platform budget optimiser Level 4: adapt within bounds
Recommendation or execution Proposes or acts	Executes after approval	Executes
Adaptive discretion Rules or adapts	Preset rules	Adapts within a stated scope
System scope Systems touched	Many systems	One system
Spend and time limits Money and period	No spend until approved · per campaign	Material spend within the platform cap · campaign length
Approval Who approves, how often	Every step	Per campaign
Reversibility What a stop undoes	Reversible until booked	Partly reversible

Every cell is a design definition from section 16; no product was tested and nothing is measured.

The task profile ILLUSTRATIVE A task profile, not a ladder: six questions a grant answers per task, each a stated condition, with three illustrative tasks compared two at a time. A content tool recommends and spends nothing, yet needs cancellation and data-use controls. An approved cross-system workflow spans many systems with approval at every step. A single-platform optimiser adapts within a stated scope, moves material spend and is partly reversible: a stop halts new actions and leaves booked inventory committed. No score or rank. Illustrative. Units: none; six stated conditions per task, not scores, ranks or timing. Scope: the paper's own design. Text equivalent: appendix "Figure text equivalents".

The authority ladder. Every cell is a design definition, not an observation of any product.

Column classifications: · ANALYSIS

LEVEL ANALYSIS	WHAT THE AGENT MAY DO ANALYSIS	TASK AND SURFACE ANALYSIS	SPEND SCOPE ANALYSIS	DURATION ANALYSIS	APPROVAL LIMIT ANALYSIS	STOP PATH ANALYSIS
1. Recommend or make content	Draft a plan, audience, copy or asset; a person edits or ignores it	Any task; no execution surface	None	Per request	A human edits or ignores every output	Not needed; nothing runs
2. Human-approved workflow	Build a proposal; a named person approves each step before anything is booked	Planning tools, buying SDKs	None until approved	Per campaign	Named approver per step; nothing books without a recorded approval	Withhold approval; queue empties
3. Execute preset rules	A trigger causes an agreed action, such as a stock-out pausing a bid	One platform or one feed	Within the rule's cap	Rule lifetime, reviewed weekly	Rule owner sets caps; exceptions reviewed	Disable the rule; confirm queued actions are cleared
4. Adapt within bounds	Change bids and allocations within set goals and scope	One platform's inventory	That platform's budget cap	Campaign length	Goal and scope control, not per bid	Pause the campaign; check booked commitments
5. Coordinate across systems	Carry state, rights and errors across hand-offs between tools	Several platforms and tools	Bounded per system	Programme	A tested grant; revocation tested before go-live	Revoke the grant; confirm every connected tool, queue and booking received the stop
6. Move material budget alone	Reallocate material budget across a portfolio without fresh approval	Cross-platform	Material	Ongoing	Verified authority, a loss budget and causal review by a party without a sales interest	Kill switch plus a named team that repairs partial actions

The task profile: six dimensions recorded separately

The task profile. Each dimension is recorded on its own for every task an agent performs; no level number implies any of them.

Column classifications: · **ANALYSIS**

DIMENSION ANALYSIS	RECORD AS ANALYSIS	WHY IT IS KEPT SEPARATE ANALYSIS
Recommendation or execution	Whether the agent proposes, or acts on a system	A content tool that cannot spend still needs cancellation and data-use controls
Rule-based or adaptive	Whether the agent follows preset rules or changes its own choices within bounds	Discretion, not integration breadth, drives the need for a loss budget
Single or multiple systems	How many platforms or tools the grant touches	Breadth drives integration and revocation work; it does not raise autonomy
Spend and time limits	The money the agent may move and the period it may run, as amounts and dates	Financial materiality is written down, never implied by a level number
Approval frequency	Per action, per campaign, per rule review, or none	A cross-system workflow approved at every step is low autonomy with broad integration

DIMENSION ANALYSIS	RECORD AS ANALYSIS	WHY IT IS KEPT SEPARATE ANALYSIS
Reversibility	What a stop undoes, and what it does not: a stopped agent halts new actions and does not cancel commitments already booked	Cancellation exposure is recorded as a sum and a list of booked items

Two profiles show why the ladder alone misleads. Take a human-approved cross-system workflow. It records as execution, rule-based, multiple systems, no spend until approved, approval per action, reversible until booked. That is low autonomy with broad integration. Now take a single-platform bidding system that moves a large budget within bounds. It records as execution, adaptive, one system, a material spend limit, approval per campaign, partly reversible. That is high discretion with narrow integration. The second needs the loss budget. The first needs the revocation test. A content tool with no spend still needs a cancellation path and data-use controls. And stopping any agent halts new actions only. It does not undo commitments already booked. That is why the stop path on the ladder names booked commitments. · ANALYSIS

What the record shows at each level

Level 1 has the strongest evidence in this register, and it is not from advertising: the 2025 study of 5,172 support agents tested advice that staff could edit or ignore [12]. Level 2 is described in the AAMP 2.0 SDK release, which includes approval before booking; no live approval log was inspected [21]. Level 3 is an advertised workflow: Publicis describes bids triggered by stock and price signals. That is a specific product description, not tested value [19]. Level 4 is documented by Google for bidding inside its own inventory [15]. Amazon's page for its automated products carries a title and a body that now give different stages, so it cannot date a rollout on its own [16]. Level 5 is described in parts by the SDK and AdCP materials; no end-to-end run was made in this review [21][22][23]. Horizon describes Blu APIs and MCP connections, which is availability, not tested compatibility [13]. Level 6 has no reviewed evidence at scale with profit proof from a party without a sales interest. The one named Horizon client use, SharkNinja, is audience collaboration with retailers, not autonomous buying [13]. [12][21][19][15][16][22][23][13] · ATTRIBUTED A060 A055 A057 A056 H005 H006

Two rules follow. First, a tool may sit at several levels, so a contract must grant levels by task, not by product. Second, the stop path must work at the pace of the spend. A level 4 grant on a platform that books commitments overnight needs a stop that reaches booked commitments, not only the bid engine. This review ran no stop test on any product. Section 28 proposes one for the buyer. · ANALYSIS

A creative lead belongs on this ladder too. At level 1, a model can make more versions of a weak idea. Cheaper versions are not a stronger idea. The creative lead judges distinctiveness, fit and rights before volume. The lead can stop a variant engine at any level. · INFERENCE A047

Protocols: a promise needs a working test

Where each protocol and product sits on the road from proposal to measured adoption, and the controls a buyer must test before trusting any of them.

IN SHORT

OpenProposal is a draft under comment. AdCP's 3.2.1 release is documented at a pinned commit. For platform agent products, search results describe beta or pilot stages, but no primary page was opened, so no stage is recorded. No protocol, tool or group has passed a compatibility, stop, revocation or exit test in any public record this review found. This review ran none.

A protocol is a shared set of rules for how systems talk to each other. It moves through six stages: proposal, specification, reference implementation, compatibility test, deployment and measured adoption. Each stage is a separate fact with its own evidence. The table places the items named in this paper at the stage the inspected record supports, as of 5 October 2026. Cells say 'not established' where this review found no primary record. That is a gap in inspection. It is not a claim that the thing does not exist. · ANALYSIS

Protocol and product stages at the cutoff. Documentation-level claims only; no implementation was exercised by this review.^{[20] [21] [22] [23] [15] [16] [13] [14] [24] [25]}

Column classifications: · ANALYSIS · ATTRIBUTED

ITEM ANALYSIS	PROPOSAL ATTRIBUTED	SPECIFICATI ON ATTRIBUTED	REFERENCE IMPLEMENT ATION ATTRIBUTED	COMPATIBILI TY TEST ANALYSIS	DEPLOYMEN T ATTRIBUTED	MEASURED ADOPTION ANALYSIS	THIS REVIEW'S OWN TEST ANALYSIS
AdCP (Ad Context Protocol)	Published	Schemas, FAQ and governance documents, pinned to commit 6adfebe2 on 5 October 2026 [22][23]	The pinned release notes state version 3.2.1 as generally available, published 30 September 2026 [25]; the FAQ marks campaign governance surfaces as experimental [23]. 3.2.1 is the version analysed; a later patch tag, v3.2.2 of 4 October 2026, is recorded at the pinned commit, its changelog entry lists documentation changes only, and the release notes at that commit carry no 3.2.2 entry [25]	Not established by this review; no cross-vendor result was opened	Not established in a primary record opened here	Not established	None run
Platform agent products (several named platforms)	Announced	Vendor-specific; not either protocol	Vendor-owned	Not applicable across vendors	Search-level results dated before the cutoff describe beta, pilot or closed-beta stages; no primary page was opened, so no stage is recorded here	Not established	None run
Google automated bidding; Amazon automated buying products	Published	Google documents goal-led bidding across its inventory [15]; Amazon's page dated 21 March 2025 describes its products, with a title and body that now give different stages [16]	Vendor-owned	Not applicable	Documented by the vendor within its own surfaces	Not established	None run

ITEM	PROPOSAL	SPECIFICATI ON	REFERENCE IMPLEMENT ATION	COMPATIBILI TY TEST	DEPLOYMEN T	MEASURED ADOPTION	THIS REVIEW'S OWN TEST
ANALYSIS	ATTRIBUTED	ATTRIBUTED	ATTRIBUTED	ANALYSIS	ATTRIBUTED	ANALYSIS	ANALYSIS
Horizon Blu APIs and MCP connections	Announced 18 June 2026 ^[13]	Stated availability ^[13]	Not inspected	Not tested; versions and revocation untested	One named audience- collaboration deployment, issuer-reported ^[13]	Participation counts are not adoption ^[14]	None run

SEE FIGURE 2 · FROM A CLAIM TO EVIDENCE

Protocol and product stages at the cutoff, in brief

The table above is also drawn as the protocol view of the evidence figure in section 2: six stages per item, in a few words each, linked only to the record and version verified. A cell that says not established marks a gap in inspection, not a zero. Shared transport is not shared meaning, rights or state.

Six stages per item, in a few words each (full wording: the section 17 table above and the figure’s text equivalent); the last column names the record and version verified

ITEM	PROPOSAL	SPECIFICAT ION	REFERENC E IMPL.	COMPATIBI LITY TEST	DEPLOYME NT	MEASURED ADOPTION	RECORD AND VERSION VERIFIED
IAB Tech Lab AAMP 3.0 with OpenPropo sal	Released for comment	Draft, pinned by commit	Not established	Not established	Not established	Not established	Cloned by git; pinned commit; commit 66ce40a8; tag v3.0-draft-1
IAB Tech Lab AAMP 2.0	Published	Reference SDKs described	Described by publisher	Not established	No approval log inspected	Not established	Stored passage; AAMP 2.0
AdCP (Ad Context Protocol)	Published	Schemas pinned by commit	3.2.1 GA; governance experimental	No cross- vendor result	Not in a primary record	Not established	Cloned by git; pinned commit; version analysed 3.2.1 at commit 6adfebe2; patch tag v3.2.2 recorded, not substituted
Platform agent products (several named platforms)	Announced	Vendor specific	Vendor-owned	Not applicable	No primary page opened	Not established	None opened
Google automated bidding; Amazon automated buying products	Published	Documented by vendor	Vendor-owned	Not applicable	In vendor surfaces	Not established	Vendor pages, stored; Google page undated and living; Amazon page dated 21 March 2025
Horizon Blu APIs and MCP connections	Announced 18 Jun 2026	Stated availability	Not inspected	Not tested	Named use: audience collab.	Counts, not adoption	Issuer releases, stored; none stated

VERSION RECORD

Version analysed: 3.2.1 at commit 6adfebe2 (2026-10-05); generally available from 2026-09-30. **Tags:** v3.2.1 (c32bd78c, 2026-10-01); v3.2.2 (e3b3aaa9, 2026-10-04). **Later patch tag inside the cutoff:** 3.2.2 of 2026-10-04: The CHANGELOG.md entry at the pinned commit lists two patch changes, both documentation: how a buyer-declared account reference behaves before provisioning, with one recovery for ACCOUNT_NOT_FOUND; and consistency fixes between the trust and verification docs and the spec (signed responses limited to the brand-claim verification tasks; the relationship_trust values and house_domain field from the schema). No schema or wire-protocol change is named; the wire pin stays 3.2. docs/reference/release-notes.mdx, versions.mdx and whats-new-in-3-2.mdx at the pinned commit carry no 3.2.2 entry; the two changesets are still present under .changeset/ at that commit. **Rule:** The version analysed stays 3.2.1 at commit 6adfebe2. The 3.2.2 patch tag of 4 October 2026 is inside the 5 October 2026 cutoff and is reported beside it; it is never substituted for the version analysed.

Tag commits and dates as the source register recorded them on 5 October 2026 (ROOT-S005, ROOT-S302). The local pinned clone is a depth-1 clone that holds no tag objects, so the editorial repair of 6 October 2026 could not re-verify them from the tags themselves. Inspected: Editorial repair, 6 October 2026, on the local pinned clone only (no network, nothing re-fetched): git rev-parse HEAD and git log -1 (6adfebe2, 2026-10-05 09:11:57 -0400); git tag (none present; shallow clone); package.json; CHANGELOG.md top entry; .changeset/; docs/reference/release-notes.mdx, versions.mdx and whats-new-in-3-2.mdx. [22] [25]

The paper names six stages (the brief lists five; the paper's proposal stage is kept). Each stage is a separate fact with its own evidence. A cell says not established where this review found no primary record; that is a gap in inspection, not a claim that the thing does not exist. Shared transport is not shared meaning, rights or state. A buyer needs three separate proofs: that messages arrive, that both sides mean the same thing by them, and that the rights and state behind them move too. Tests run by this review: None. The thirteen controls and the pair test are proposed for a buyer.

The AdCP row carries one further disclosure. The author co-leads a working group of the body that publishes AdCP; see the disclosure at the end of this paper. Nothing in this paper treats AdCP membership, listings or announcements as evidence of compatibility or adoption. The author's own local work on the protocol is not cited. • **FACT**

Shared transport is not shared meaning, rights or state

Transport carries a message. A domain specification gives the message meaning. Two tools can share a transport and still treat budgets, units, consent, errors and state differently. AdCP's own FAQ separates task meaning from the MCP or A2A transport beneath it [23]. Even a sound technical link can lack the legal right or business term needed to use it. So a buyer needs three separate proofs: that the messages arrive, that both sides mean the same thing by them, and that the rights and state behind them move too. [23] • **INFERENCE** A056

The controls a buyer must test

The checklist below is proposed, not run. It turns each control named in this paper into a question with a pass condition. This review ran no live media spending or customer-data experiments to back any of them. It does not claim that any named vendor has passed them. • **ANALYSIS**

Thirteen controls as questions. Proposed for a buyer; none was run by this review.

CONTROL ANALYSIS	THE QUESTION TO TEST ANALYSIS	WHAT A PASS LOOKS LIKE ANALYSIS
Permission	Can the grant name tasks, accounts, sums and time windows, and nothing else?	A request outside each bound is refused and logged
Delegation	Can a grant be passed to a sub-agent or partner without widening it?	The delegated grant is no wider than the parent and is visible in the log
Revocation	Can the grant be withdrawn mid-run?	Every connected tool, queue and booked commitment receives and confirms the withdrawal
Stop	Can a named person halt all activity at the pace of the spend?	Spend stops within the agreed window; partial actions are listed for repair
Identity	Can each data input be traced to its source and owner?	Provenance recorded per input; unknown sources rejected
Consent and permissible use	Is each input used only for its permitted purpose and expiry?	Purpose and expiry are enforced, not only documented
Logs	Does each action record input, model or tool version, rule, approver, time and result?	A sampled action can be reconstructed from the log alone
Audit access	Can the brand's proof function read logs directly, not through the operator's dashboard?	Direct read access exercised; an adverse review delivered
Memory	Is learned state listed, owned and bounded per client?	No client's state is pooled with another's without written consent
State export	Can goals, rules, test history, settings and notes be exported in a named format?	An export is read and reused by a second implementation
Recovery	Are retries prevented from duplicating spend, and who repairs partial actions?	A forced failure and retry produce one booking and one named repairer
Liability	Who pays for wrong buys, data misuse, outages and missed stops?	The deal names the payer for each; a cap is not a liability ceiling
Proof delay	Are exposure limits set while results mature?	A fast proxy never counts as a causal result; spend before proof is bounded and funded by a named party

The buyer's compatibility test, proposed: name the versions and implementations on both sides. Send a failed request, a retry, a revoked grant and a duplicate booking attempt. Record which record survives on each side, who is liable for the duplicate, and whether the revocation reached queued and booked commitments. Then export the state and rebuild one workflow with a second implementation. Pass or fail is per pair and per version. It is not a property of the protocol. · ANALYSIS

Tests run by this review: none. Tests proposed for a buyer: the thirteen questions and the pair test above. The paper claims that these controls must be tested. It does not claim that any protocol provides them. A specification can describe a field for a grant without any implementation enforcing it. · FACT

Who earns, who saves, who carries risk

Seven fee routes, six actors, the gap between spending and proof, and one task model with one boundary.

IN SHORT

Fees are a design choice; none creates neutrality by itself. On every route, name who funds spend before proof, who funds failed tests and who keeps a saving. The task model shows that a saving can go to the agency, to the client or be over-passed. It does not show added sales.

Section 14 named nine mechanisms of value capture. This section follows the money through seven fee routes and six actors. It does not rank the routes. For each it shows who earns, who saves, who carries the risk, and what a buyer must test. Then it prices one task in one model with one boundary. · **ANALYSIS**

FIGURE 8 ILLUSTRATIVE

Seven fee routes


On each route the brand pays. Who earns as volume rises, who carries the cost and what to test differ, and every client saving is conditional.

Agency buying versus principal resale


The same three relationships on both routes. The reseller's spread sits inside the bundled price; it is never a second line.

→ Money - - -> Advice ···> Evidence Boxes and arrows carry no amounts.

Agency buying A supplier acts for the client under a mandate.



Principal resale A supplier buys media on its own account and resells it. The contract sets price, risk and rights.



The brand pays one bundled price; the spread is inside it, disclosed, never added again as a second line.

Principal and agency buying may coexist inside one mandate, stated line by line. Principal buying is acceptable for a defined scope with an audited comparison. No flow is drawn with a quantitative width. No comparable, sourced flow values exist in the record.

Record: 58% of 114 qualified client-side ANA members had used principal media in the past year, up from 47% in the prior wave. A member sample, not a share of spend.

Seven fee routes ILLUSTRATIVE Section 18's entry to the commercial-model explorer: agency buying beside principal resale, with money, advice and evidence drawn as three relationships and the reseller's spread inside the bundled price, never a line on top. Select any of the seven fee routes for payer, earner, cost bearer, benefit condition, conflict and evidence. A fee route proves neither neutrality nor abuse. Status: illustrative (analysis). Units: none. Scope: the paper's own design. A058 A059 V3-097 A061 A062 A037 A028 A029 A069 A071 V3-094 Text equivalent: appendix "Figure text equivalents".

Seven fee routes. Who earns, who saves, who carries risk, and the test that separates value created from value captured.^{[49][50][8][27][7][11]}

Column classifications: · **ANALYSIS**

ROUTE ANALYSIS	WHO EARNS AS VOLUME RISES ANALYSIS	WHO CAN SAVE ANALYSIS	WHO CARRIES THE RISK ANALYSIS	WHAT TO TEST ANALYSIS
Labour or service fees	The agency, when billable scope or hours rise	The brand, if the work changes decisions at a fee below its own cost of doing it	The agency carries delivery cost; the brand carries rework it cannot see	Approved useful output per total hour; quality held; rework falling; fee against retained-team cost
Spend-linked fees	The agency, when billable spend rises	The brand only if spend growth is profitable at the margin	The brand carries unprofitable spend	Marginal client profit per additional unit of spend, after returns and discounts; disclosed CMO pay plans reward company profit, not spend [49][50]
Principal media spread	The reseller, before its service cost and inventory risk	The brand, if the bundled price for comparable inventory is below the agency route and it can assess the product	The reseller carries inventory risk it chose; the brand carries the price it cannot compare	Role disclosed per line; comparison on equal inventory, timing, quality, data rights and payment terms; use rose from 47% to 58% in ANA's samples [8]
Technology subscriptions and platform fees	The provider, from seats, usage and renewals	The brand, if the tool replaces a dearer process and exit charges are bounded	The brand carries overlapping licences and exit cost	True usage; overlapping licences removed; exit cost priced; committed revenue is not recognised revenue [27]
Data services	The supplier, from permitted data work and access	The brand, if match quality and lawful use improve targeting or measurement and the data can move	The brand carries unlawful use and lost provenance	Provenance, match rate, use rights, portability; the one vendor case is a before-and-after without a control [7]
Strategic analytics and advice	The partner, for analysis and advice	The brand, if the advice changes a decision for the better	The brand carries advice shaped by a sales interest	Separate the sales incentive from the method, the budget and the reviewer
Performance or outcome fees	The partner, when the agreed metric clears the hurdle	The brand, if the metric is incremental profit over a matching window	The brand carries fees paid on harvested or shifted sales	Fee tracks added profit after substitution, discounting and brand effects; refunds and caps exercised; 63% of 69 WFA respondents expect more such fees, still a small share [11]

Six value-chain actors: what each gains and what each carries.

Column classifications: · **ANALYSIS**

ACTOR ANALYSIS	WHAT IT GAINS ANALYSIS	WHAT IT CARRIES ANALYSIS
Brand	More profitable demand, time saved, better learning	Funds spend, retained staff, tests and weak outcomes unless the contract shifts them
Group	Shared tools, broader client access, scale and retained margin	Capital, integration, staff, delivery and some contract risk

ACTOR ANALYSIS	WHAT IT GAINS ANALYSIS	WHAT IT CARRIES ANALYSIS
Platform	Media income, tool use and more of the client workflow	Infrastructure and service duties under its own terms; the brand's profit risk stays with the brand
Publisher	Demand, yield and a route to buyers	Production, audience trust, payment terms and loss of direct customer ties
Technology provider	Usage, licence and service income	Compute, product upkeep and agreed service risk
Specialist or proof team	Fees for craft, tests or assurance	Skill and delivery costs; a reviewer's freedom needs its own incentives

The spend-to-proof gap

Fast signals arrive in hours. Causal results arrive in weeks. Brand effects can take longer. Money leaves in between. Illustration only: \$10,000 spent each day for seven days is \$70,000 before the first delayed read arrives. That figure is not a forecast loss and not a hard risk ceiling. Booked commitments, reporting delay and stop failure can add exposure beyond it. Spend awaiting proof is not a loss; it is spend whose return is not yet known. The party that funds the spend carries that gap unless the deal moves it. · **ILLUSTRATIVE**

Two wording rules follow for every contract and every chart in this paper. A spending cap is a brand-set rule about exposure. It is not a liability ceiling, because commitments and a failed stop can exceed it. And spend before proof is never described as certain loss. Name the payer for failed tests before launch. A brand can fund a ring-fenced learning budget. A partner can share the risk for a fee. Either way, log failures as learning. Do not let an outcome fee reward only the easy wins while the client funds all the hard tests. · **ANALYSIS**


FIGURE 9 ILLUSTRATIVE

Spending to proof: two clocks

Money leaves every day. Causal evidence comes later. The gap is spend whose return is not yet known. It is not a loss, and a cap is not a ceiling on what the brand may owe.

Two clocks: money leaves every day, evidence arrives later

Each column is one day. The dark part is money out with no read yet. A hard pre-spend cap clips each day's spend to the room left; it does not cap what the brand may owe.



At the first read (day 7): spent \$70,000, all of it awaiting evidence · commitments \$0 · remaining authority \$20,000.

Stop point: day 9, when spent plus commitments reach the cap of \$90,000 (that day's allowed spend \$10,000, clipped by the pre-spend check). Further spend needs an approval.

Exposure beyond the cap at these inputs: \$0. A slow stop or commitments above the cap take it past the cap; the stop-delay and commitments inputs show that as a number.

Read schedule: the first read lands on day 7; after it each day's spend is read 7 days later (a rolling lag), not one assessment of all earlier spend.

allowed spend each day = $\max(0, \min(\text{proposed daily spend}, \text{cap} - \text{spent} - \text{commitments}))$ · awaiting = spend whose read is not yet due (the last 7 days) · remaining authority = $\text{cap} - \text{spent} - \text{commitments}$, never below 0 · excess = spend after the stop was due (slow-stop assumption, default none) · beyond the cap = $\text{spent} + \text{commitments} - \text{cap}$, when positive

Assumed: \$10,000 a day, divisible, under a hard pre-spend cap · first read after 7 days, then a rolling read 7 days after each day's spend · cap \$90,000 (illustrative default, borrowed from the section 26 mandate, not a recommended brand limit; section 18 states no cap) · commitments \$0 · the first read lands on time · the stop bites the same day (no slow-stop overshoot assumed). Media only; agency fees sit on a separate line.

Spend awaiting evidence is not a loss. A cap is not a liability ceiling. Cash exposure is not a financing expense.

Spending to proof: two clocks ILLUSTRATIVE Illustration only: \$10,000 a day under a hard pre-spend cap; first read after seven days, then a rolling read. \$70,000 is out before the first read. Spent, commitments, remaining authority and spend past the stop are separate lines, counted once. At the borrowed \$90,000 cap the stop falls on day 9. Not a loss, not a liability ceiling. Status: illustrative. Units: USD, illustrative; days. Scope: One bounded exposure illustration with assumed inputs; media only; not a profit model, not an estimate for any brand and not a forecast of loss. Text equivalent: appendix "Figure text equivalents".


One task, one model, one boundary

The model below prices one task. The base fee is 100 units. Labour costs 60 and other costs 20, so profit is 20. Suppose a new way of working cuts labour by 25% and adds 5 units of tool and review cost. Cost falls to 70. At a fixed fee the agency's profit rises to 30 and the client saves nothing. At a fee of 85 the client saves 15, more than the 10-unit net saving, and the agency's profit falls to 15, below where it started. At a fee of 90 the client takes the whole net saving and the agency is back at 20. With no labour saving at all, the added cost is a pure loss to whoever bears it. All inputs are assumptions in arbitrary units. · **ILLUSTRATIVE**

FIGURE 10 ILLUSTRATIVE

Who keeps a task saving?

A labour saving on one task can go to the agency at a held fee, or to the client at a cut fee, or be eaten by new cost. Nothing here measures added sales.



$\text{cost} = \text{labour} \times (1 - \text{saving_pct} \div 100) + \text{other} + \text{new cost} \cdot \text{profit} = \text{fee} - \text{cost} \cdot \text{client fee saving} = 100 - \text{fee}$
Assumed: fee 100 · labour 60 · other 20 · saving 25% · new cost 5 (the added-cost case uses 15). Arbitrary units; base fee 100.
Output and quality held fixed. Nothing here measures added sales.

Who keeps a task saving? ILLUSTRATIVE One assumed task, arbitrary units: fee 100, labour 60, other 20; a 25% labour saving plus 5 new cost cuts cost 80 to 70. Four presets at those inputs: the agency keeps the 10-unit net saving at a held fee (profit 30); the client keeps it at a fee of 90 (profit 20); an equal split at a fee of 95 (profit 25, client 5; a new illustrative preset added 6 October 2026); and new cost of 15 absorbs the saving (cost 80, profit 20). Output and quality fixed; no added sales; not a forecast of loss. Status: illustrative. Units: arbitrary units; task model base fee = 100. Scope: Sensitivity illustration with assumed inputs; not an estimate of agency prices, achievable AI gains or any firm's spread. Text equivalent: appendix "Figure text equivalents".

Who keeps a task saving? A sensitivity model

Every input is an assumption in arbitrary units; the base fee stays 100. Formula: cost = labour × (1 – saving ÷ 100) + other + new cost; profit = fee – cost; client fee saving = base fee – fee. The chart, the table and the sentence are one calculation.

Client fee

100

Assumption: The base fee is 100 units by assumption. A lower fee passes part of any saving to the client; at 90 the client takes exactly the default net saving.

Labour cost at the base

60

Assumption: The labour share of the base cost. The saving applies only to this line.

Other cost

20

Assumption: Non-labour cost that the saving does not touch.

Labour saving

25

Assumption: The share of labour cost removed by the new way of working. Zero means no saving, and the new cost then becomes a pure loss to whoever bears it.

New tool and review cost


5

Assumption: Licences, integration and the human review the new way of working adds. At 15 it swallows the default saving exactly.

A 25% labour saving less 5.0 of new cost takes cost from 80.0 to 70.0. At a fee of 100.0 the modelled margin is 30.0 (base 20.0); the client’s fee is unchanged. No added sales are assumed.

Assumed: fee 100 · labour 60 · other 20 · saving 25% · new cost 5 · base fee 100 (arbitrary units)

MODEL BOUNDARY One task priced at an assumed fee of 100 units, with output and quality held fixed and no added sales. It is a sensitivity test in arbitrary units, not an estimate of agency prices or achievable AI gains. It excludes tax, financing, media spend, working capital and quality change. A negative result is an arithmetic consequence of the chosen inputs, not a forecast of loss. This is the only cost model in this edition.



CASE	LABOUR	OTHER	NEW TOOL / REVIEW	COST	FEE	MARGIN	CLIENT FEE SAVING
Base	60.0	20.0	0.0	80.0	100.0	20.0	0.0

CASE	LABOUR	OTHER	NEW TOOL / REVIEW	COST	FEE	MARGIN	CLIENT FEE SAVING
Chosen	45.0	20.0	5.0	70.0	100.0	30.0	0.0

The model holds output and quality fixed and assumes no added sales. It is a sensitivity test in arbitrary units, not an estimate of agency prices or achievable AI gains. It leaves out tax, financing, media spend, working capital and quality change. A negative result is an arithmetic consequence of the chosen inputs, not a forecast of loss. Labour savings have value only when they are cash savings or useful redeployed capacity. This is the only cost model in this edition. No whole-budget model sits beside it. • **ASSUMPTION**

One honest task test goes with the model. Measure approved useful work per total hour, not generated items per prompt. Count the briefing, the review, the corrections, the legal check, the integration and the failed output. Record which team absorbed new work. Then record where the saving went: kept by the agency, passed to the client, or lost to new cost. • **ANALYSIS**

SECTION 19

The CMO must manage two time horizons

Annual targets and long-term equity can pull in different ways. Creative judgment, brands and long-term demand must stay visible in any operating model.

IN SHORT

Fast loops reward what responds this week. Brand effects need a longer window. Four kinds of false gain, harvesting, channel shifting, discounting and brand erosion, can make a fast system look like growth. Run harvest and build as a portfolio, not a binary period choice. Keep a creative lead with real authority. Control rights protect the relationship; brief quality, timely decisions and useful feedback are what make the work better.

Two disclosed pay plans show the problem. Coca-Cola's 2026 proxy separates annual growth measures from multi-year share awards; the latter carry a relative total shareholder return modifier ^[49]. Kohl's fiscal 2025 annual incentive weighted merchandise sales 60% and operating margin 40%, with long-term equity awards separate ^[50]. These are two company cases, not a norm for CMO pay. They show that the annual plan and the long-term award are measured on different clocks. Neither is a media return. ^{[49][50]} • **ATTRIBUTED** A058 A059

Share price is the price of one share. Market capitalisation also depends on share count. Total shareholder return includes distributions over a period under a stated method. They are not interchangeable goals. A stock move alone does not prove what marketing caused. · **INFERENCE**

Harvest and build, as a portfolio

A CMO can meet an annual cost target while weakening future demand. The reverse risk also matters: a vague brand claim can shield waste. So the allocation in any period is explicit, and it is a portfolio, not a binary choice. Some products, audiences and markets are in a harvest phase, in which the brand takes the demand it has built at the best margin it can. Others are in a build phase, in which it accepts a lower near-term return to grow demand it will harvest later. Each gets its own budget, horizon and measure, and brand building is part of the ongoing allocation, not a residual left after acquisition. A fast system cannot tell the two apart on its own. A system driven by fast proxies may favour immediately measurable demand, because harvested sales respond this week and built demand does not; the objective, the signals and the constraints it is given decide how far that bias goes. · **ANALYSIS**

Four ways an apparent gain can be false. First, harvesting: an automated system claims sales that would have happened anyway. The 2019 comparison of 15 US Facebook experiments with observational methods found the estimates often differed ^[51]. Treat every attributed gain, credit assigned by a rule, as credit, not cause. Second, channel shifting: each platform claims the same sale, so adding attributed sales across silos overstates the total. Test substitution across channels, not only within one. Third, discounting profit away: an offer engine can raise conversion by lowering margin. So the outcome base must be profit after discounts, returns and fulfilment, not revenue. Fourth, weakening brand health: fast feedback pushes spend toward what responds now, and brand effects need a longer window. ^[51] · **INFERENCE**

Creative judgment is a separate test of value

Keep a creative lead with real authority. A model can make more versions of a weak idea. It can also help teams test good work faster. Cheaper versions are not a stronger idea. Judge distinctiveness, fit, rights and quality before volume. Keep room for a new idea whose value will not show in this week's conversion report. The creative lead can reopen a brief, reject an offer on brand grounds and stop a variant engine. That holds in every one of the three paths in Part III. · **INFERENCE**

A047

This is the part of the loop that cheaper production does not automate. Section 1 said it: a brand must keep that judgment, or knowingly buy it with a right to overrule. Each path in Part III names where that right sits. · **ANALYSIS**

Measure long-term demand on purpose

Set near-term and long-term tests in advance. Use margin and cash alongside reach, customer health and demand measures suited to the business. Use attribution to track and assign credit where useful. Use sound tests to estimate added effect. Use longer models and judgment where tests have limits. Keep the model version, design and confidence range, so a later test can check them. Reconcile the three. Never assign a number to a brand effect the paper has not measured; this paper assigns none. · **INFERENCE**

Scope beyond paid media matters here. CRM, experience, retail operations, sports, creators, PR and production reach different parts of demand on different windows. A creator programme cannot be judged by a retail search signal. Any operating model should manage those differences rather than force all work into the easiest reported metric. · **ANALYSIS**

What makes the work better: the value-production loop

Control rights decide who can set, stop, test and keep. They do not by themselves make the work better. A second loop does that, and it runs in every path in Part III. A clear brief states the business problem, the audience and the constraints. A shared understanding of the customer gives both sides the same picture; it is built from the brand's facts and the agency's research. Creative choices are made and defended by a named lead. Client decisions arrive on time, because a late approval wastes the work that depends on it. Feedback is specific and usable, so the next round improves rather than restarts. Learning is kept: what was tried, what worked, and why. Control rights protect this loop. The quality of these interactions is what produces the value. A brand should test both. Section 28 adds brief quality, approval delay and rework to the measures. · **ANALYSIS**

- Brief: the business problem, the audience, the constraints and the measure of success, written before work starts
- Shared customer understanding: one picture of the customer that both sides can see and challenge
- Creative choices: made and defended by a named lead, with the right to reject volume
- Timely client decisions: approval windows agreed in advance and measured
- Useful feedback: specific, reasoned and recorded, so the next round improves
- Learning kept: what was tried, what worked and why, in a form the next brief can use

FIGURE 11 ILLUSTRATIVE


How judgment creates value

Control rights decide who can set, stop, test and keep. They do not make the work better. A clear brief, a shared customer picture, creative choices, timely decisions, useful feedback and kept learning do that.

How judgment creates value

illustrative

The paper's proposed operating logic, not a measured causal growth loop. Arrows show reading order, not cause. Control rights protect this loop; the quality of these interactions is what produces the value. A brand should test both.



One measure, or one proposed diagnostic question, per stage

Select a stage to read its measure or question; the table lists all six. Measures from the paper: 1 brief, 3 creative choices, 4 timely client decisions, 5 useful feedback. Proposed diagnostic questions, not measures (the paper names none): 2 shared customer understanding, 6 learning kept.

Near-term delivery

- Approved useful work per total hour
- Brief quality, approval delay, rework share
- Attributed results, as credit not cause

Long-term brand value

- The brand-health read, kept beside delivery metrics
- Harvest and build run as a portfolio, each with its own budget, horizon and measure
- No number assigned to a brand effect the paper has not measured

Four ways a gain can be false: harvesting; channel shifting; discounting profit away; weakening brand health. No number is assigned to a brand effect the paper has not measured. Rankings, not causal effects: WFA respondents ranked briefing, trust and feedback above money incentives (attributed; the source was not re-opened).

How judgment creates value ILLUSTRATIVE Six interactions the paper says make the work better, in order: a clear brief, a shared customer picture, creative choices by a named lead, timely client decisions, useful feedback and learning kept. Four stages carry a measure from the paper; two carry a proposed diagnostic question, not a measure, where the paper names none. Near-term delivery sits beside long-term brand value. Proposed logic, not a measured loop. Units: none (stages in reading order; measures named, not scored). Scope: The paper's own design from section 19, with the measures section 28 and section 24 add; two diagnostic questions are proposals by the figure data, labelled as questions and not as measures; no measured link from any stage to a business result is held.^[11] v3-118 Text equivalent: appendix "Figure text equivalents".

The paper's own WFA source gives this loop support of a bounded kind. The external review of 6 October 2026 reports that the WFA's respondents ranked briefing, trust and feedback above financial incentives as contributors to agency performance ^[11]. Those are respondent rankings, not causal effects, and this edition did not re-open the source; its stored passage covers the fee findings, not the ranking. The ranking is a credible rival to the view that contracts and audit alone decide agency performance, and this paper treats it as one. ^[11] · **ATTRIBUTED**

V3-118

SECTION 20

The present-day baseline

Eight tasks, seven stages, two clocks. An illustrative fragmented workflow, not a measured account of how teams work today.

IN SHORT

Today's loop runs brief, planning, activation, measurement, attribution, reporting, optimisation and a revised plan. Fast delivery checks and slow causal evidence are different loops. Attribution assigns credit; it does not prove incremental growth. An illustrative fragmented workflow may lack an owner for decision history, a proof function separate from the spender, a stated authority level per task, a tested stop and a priced exit; mature teams may already hold some or all of these.

Before the three paths, the present. The work between a brand's goal and its next decision runs through eight tasks: brief, planning, activation, measurement, attribution, reporting, optimisation and a revised plan. The tasks overlap. A team may report while it is still buying. Creative work can reopen the brief. A test can pause a campaign. A later brand result can change next quarter's goal. The figure shows the eight tasks as seven stages. Attribution and reporting are split from measurement so that credit and cause stay distinct. The baseline is an illustrative fragmented workflow. It is not a measured account of prevailing practice, and this review holds no prevalence data on how many teams work this way. · **ANALYSIS**

FIGURE 12 ILLUSTRATIVE

The present-day loop and the three paths, in one layout


Seven stages from brief to revised plan. Fast checks run on proxies. Causal evidence comes later, and only if tests are run. Attribution gives credit; it does not prove growth.

Baseline: the loop today

illustrative

Eight tasks as seven stages: Brief 1 · Planning 2 · Activation 3 · Measurement 4 · Attribution 5 · Reporting 5 · Optimisation 6 · Revised plan 7. Attribution and reporting are split from measurement so that credit and cause stay distinct.

Each box: who runs the stage, then "Approval:" whose sign-off its output needs. "Within the grant" or "within the approved plan" means no new sign-off inside limits already granted.



Attribution assigns credit; it does not prove added growth. Causal and brand evidence come later, and only if tests are run. Illustrative, not a measured sequence. No arrow shows timing, chance or size of effect; a drawn flow is not proof of cause. The baseline is an illustrative fragmented workflow, not a measured account of how teams work. No prevalence data is held.

Who sets goals, runs the work, approves more spend, tests growth, retains learning and carries costs

QUESTION	BASELINE (SHOWN)	A. BRAND-CONTROLLED	B. AGENCY-LED	C. PLATFORM-DOMINATED
Who sets goals?	Brand marketing lead	CMO with finance and data leads	Brand goal owner; the lead agency designs the service	Brand with finance or an adviser
Who runs the work?	Agency or in-house operator; platform systems	Named specialist within the grant	Named agency team or agent under contract	Each platform within its own system

QUESTION	BASELINE (SHOWN)	A. BRAND-CONTROLLED	B. AGENCY-LED	C. PLATFORM-DOMINATED
Who approves more spend?	Budget changes beyond the approved limits need the agreed approval (brand)	Brand budget owner; any wider grant returns to stage 3	Brand reserves large moves; changes beyond the grant pass through stage 3	Brand or delegated agency owner, per platform; rebalancing returns to stage 3
Who tests growth?	Causal and brand evidence arrive later, if at all; attribution assigns credit	Brand-funded proof team separate from sellers	Proof function with its own mandate, budget and access	Brand-funded team separate from media sellers
Who retains learning?	Decision history may not be kept as an asset	Brand: decision history and tests; vendor model state stays behind unless licensed	Brand: client facts and decision history; agency: methods; vendors: model state	Brand: the portfolio decision history; hidden state stays behind
Who carries costs?	Not set per stage. The party that funds the spend carries the spend-to-proof gap unless the deal moves it.	The brand pays retained staff, integration, data and technology, measurement and exit directly; buys services and production at disclosed prices; funds failed tests from a ring-fenced budget and carries spend before proof	Services, production, technology and integration paid through the partner at itemised prices; retained staff, proof and exit stay with the brand; the brand bears the cash-to-proof gap unless the partner accepts part of it	Platform spend counted once as media; platform fees as technology; the proof team, the portfolio owner and exit cost beside them; the brand funds most failed demand tests and carries delayed-result risk unless a contract says otherwise

Owners as the paper’s tables state them; a scenario column is a condition of that design, not a forecast. Open interfaces do not by themselves make learning portable; a closed system is not worse by definition.

The present-day loop and the three paths, in one layout

ILLUSTRATIVE

The present-day loop as seven stages with two clocks; the same layout then shows paths A, B and C, with a comparison table beneath. Fast checks run on proxies inside the approved limits; causal evidence reaches the plan later, and only if tests are run. Attribution assigns credit, not added growth. Illustrative baseline; the paths are conditional designs, not forecasts. Units: none (process diagram). Scope: A guide to roles and loops in the present-day workflow; not a claim that all agencies work in one sequence, and not a measured sequence.

[51] A052 Text equivalent: appendix “Figure text equivalents”.

Eight tasks mapped to seven stages in an illustrative fragmented workflow: an owner each stage may have, who may hold the data and learning, what is measured, and which loop the stage belongs to.

Column classifications: **ANALYSIS**

STAGE (BASELINE TASK)	DECISION OWNER, IN THE ILLUSTRATION	DATA AND LEARNING HOLDER	MEASUREMENT ROLE	LOOP
ANALYSIS	ANALYSIS	ANALYSIS	ANALYSIS	ANALYSIS
1. Brief	Brand marketing lead	Brand holds the brief; the agency holds its interpretation	None	Input to the planning loop
2. Planning	Agency planner with brand approval	Agency tools and data; platform planning data	Test plan may be unfixed at launch	Planning loop

STAGE (BASELINE TASK)	DECISION OWNER, IN THE ILLUSTRATION	DATA AND LEARNING HOLDER	MEASUREMENT ROLE	LOOP
ANALYSIS	ANALYSIS	ANALYSIS	ANALYSIS	ANALYSIS
3. Activation	Agency or in-house operator; platform systems	Platform and agency logs; principal routes recorded by the reseller	Delivery checks	Fast loop
4. Measurement	Agency analytics; brand analyst	Mixed; raw access may be incomplete	Delivery metrics and attributed results	Fast loop output
5. Attribution and reporting	Agency or platform reporting	Each platform's own credit rule; agency reconciliation	Assigns credit; does not prove incremental growth [51]	Fast loop; can be mistaken for learning
6. Optimisation	Operator within platform or agency rules	Platform model state; agency rules	Acts on proxies	Fast loop, back to activation
7. Revised plan	Brand with agency	Decision history may not be kept as an asset	Causal and brand evidence arrive later, if at all	Learning and planning loops, back to the brief

Two clocks run through this loop. Fast operational feedback runs from measurement through optimisation back to activation, inside the limits already approved. Delayed causal and brand evidence runs from measurement through tests to the revised plan and the brief. Budget changes beyond the approved limits need the agreed approval. Attribution assigns credit; it does not prove incremental growth, the growth that would not have happened without the action. A fast system may act on a provisional signal. Its log must keep that fact, so that later proof can reverse the plan. • **INFERENCE** A052

What the illustrative workflow may lack, and what each path in the next three sections names as a right to secure: an owner for decision history; a proof function separate from the party that spends; a stated authority level per task; a tested stop; and a priced exit. None of these is a technology. Each is a right that a contract either grants or does not. Mature teams may already hold some or all of them. The paths do not invent these controls. They make them explicit. The baseline and the paths face the same standard of evidence. • **ANALYSIS**


A shorthand helps before the paths. This paper calls today's loop PAM: plan, activate, measure. It is one wheel. It turns once per campaign, with the fast check and the slow read inside it. The figure below sets it beside a possible future. The loop survives. Three things change. First, it runs as three loops on three clocks. A fast loop acts within a grant and never counts as proof. A learning loop turns on evidence. A planning loop revises the goal. Only the last two may change the mandate, and only through a logged approval, as section 26 sets out. Second, grant, record and proof become steps of their own. Third, the stop works at any point; it does not wait its turn. The five rights of section 1 sit at the centre, whoever holds them. In path A the brand holds the goals, the grants and the record. In path B the agency runs the loop, and the brand keeps the goals, the stop and the big budget calls. In path C each platform runs its own fast loop. This is a design, not a forecast. The motion is a reading aid, and the speeds are relative, not to scale. • **ANALYSIS**

FIGURE 13 CONDITIONAL

From PAM to the governed loop

The loop survives. In the possible future it runs as three loops on three clocks around the five rights, with grant, record and proof as steps of their own and a stop that works at any point.

Past · PAM: plan, activate, measure



Where PAM breaks

It closes on credit, not cause.

Attribution assigns credit, not added growth, yet its report usually feeds the next plan.

One wheel hides two clocks.

Fast checks on proxies and slow causal reads are different loops; drawn as one, the fast passes for proof.

The learning may leave.

Platform model state and agency rules may hold it; decision history may not be kept as an asset.

Possible future · the governed loop


Three loops on three clocks around the five rights. A conditional design, not a forecast.

--- Fast loop · hours to days

act, signal, adjust within the grant; never proof

Learning loop · weeks ·····

proof changes what agents may do next



From PAM to the governed loop **CONDITIONAL** The present-day loop (PAM: plan, activate, measure) beside a possible future: the same work as three loops on three clocks around the five rights. The fast loop acts within the grant and never counts as proof; only the learning and planning loops change the mandate, and only through a logged approval; a stop reaches every loop at any time. Motion is a reading aid: speeds are relative, not to scale. Conditional design, not a forecast. Units: None; motion speeds are relative, not to scale. Scope: The section 20 baseline and the loop rules of section 26, restated for any of the three paths; no brand, agency or platform is shown running it. A052 Text equivalent: appendix "Figure text equivalents".

One cost boundary for comparing the paths

The three paths are compared on one brand-side total-cost boundary. It is a comparison rule, not a model. It sits beside the single task model in section 18. The boundary holds nine lines. Each cost is counted once, on the line where the brand pays it. The invoice may say in-house, agency or platform; the line is the same. Media and inventory are counted at the price the brand pays. On an agency route that is the net media cost plus the disclosed fee. On a principal route it is the bundled price the brand pays. The reseller's spread sits inside that price and is disclosed. It is never added again as a second line. Vendors' internal costs are left out unless they are passed to the brand. The comparison holds scope, quality and time horizon equal on every route. A route that delivers less, or worse, or for a shorter period, is not cheaper. • **ANALYSIS**

- Media and inventory, at the price the brand pays, counted once
- Agency and specialist services: fees, retainers, output prices, outcome pools
- Retained staff: the brand's own people who direct, operate, steward data and prove results
- Creative production and rights: assets, revisions, usage rights, rejected volume
- Data, technology and compute: licences, seats, usage charges, clean rooms, model and tool costs passed to the brand
- Integration and maintenance: building and keeping the joins between systems, security, change control
- Measurement and experiments: the proof function, test budgets, holdouts, brand tracking
- Financing and commitments: spend before proof, booked commitments, payment terms, float given or taken
- Exit and switching: drills, transition periods, lost state, re-licensing, retraining

The same task and the same inputs give the same total on every route. The orchestration layer can be shown as a separate subtotal where that helps a decision; it is mainly the integration line and the data and technology line. Sections 21 to 23 map each path onto these nine lines. Section 24 uses them as the full cost to count. • **ANALYSIS**

Three paths follow. They are conditional designs, not forecasts. No probability is assigned to any of them, because no defensible method exists for one in this evidence set. They coexist by task, channel, market and brand. The selector below lets a reader compare who owns each stage, what rights the brand keeps and what would show the path is working. · **SCENARIO**

The three paths are compared in the web edition's selector and in Figure 12; in print the comparison is drawn once (Figure 12, with its comparison table) and each path once, as a compact chapter view with its rights table, in its own section.

SECTION 21

A. Brand-controlled, open orchestration

A conditional model. The brand funds the decision core and contracts for usable rights. Openness is a variable, not part of the definition.

IN SHORT


The brand owns the goals, the grants and the record of choices and tests. Agencies and specialists supply strategy, craft, experiments and execution. Every right named here is a contract condition to prove, not a current entitlement. The path is feasible in parts today and unproven as a whole.

The brand owns the goals and the grants to act. It keeps a usable record of choices and tests. Agencies supply strategy, craft and skilled work. Open interfaces help join the parts where they work. But the path is defined by control, not by openness. A brand can run this core with a closed tool set, and an open ecosystem can leave a brand with weak control. This model assumes rights that must be secured. It does not claim them for every brand today. · **SCENARIO** A051

FIGURE 14 CONDITIONAL

A. Brand-controlled, open orchestration

The brand sets goals and rights, approves the plan and issues the grant. Records, proof and learning stay brand-held by contract. Conditional: it holds only if the brand contracts these rights.



Brand-controlled core: Rights and portable client state are contractual conditions, not current entitlements.
Conditional, not a forecast. Dotted green = fast checks (5 → 4) and delayed evidence (6 → 7 → 1); grey = reading order.

A. Brand-controlled, open orchestration CONDITIONAL A. Brand-controlled, open orchestration, chapter view: seven stages in the shared layout; records, proof and learning sit in a brand-held band by contract; fast checks return 5 to 4 within the grant, causal evidence runs 6 to 7 to 1, later. Conditional: rights and portable client state are contract conditions, not entitlements. The download holds the full loop. Units: none (process diagram). Scope: Conditional design; rights are contract conditions; feasible in parts today, unproven as a whole.^{[9][10][21]} A064 A066 A067 Text equivalent: appendix “Figure text equivalents”.

Scenario A: seven stages and their owners. A proposed allocation of rights, not a legal finding.

Column classifications: · SCENARIO

STAGE SCENARIO	OWNER OF DECISIONS SCENARIO	OWNER OF DATA SCENARIO	OWNER OF LEARNING SCENARIO	OWNER OF MEASUREMENT SCENARIO
1. Set goals and rights	CMO with finance and data leads	Brand controls its business context; each outside input has a licence	Brand: goals, limits and the reason for each	Finance sets success and cost boundaries
2. Plan and create	Brand lead; agency creative lead for craft	Brand gets agreed briefs, source files and use rights	Agency keeps its methods; brand keeps the brief and the decision	Tests are planned before launch
3. Approve the plan	Brand budget owner	Brand keeps grants, limits and approval history	Brand: approval history is an asset	Risk owner checks the proposed exposure
4. Run bounded work	Named specialist within the grant	Platform data remain under platform terms	Specialist learns its craft; client-specific state is contracted	Delivery checks can stop unsafe work
5. Collect signals	Brand data steward	Client log store holds permitted events and provenance	Brand: logs are the raw material of decision history	Delivery and credit reports stay distinct

STAGE SCENARIO	OWNER OF DECISIONS SCENARIO	OWNER OF DATA SCENARIO	OWNER OF LEARNING SCENARIO	OWNER OF MEASUREMENT SCENARIO
6. Test and review	Brand-funded proof team separate from sellers	Brand retains test design, results and failure history	Brand: null results kept	Causal tests and brand review can reject proxy gains
7. Keep learning; replan	Brand plan owner	Contract grants portable client state; third-party IP stays licensed	Brand: decision history and tests; vendor model state stays behind unless licensed	Later evidence updates the next plan

Feedback paths. Stage 2 sends a proposal to stage 3; stage 3 grants the scope, cap and stop right; stage 4 runs the work. Fast delivery feedback returns from stage 5 to stage 4 within the existing grant. Delayed test evidence runs from stage 6 to stage 7. Stage 7 informs stage 1 and a new plan. Any wider grant returns to stage 3 for approval. Baseline mapping: brief is stage 1; planning is stages 2 and 3; activation is stage 4; measurement is stage 5. Attribution and reporting sit in stage 5 as credit, kept distinct from stage 6 as cause. Optimisation is the 5-to-4 loop. The revised plan is 7 to 1. • **SCENARIO**

What changes

- Planning sits with the brand core and specialist teams. The core funds the joins between them.
- Buying is split by task. Principal buying may still be chosen for a defined scope, with the role disclosed per line and a comparison right.
- Reporting is one reconciled report that the brand's data steward owns. Delivery checks and credit reports stay separate documents.
- Attribution tracks credit. A separate, pre-registered test plan checks added value. The two are never blended.

Agency role and income; the group's advantage

Agencies earn for craft, service, tools and tests. A group may sell shared assets or a managed module into this core. Its edge would be skilled people, tested links, buying terms and repeatable work. It need not own the whole client system. A group that can show a working join at full cost has a product here. A group that can show only a tour of product tiles does not. • **SCENARIO**

Who integrates

The seam owner can be one of three parties: the brand's own product lead, a lead specialist appointed as integrator, or a group selling a managed module. The path is still A as long as the brand holds the goal, the stop, the proof and the exit. A lead-specialist integrator is the modular variant: a named custodian for shared logs and permissioned signals who is neither the brand nor a group. Two cautions

travel with it. Common vocabulary without working transfers is weak evidence that the modular variant works. And better interfaces lower integration cost for the largest groups too, so open standards do not imply their decline. · **SCENARIO**

CMO skills and retained duties

The CMO needs a product or operations lead, data stewardship, creative direction, finance and a strong proof function. Staff can be shared or bought, but one person must own each join. The brand keeps goals, major budget choices, the creative veto and the right to stop. Section 27 sets out the roles. · **SCENARIO**

Necessary assumptions

- The brand has enough skill, stable demand and funds to run the core.
- Suppliers grant usable exports, and the exports can be read elsewhere.
- Data use is lawful in each market where the core operates.
- Tools preserve meaning and state across hand-offs.
- An exit drill works before broad rollout.

Favourable and adverse outcomes

Favourable: the brand gets better context and choice, and can change one supplier with less rework. Adverse: the brand builds a brittle set of links. Staff spend more time fixing them than helping customers. It copies dashboards without a common goal. A supplier exits and state is lost. The brand now owns the failure while outside fees remain. · **SCENARIO**

Full costs and risks

On the nine-line boundary of section 20. Lines the brand pays directly in this path: retained staff, the decision core; integration and maintenance, the joins the core funds; data, technology and compute, its own licences and security; measurement and experiments, the proof team and the learning budget; exit and switching, drills and transition. Lines bought from outside at a disclosed price: agency and specialist services, and creative production and rights. Media and inventory are counted once at the price the brand pays, with any principal scope disclosed per line. Financing and commitments sit with the brand unless a deal moves them. Nothing is left outside the boundary; a cost that moves in-house moves between lines, it does not disappear. Open systems can raise the integration line when standards are thin or tools change often. The risk that decides this path is total cost staying higher than the managed routes at equal scope, quality and horizon, with no gain in control. · **SCENARIO**

Conflicts, failed tests, liability and the spend-to-proof gap

Conflicts: staff defend their own build; specialists optimise their own task; usage-priced tools reward more calls. Failed tests: brand-funded, from a ring-fenced learning budget set before launch. Liability: the brand carries wrong buys and outages inside its own core unless a vendor's agreement says otherwise; each agency carries its delegated execution. Spend-to-proof: the brand carries it; exposure limits per authority level bound it but do not remove it. · **SCENARIO**

Creative judgment and long-term demand

The brand's creative lead holds a veto on what is said and can reopen the brief. Brand-window measures sit in stage 6 beside the causal tests, and the proof team can reject a gain that is harvesting, channel shifting or discounting. · **SCENARIO**

What the record supports today, and what is missing

In-house skills and continued agency use coexist: 82% of 162 ANA respondents had an in-house agency and 92% still used outside agencies in 2023 ^[9]; 66% of 45 WFA/Observatory companies had one and all used outside agencies ^[10]. Protocol and SDK work supplies parts of a possible link layer ^{[21][22]}. These are building blocks, not proof of the whole model at scale. ^{[9][10][21][22]} · **ATTRIBUTED**

A064 A065 A066 A067 A068 A055

Missing proof: repeated multi-vendor exit tests, full operating cost, portable learned state, and gains across markets measured by a party without a sales interest. No reviewed source shows that an open brand core is cheapest for all firms. · **SCENARIO**

Scenario A signposts, with denominators and windows.

Column classifications: · **SCENARIO**

SIGNPOST SCENARIO	DEFINITION SCENARIO	WINDOW SCENARIO	WHAT WOULD WEAKEN THE CASE SCENARIO
Client-owned accounts	Platform and tool accounts owned by the brand ÷ all accounts used	Per year	Share falls as scope grows
Cross-vendor tests	Completed cross-vendor tests ÷ attempted	Per year	Attempts rise, completions do not
State exports	Working state exports ÷ exits attempted	Per exit	Exports delivered that the next team cannot use
Supplier change time	Median and range from decision to stable use	Per change	Time rises with each added link
Repeated gains	Tested gains after retained staff cost ÷ tests run	Per quarter	Gains vanish after full cost

Count actual migrations, not standards memberships. A logo on a membership list is not a signpost. · **ANALYSIS**

What would weaken this path

Total cost higher with no gain in quality or control; failed migrations; scarce internal skill; persistent audit gaps. Any of these would favour a managed route for those tasks. · [SCENARIO](#)

Where it does not travel

Path A needs usable exports and lawful data in each market. Where local data rights or talent are thin, A does not travel; use B or C there under common goals. A brand can run A in its home market and buy a managed service in three markets where it has no team. · [SCENARIO](#)

One CMO decision

Keep a small decision core first. Buy one bounded service through it. Compare cost and quality with a managed option. Try a supplier switch before spending more. The CMO funds tests that fail and carries spend before proof unless the deal shares that risk. Paths B and C can run bounded tasks inside this core at the same time. The choice is by task, not for the whole brand. · [SCENARIO](#)

SECTION 22

B. Agency-led, managed orchestration

A conditional model. The partner runs the service. The brand keeps reserved decisions, audit and stop rights, and agreed client state. Horizon's public design is the closest case.

IN SHORT


The agency or group runs the system and takes on service duties. The brand reserves goals, stop rights, proof and major budget choices. Principal and agency buying may coexist. Horizon's HorizonOS, Blu and Labs releases describe a proposed version of this path. They show what Horizon said, not client outcomes.

The agency or group runs the system and takes on service duties. The brand keeps agreed goals, stop rights and major budget choices. The agency joins tools, data and staff. This could work even when the brand lacks the scale to build its own core. Principal and agency buying may coexist inside one mandate, stated line by line. · [SCENARIO](#)

FIGURE 15 CONDITIONAL

B. Agency-led, managed orchestration

The agency runs the service inside a grant of scope. The brand keeps the goal, the budget, the proof and the exit, and checks the full cost. Conditional: the full-cost report is a condition to prove.



Agency-operated service: The brand reserves agreed goals, budget, proof and exit rights. Principal and agency buying may coexist, stated line by line.
Conditional, not a forecast. Dotted green = fast checks (5 → 4) and delayed evidence (6 → 7 → 1); grey = reading order.

B. Agency-led, managed orchestration CONDITIONAL B. Agency-led, managed orchestration, chapter view: the agency plans and runs the managed service in a band of its own, within a grant of scope; the brand reserves the goal, the budget, the proof and the exit, and reconciles the full-cost report. Principal and agency buying may coexist line by line. Conditional; the download holds the full loop. Units: none (process diagram). Scope: Conditional design; pieces of managed coordination exist in the record; broad autonomous authority and group-wide client gain do not.^{[39][19][7][13]} A007 A028 A029 H020 Text equivalent: appendix “Figure text equivalents”.

Scenario B: seven stages and their owners. A proposed allocation of rights, not a legal finding.

Column classifications: SCENARIO

STAGE SCENARIO	OWNER OF DECISIONS SCENARIO	OWNER OF DATA SCENARIO	OWNER OF LEARNING SCENARIO	OWNER OF MEASUREMENT SCENARIO
1. Agree the mandate	Brand goal owner; lead agency designs the service	Brand context shared under a limited purpose	Brand: goals, reserved choices and the mandate text	Both sides agree cost and outcome definitions
2. Plan and create	Agency lead; brand approves strategy	Agency tools and IP may stay agency-owned; client assets need agreed rights	Agency: reusable methods; brand: the decision and its reasons	Test plan exposes conflicts before buying
3. Grant authority	Brand reserves large moves; agency acts within scope	Both keep approval and grant records	Brand: the grant history	Risk limits are checked outside the sales team
4. Run managed work	Named agency team or agent under contract	Group uses licensed data; platform rights remain separate	Operator: integration patterns; vendor: model state	Delivery logs cover principal and agency routes
5. Report full cost	Agency supplies; brand reconciles	Brand has access to itemised fees, spend and agreed logs	Brand: the reconciled record	Attribution never blended with causal claims

STAGE SCENARIO	OWNER OF DECISIONS SCENARIO	OWNER OF DATA SCENARIO	OWNER OF LEARNING SCENARIO	OWNER OF MEASUREMENT SCENARIO
6. Challenge results	Proof function with its own mandate, budget and access	Brand keeps test results and may use another reviewer	Brand: designs, results and null results	Reviewer can challenge media the agency sells
7. Review and renew	Brand budget owner with agency counsel	Client-specific learning export is agreed; proprietary models may remain	Brand: client facts and decision history; agency: methods; vendors: model state	Delayed effects can change fees, scope or partner

Feedback paths. Stage 2 sends a proposal to stage 3; stage 3 grants the scope; stage 4 runs the managed work. Service signals return from stage 5 to stage 4 within the approved mandate. Delayed proof flows from stage 6 to stage 7. Renewal decisions return to stage 1. Changes beyond the grant pass through stage 3. Baseline mapping: brief is stage 1; planning is stage 2; activation is stage 4 after the grant in 3; measurement and reporting are stage 5. Attribution sits in stage 5 as a shared service record, kept distinct from stage 6. Optimisation is the 5-to-4 loop. The revised plan is 7 to 1. • **SCENARIO**

What changes

- Planning and reporting move into one managed process run by the partner.
- Buying mixes agency and principal routes. Each is stated in the deal, with the role disclosed per line and an audited comparison right.
- Reporting becomes a shared service record. The brand reconciles it against itemised fees, spend and logs.
- Attribution becomes part of that service record. Causal review keeps a route to challenge the record and the media choices behind it.

Agency role and income; the group's advantage

Income may mix service fees, subscriptions, data charges and principal spreads, with a limited outcome fee on top. The group's edge would be shared skill, scale, service cover and tested integration. Each income stream needs its own scope and price test. A group that bundles the streams into one price makes each of them harder to test. • **SCENARIO**

Who integrates

In B the operator integrates by definition. The variable is how far the operator's integrations are its own property. A managed module built on open interfaces lets the brand swap a vendor. A managed service built on the operator's own links does not. Either is still B. What moves a brand toward A is holding workable decision rights and portable state. What moves it toward C is an operator that depends on platform model state it cannot export. • **SCENARIO**

CMO skills and retained duties

The CMO needs a strong client lead, procurement, creative judgment, data oversight and proof. The team needs fewer operators and more people who can direct, test and challenge the partner. The brand still owns the business risk and the reserved choices. It holds an explicit stop right over the mandate. · **SCENARIO**

Necessary assumptions

- The partner can join real work, not just a pitch.
- The partner's internal rewards support the whole client scope, not the products it owns.
- Costs and conflicts are visible line by line.
- Stop rights work at the pace of the spend.
- The proof function has access and freedom.
- Exit rights cover useful client state, in named formats.

Favourable and adverse outcomes

Favourable: broad service, fewer hand-offs, shared fixed cost and one owner for repairs. Adverse: the partner steers spend to its own products. Cost cuts harm craft or service. A shared score rewards easy sales and hides long-term damage. The brand cannot rebuild the process on exit. A single interface masks several unresolved teams. · **SCENARIO**

Full costs and risks

On the nine-line boundary of section 20. Lines paid through the partner at itemised prices: agency and specialist services, the managed fee, any bundle premium and the outcome pool; creative production and rights; data, technology and compute, licences passed through at cost or a disclosed margin; integration and maintenance, the operator's, charged as a line rather than hidden in a bundle. Lines the brand still pays directly: retained staff, the directing and proof team; measurement and experiments, the proof function's own budget; exit and switching, the priced exit and the annual drill. Media and inventory are counted once at the price the brand pays; on a principal route the spread sits inside that price and is disclosed, not added again. Financing and commitments stay with the brand unless the mandate moves them. A bundle that hides any line fails the boundary test, because the brand cannot then see what it pays for. The deciding risk is supplier concentration and a harder exit. A closed component can still be good value if the terms and the result are sound at equal scope, quality and horizon. · **SCENARIO**

Conflicts, failed tests, liability and the spend-to-proof gap

Conflicts: the agency gains from more spend, more product use or less labour; the brand demands savings while adding scope. Separate pay for media selection from income from the media selected where feasible. Where not, hold a comparison right on equal inventory. Failed tests: the deal must name the payer. A partner may share the risk for a stated premium, and the record of the test still belongs to the brand. Liability: the operator carries delegated execution, integrations and incident duties. Wrong buys and missed stops are allocated in the deal. The brand still carries business risk. Spend-to-proof: the brand bears the cash-to-proof gap unless the partner explicitly accepts part of it. · SCENARIO

Creative judgment and long-term demand

The agency's creative lead leads. The brand's creative lead can overrule on brand grounds and reopen the brief. The proof function carries a brand window and can reject proxy gains, including gains from the operator's own products. · SCENARIO

What the record supports today, and what is missing

Pieces of managed coordination exist. Publicis placed Chief Client Officers across disciplines in December 2015, so the structure predates current AI [39]. Publicis describes a stock-to-bid workflow [19]. The Acxiom case names a client integration with a reported recognition result and no disclosed control [7]. Horizon's releases describe a proposed version of this path, below. None of these shows broad autonomous budget authority or group-wide client gain. [39][19][7] · ATTRIBUTED A007 A028 A029 H020

Missing proof: matched-scope client profit, true total service cost, failure rates and repeatable exit outcomes. Deal synergies and group margins do not fill those gaps. A core margin on a selected perimeter cannot separate scale, breadth and cost removal [45][45] · SCENARIO A048

Scenario B signposts, with denominators and windows.

Column classifications: · SCENARIO

SIGNPOST SCENARIO	DEFINITION SCENARIO	WINDOW SCENARIO	WHAT WOULD WEAKEN THE CASE SCENARIO
Repeated cross-service use	Accounts with repeated cross-service workflows ÷ eligible accounts	Per year	One showcase account; others unchanged
Logged authority	Delegated actions with a recorded grant ÷ all delegated actions	Per quarter	Actions found with no grant behind them
Failed hand-offs	Failed hand-offs per campaign, trend	Per quarter	Flat or rising with scope
Cost bridges	Full-cost bridges delivered ÷ bridges requested	Per quarter	Requests answered with dashboards, not bridges

SIGNPOST SCENARIO	DEFINITION SCENARIO	WINDOW SCENARIO	WHAT WOULD WEAKEN THE CASE SCENARIO
Tests of proprietary media	Tests by the proof function on the operator's own products ÷ proprietary products sold to the brand	Per year	Products sold but never tested; adverse results never acted on
Successful transitions	Exits that moved usable state ÷ exits attempted	Per exit	Renewals used as the only proof of value

What would weaken this path

No net client value after costs; proof blocked or captured; shallow tool adoption behind a wide marketplace; high group margin with declining client quality. · **SCENARIO**

Where it does not travel

Legal buying roles and local team depth change the terms by country. A global mandate cannot substitute for country evidence. A group's global platform sits on uneven local economics, so a managed mandate must be tested market by market. This review opened no primary record on local buying law. · **SCENARIO**

One CMO decision

Buy a bounded managed outcome with a clear service owner. Reserve budget, proof and stop rights. Price the exit before entry. Agree who pays for failed tests and wrong buys. Keep a skilled brand lead even if the operator supplies most of the staff. Path A can hold the core while B runs this one outcome. C can run bounded auction tasks inside it. · **SCENARIO**

The Horizon case: a proposed operating model

Horizon is in this paper as a case of a proposed operating model. It is not evidence that agencies outside the groups have won. Three issuer releases show what Horizon said on 18 December 2025, 18 June 2026 and 1 October 2026 ^{[34][13][14]}. They do not show client outcomes. The 20 supplement assessments behind this case are separate from the 72 claims in the original audit. They are not a full audit of Horizon or its partners.^{[34][13][14]} · **ATTRIBUTED** H001 H002 H003 H004

Entity roles in the Horizon record: what each release establishes and what it does not.^{[34][13][14][33]}

Column classifications: · **ATTRIBUTED** · **ANALYSIS**

ENTITY ATTRIBUTED	WHAT THE RECORD ESTABLISHES ATTRIBUTED	WHAT IT DOES NOT ESTABLISH ANALYSIS
Horizon Media Holdings	A private parent ^[34]	An operating-model classification; a capitalisation table

ENTITY ATTRIBUTED	WHAT THE RECORD ESTABLISHES ATTRIBUTED	WHAT IT DOES NOT ESTABLISH ANALYSIS
Horizon Media	The agency issuing the releases; described specialist staff [34]	Staffing depth or client access
HorizonOS	An operating system launched 18 December 2025, described with client pilots and a partner marketplace [34]	Working substitution between partners; a matched profit study
Blu	A layer announced 18 June 2026 with APIs, MCP connections and agentic buying capabilities [13]	Tested compatibility, revocation, an authority log, or an execution test by an outside party
HorizonOS Labs	A partner programme; a third cohort of 11 partners chosen from more than 600 submissions; more than 50 participants to date; next steps were client meetings and pilot scoping [14]	Live integrations, successful pilots or client gains
Individual vendors	Selected partners with their own assertions about returns, interface coverage and error detection, all unverified [14]	Any underlying method, sample, baseline or failure rate
Horizon Global	A 50/50 commercial partnership with Havas Media Network announced September 2025 [34] [33]	A merger of the parents; Havas revenue from combined billings

Four links must stay distinct in any map: an equity partnership (Horizon Global), a commercial partnership, a selected vendor (a Labs cohort) and a deployed integration (the SharkNinja audience collaboration). · INFERENCE H017 H018 A032

Evidence boundaries

- The SharkNinja example is audience collaboration with retailers. It is an issuer-reported deployment of one task [13]. It is not proof of autonomous buying, delegated budget authority or profit.
- Selection and participation counts (11; more than 600; more than 50) measure programme activity, not live adoption [14]. Pilot scoping is a stage, not a client-value result.
- A stated selection criterion, such as evidence of business outcomes, is not the underlying evidence for that criterion. The tests, samples and results were not supplied [14].
- An ecosystem described as free of vendor lock-in is an unverified commercial statement. Technical interfaces and client exit rights are separate things [14].
- API or MCP access does not show portability, shared meaning, permission enforcement, revocation or liability [13].
- Absence of inspected evidence is stated as such. Nothing here claims that a capability, contract or result does not exist.
- No Horizon principal margins, partner commissions or fees are inferred anywhere in this paper.

The hypothesis and its rival

Hypothesis: an agency earns value by finding, testing, connecting and operating useful capabilities around a client problem. Repeated experience improves its judgment and lowers its delivery cost. That is a conditional strategic hypothesis, unverified. Rival: an expanded reseller or systems-integration business that connects and resells third-party tools. It carries high ongoing coordination and maintenance cost. It creates a new dependence on the operator that runs the connections. Both can coexist with media buying and principal economics. The releases do not settle which description fits, and nothing in them disclaims buying income. The observable difference lies in the signposts below: whether reuse lowers comparable cost, whether pilots reach sustained use, and whether clients can leave with usable state. • **INFERENCE** H019


Two linked loops

The figure is an interpretive design inspired by the strategy, not Horizon's implemented architecture. It implies no instant causal measurement. The client loop runs: the brand sets the business goal, context and limits; agency and creative specialists form choices; the brand or a named delegate approves the bounded scope; operators or agents run the approved work; spend, delivery, quality and customer signals are recorded; added value is tested, decisions revised and the goal revisited. The capability loop runs: define a client gap or failed hand-off; select candidate partners against credible alternatives; approve a pilot, budget, owner and success rule; run a bounded test; judge evidence, quality, total cost and failure modes; adopt, revise or stop, and feed the result into the next gap. • **SCENARIO**

FIGURE 16 **CONDITIONAL**

Two linked loops: client and capability

A client loop and a capability loop, each gated at stage 3. Three information links join them. This is a design to test, not Horizon’s built system.



Interpretive design: not Horizon’s implemented system, not proven compounding. No arrow encodes timing or effect size.

Two linked loops: client and capability **CONDITIONAL** Interpretive design inspired by Horizon’s strategy, not its implemented architecture: a client loop and a capability loop, six stages each, drawn in the same orientation and gated at stage 3. The page opens on the client loop; the capability loop, the three cross-links (needs, evidence, options only) and a complete overview are separate views, with a step-through from client need to adopt or stop. Fast checks (5 to 4) stay inside the grant; delayed evidence (6 to 1 and 2) can change the mandate. No arrow encodes timing or effect; not proven compounding. Status: conditional. Units: none (process diagram). Scope: Interpretive design inspired by the stated strategy; not a verified Horizon workflow, contract or measured network effect.^{[34][13][14]} H020 H006 H008 H010 H011 Text equivalent: appendix “Figure text equivalents”.

The cross-links are explicit. Client needs and evidence (client stages 2 and 6) define the gap (capability stage 1). Evaluated capabilities (capability stage 6) become options in the next client plan (client stage 2), and only options: using one needs client stage 3 approval for that account. Adoption in one pilot does not grant authority across accounts. Fast operational feedback (client 5 to 4) stays inside the existing grant. Delayed causal and brand evidence (client 6 to 1 and 2) can change the mandate. New scope needs approval. Stopped pilots stay in the capability-loop denominator and useful negative results are kept. Both loops close: client 6 returns to 1 and capability 6 returns to 1. **SCENARIO**

Three mechanisms, kept separate. Each needs its own evidence; a partner count shows none of them.

Column classifications: · **ANALYSIS**

MECHANISM ANALYSIS	WHAT IT WOULD MEAN ANALYSIS	EVIDENCE IT NEEDS ANALYSIS	WHAT A PARTNER COUNT SHOWS ANALYSIS
Learning from repeated work	The agency's judgment improves with each pilot and client	Fewer failed pilots over time in matched problems; faster, cheaper selection	Nothing
Economies from reusing integrations	Fixed integration cost spread across clients	Lower comparable deployment cost and time on the second and later uses, with quality held	Nothing
Network effects	Additional participants make the system more useful to other participants	A measured benefit to existing participants from each new one	Nothing; more partners can also mean more coordination cost, security exposure and conflicting incentives

The second loop must not be called a proven network effect. A circle on a diagram is not proof of compounding value. The links may be too bespoke to reuse. New tools may add more upkeep than benefit. Successful clients may be selected for publicity. Track failures and idle integrations as well as launches. · **INFERENCE**

Horizon-case signposts, with denominators and windows. Participation counts, submissions and cohort sizes are not signposts.

Column classifications: · **SCENARIO**

SIGNPOST SCENARIO	DEFINITION SCENARIO	WINDOW SCENARIO	WHAT WOULD WEAKEN THE CASE SCENARIO
Pilot conversion	Sustained live deployments ÷ completed, eligible pilots in the same cohort; stops and pending cases reported separately	A fixed observation window after cohort completion, for example twelve months (illustrative)	Many demos, few sustained uses; stops unreported. A high rate alone is not value
Deployment time and total cost	Median and range from approved brief to stable use, on comparable scope; includes brand labour, licences, maintenance and decommissioning	Per cohort, year on year	Reuse fails to lower comparable cost or time, or quality falls
Client value	Incremental profit, quality and long-term demand on matched scopes, after discounting, channel substitution and the full operating cost	Test window plus a brand window	Gains vanish after full cost, or are substitution
Portability	A completed vendor substitution and an operator exit drill; lost state, cost, downtime and licences documented	Before scale-up; repeated annually	Useful rules or test history cannot move on the promised terms
Repeatability	Relevant repeated uses across clients or tasks within lawful data boundaries	Rolling	One promotional success that does not recur

Falsification. The advantage weakens if reuse does not lower comparable cost; if client gains vanish after full cost; if few pilots reach sustained use; if proof cannot challenge sales incentives; or if promised exits fail. The rival case must be tested honestly too. A simpler brand-controlled core or a platform-led route may work better for the same task. Do not compare a well-funded managed service with an underfunded foil. • **SCENARIO**

Classification. Horizon's public strategic design is closest to Scenario B: agency-led managed orchestration with an open partner ecosystem. It could enable Scenario A if the brand holds workable decision rights and portable state; the releases do not settle those rights. It may depend on Scenario C within particular media surfaces where platforms keep model state. It is not a fourth scenario, and this classification is not a verified client contract. What stays unavailable to this review: client contracts, fee mix, principal and partner economics, authority logs, compatibility results, failure rates, pilot conversion by cohort, and any client outcome measured by a party without a sales interest. These are stated as unavailable, not as absent. • **INFERENCE** H020 H012 H005

SECTION 23

C. Platform-dominated, fragmented autonomy

A conditional model. Platforms automate bounded tasks inside their own environments. The brand still needs a cross-platform budget view, creative direction and evidence.

IN SHORT

Local loops run fast and can work well. The cross-platform picture stays incomplete, exports can be partial, and model state stays behind. 'Fragmented' is a condition to test, not a verdict. A closed service can still deliver strong value on a defined task.

Platforms automate more execution within their own environments. The brand gives goals, assets and budget bounds. Local systems can work well while the cross-platform picture stays incomplete. The platform boundary may include third-party supply; 'closed' refers to control and state, not to who owns each impression. 'Fragmented' describes the condition in which each platform holds its own signals, credit rules and model state. It is a condition to test, not a judgment about any platform. • **SCENARIO**


FIGURE 17 CONDITIONAL

C. Platform-dominated, fragmented autonomy

Each platform runs its own local loop and keeps its own signals and model state. Goals, cross-channel tests and the budget split stay outside. Conditional: a design to test, not a forecast.

C. Platform-dominated, fragmented autonomy

conditional



Platform-local execution: Generic platform lanes, local model state, and a brand portfolio and proof layer outside the platforms. Local value can be strong where portability is limited. Conditional, not a forecast. Dotted green = fast checks (5 → 4) and delayed evidence (6 → 7 → 1); grey = reading order.

C. Platform-dominated, fragmented autonomy CONDITIONAL C. Platform-dominated, fragmented autonomy, chapter view: each platform runs and reads its own loop inside a lane; local runs, siloed signals and model state stay inside it, in words as well as colour. Goals, cross-channel tests and the budget split sit outside. Local value can be strong where portability is limited. Conditional; the download holds the full loop. Units: none (process diagram). Scope: Conditional design; bounded platform control is documented by vendors; the whole scenario is unproven.^{[15][16]} A052 A057 Text equivalent: appendix “Figure text equivalents”.

Scenario C: seven stages and their owners. A proposed allocation of rights, not a legal finding.

Column classifications: · SCENARIO

STAGE SCENARIO	OWNER OF DECISIONS SCENARIO	OWNER OF DATA SCENARIO	OWNER OF LEARNING SCENARIO	OWNER OF MEASUREMENT SCENARIO
1. Set portfolio goals	Brand with finance or an adviser	Brand owns goals and business records it can lawfully use	Brand: portfolio goals and the reasons for the split	Portfolio success is defined outside the platforms
2. Build assets and rules	Creative lead; platform setup teams	Usage rights and feeds are granted by scope	Creative lead: what the brand says; setup teams: platform craft	Content quality and fit checked before scale
3. Approve each scope	Brand or delegated agency owner	Each platform receives a bounded budget and goal	Brand: the grant history per platform	Cross-platform risk cap remains outside
4. Optimise locally	Platform within its own system	Platform retains proprietary model state; client rights depend on terms	Platform: learned state stays with the platform	Local delivery checks operate fast
5. Read siloed signals	Platform reports; brand checks	Exports can be partial or use different definitions	Brand: what it can derive from exports	Attribution in each silo can overlap

STAGE SCENARIO	OWNER OF DECISIONS SCENARIO	OWNER OF DATA SCENARIO	OWNER OF LEARNING SCENARIO	OWNER OF MEASUREMENT SCENARIO
6. Test across channels	Brand-funded team separate from media sellers	Brand retains test plans, outcomes and agreed raw data	Brand: tests, null results and substitution estimates	Tests assess substitution, overlap and longer effects
7. Rebalance; learn	Brand budget owner or adviser	Brand stores cross-platform learning it can derive; hidden state stays behind	Brand: the portfolio decision history	Delayed evidence changes platform budgets and goals

Feedback paths. Stage 2 supplies assets, feeds and rules to stage 3; stage 3 approves a bounded budget and goal per platform; stage 4 runs locally. Stage 5 feeds local changes back to stage 4 within preset rules. Stage 6 supplies delayed portfolio evidence to stage 7. Rebalancing returns to stage 3 for approval, while stage 7 informs the next stage 1 goal. No instant causal loop is assumed. Baseline mapping: brief is stage 1; planning is stage 1 for the portfolio and stage 2 for each platform; activation is stage 4 after the scope in 3; measurement is stage 5 per silo. Attribution and reporting sit in stage 5 and overlap across silos. Optimisation is the 5-to-4 loop within preset rules. The revised plan is 7 to 1. · **SCENARIO**

What changes

- Planning splits into an outside portfolio choice and local platform plans.
- Buying needs less manual work inside each silo. The platform's own system bids toward the goal it was given.
- Reports arrive fast but use different credit rules. An outside layer must reconcile them.
- Attribution overlaps across silos. The brand cannot safely add all attributed sales as unique added sales.

Agency role and income; the group's advantage

The agency earns for creative work, feeds, setup, portfolio advice and tests. A group can add scale, specialist skill and cross-market service. Its value shrinks if it merely repeats platform reports. Platform use can be an agency input as well as an agency substitute. · **SCENARIO**

CMO skills and retained duties

The CMO needs strong goals, creative craft, feed and data quality, finance and cross-platform proof. The team needs fewer bid operators. It needs someone who can reject a platform's local optimum when the portfolio or the brand suffers. And it needs a stop right over each platform budget. · **SCENARIO**

Necessary assumptions

- Platform signals are useful within their scope.
- The brand can set usable goals and constraints for each platform.
- Some tests outside the platforms remain feasible.
- Enough reporting leaves each silo to guide a portfolio choice.
- Platform incentives and data rights stay bounded by terms the brand has read.

Favourable and adverse outcomes

Favourable: strong local performance with low setup effort; a smaller brand gains skills it could not fund alone. Adverse: each platform claims the same sale. The systems harvest existing demand or overuse discounts. Fast feedback pushes spend before returns and brand effects appear. The brand loses a shared view of reach and cannot move learned state. Cheap execution becomes costly growth. · **SCENARIO**

Full costs and risks

On the nine-line boundary of section 20. Lines in this path: media and inventory, the platform spend, counted once at the price the brand pays; data, technology and compute, platform fees and any paid tools; creative production and rights, feeds and assets; agency and specialist services, any agency fee on top for setup, feeds and advice; integration and maintenance, the outside reconciliation layer; retained staff, the portfolio owner; measurement and experiments, the cross-platform proof team; financing and commitments, spend before delayed results; exit and switching, the cost of leaving each platform, including lost model state. Excluded, as on every path: the platform's own internal costs, which the brand does not pay directly and cannot see. The deciding risk is reliance on platforms with partial state and weaker cross-platform learning. A closed system can still deliver better value on a well-defined task at equal scope, quality and horizon. · **SCENARIO**

Conflicts, failed tests, liability and the spend-to-proof gap

Conflicts: platforms earn as media or paid use rises; agencies may charge on top; the brand may prefer easy reported wins. Failed tests: the brand funds most failed demand tests. Liability: platforms carry infrastructure and service duties under their own terms; the brand's profit risk stays with the brand. Wrong local actions within the platform's rules are the brand's cost. Spend-to-proof: the brand carries delayed-result risk unless a contract or refund term states otherwise. · **SCENARIO**

Creative judgment and long-term demand

The creative lead sits outside the platforms, approves assets before scale and can stop a variant engine. Portfolio proof carries a brand window and tests

substitution between channels. So a local win that moves a sale from one platform to another does not count as growth. · **SCENARIO**

What the record supports today, and what is missing

Google documents goal-led bidding across its own inventory ^[15]. Amazon's page dated 21 March 2025 describes automated buying products. Its title and body now give different stages, so the page cannot date the product's stage on its own ^[16]. These make bounded platform control feasible. Neither proves the whole scenario. ^{[15][16]} · **ATTRIBUTED** A057

Missing proof: cross-platform profit measured by a party without a sales interest, stable learning exports, and evidence of full portfolio authority. Local case results cannot settle overlap or the right split across channels. A platform-heavy route needs outside cross-market and brand measurement. · **SCENARIO** A052

Scenario C signposts, with denominators and windows.

Column classifications: · **SCENARIO**

SIGNPOST SCENARIO	DEFINITION SCENARIO	WINDOW SCENARIO	WHAT WOULD WEAKEN THE CASE SCENARIO
Automated share	Spend in bounded automated tools ÷ total media spend	Per quarter	Share rises while portfolio proof falls
Bid labour	Bid-operator hours ÷ spend	Per quarter	Hours fall but reconciliation hours rise faster
Usable exports	Platforms with a usable export ÷ platforms used	Per year	Exports partial or defined differently
Budget tests	Tests of budget changes against a credible alternative ÷ platforms with material spend	Per year	Material spend never tested outside its platform
Fee mix	Agency fees for craft and proof ÷ all agency fees	Per year	Fees track platform reports instead of judgment
Exit cost	Exit cost and data loss measured per platform	Per exit	Dependence grows without a price

What would weaken this path

The fragmented condition weakens if portable state and common proof become widely usable. The path loses appeal if tests outside the platforms show poor net value, or local gains that only move sales between channels. · **SCENARIO**

Where it does not travel

Platform mix, commerce features and privacy rules differ by market. This review did not open primary records on them. A platform-led route that works in one

market is not evidence for another. Keep the portfolio goal common and the platform mix local. · **SCENARIO**

One CMO decision

Automate a clear local task first. Keep portfolio allocation and proof outside it. Test budget changes against a credible alternative. The brand funds most failed demand tests and carries delayed-result risk unless a contract or refund term states otherwise. This can run inside an A core or a B mandate. It is a choice for a task, not for the brand. · **SCENARIO**

SECTION 24

Choose by task: make, buy or partner

A decision by task, skill, full cost and deciding risk, plus a four-cell matrix that keeps openness and client control apart.

IN SHORT


The question is not agency or in-house. It is which tasks the brand must own to keep four rights, goal, stop, proof and exit, and which it can buy at a lower full cost without losing them. This paper recommends keeping three tasks in-house whatever the brand's size; that is a design choice, not an empirical rule. The rest is a cost-and-risk comparison that differs by brand, market and task.

Start with the task, not the supplier. For each task, name the skill it needs, count the whole cost on each route, and name the risk that would decide the choice if it went wrong. Then pick a default and write down the condition that would flip it. A default is a starting position; the brand's size, market and skills can overturn it. 'Full cost' is the nine-line brand-side boundary of section 20: media and inventory; agency and specialist services; retained staff; creative production and rights; data, technology and compute; integration and maintenance; measurement and experiments; financing and commitments; exit and switching. Each is counted once, at equal scope, quality and horizon. The table is a proposal. · **ANALYSIS**

FIGURE 18 ILLUSTRATIVE

Openness and client control

Openness and client control are two questions, not one. Four configurations, described and tested, not scored. No cell is good or bad by itself.



Openness is the ability to connect and replace providers, not a vendor count. A file export alone is not a strong exit. Use the matrix per task, not per supplier.

Openness and client control ILLUSTRATIVE From section 24: openness (can a technology provider be connected and replaced?) and client control (can the brand use its rights and move its operating state?) are separate questions. Openness is not a vendor count. Four cells, described, not scored; no firm is placed; no cell is good or bad. A file export alone is not a strong exit. Illustrative. Units: none. Scope: the paper's own design. Text equivalent: appendix "Figure text equivalents".

Make, buy or partner, by task. A proposal; every default has a flip condition.^{[17][8][15]}

Column classifications: ANALYSIS

TASK ANALYSIS	SKILL NEEDED ANALYSIS	FULL COST TO COUNT ANALYSIS	DECIDING RISK ANALYSIS	DEFAULT ROUTE, AND THE CONDITION THAT FLIPS IT ANALYSIS
Set the business goal, margin floor, budget caps and customer exclusions	Judgment about the business; finance; legal	Senior time; finance model; legal review	A wrong goal scales the wrong action and cannot be recovered by delegation	Make; this paper's recommendation for every brand. A partner may advise; the brand signs.

TASK ANALYSIS	SKILL NEEDED ANALYSIS	FULL COST TO COUNT ANALYSIS	DECIDING RISK ANALYSIS	DEFAULT ROUTE, AND THE CONDITION THAT FLIPS IT ANALYSIS
Brand strategy and creative direction	Distinctiveness, taste, category knowledge, rights	Creative lead; research; rights clearance; the cost of rejected volume	Cheaper versions of a weak idea; loss of distinct brand meaning	Make the direction; buy the craft. Flip to buying the direction only for a small brand with no creative lead, with an explicit overrule right.
Creative production at volume	Production craft, AI tooling, review	Per approved asset including review, legal and rework; never per prompt	Volume targets reward needless output	Buy or partner. Flip to make when volume is steady and review cost is the largest line.
Customer data stewardship and consent	Data engineering, privacy law, provenance	Stewards; tooling; audits; reviews per jurisdiction	Unlawful use; loss of exportable records	Make the stewardship; partner for infrastructure. Platform data stays under platform terms either way.
Identity resolution and data collaboration	Match engineering; clean-room operation (shared data matching in which neither party sees the other's raw records)	Licence; match-quality tests; the cost of leaving the vendor	Owner neutrality promised, not proven ^[17] ; profiles are not people	Partner, with rights to leave. Flip to buy only if the brand runs collaboration at scale across many partners.
Media planning across platforms	Portfolio allocation, negotiation, category experience	Planners; data; mix models	Local optima; platform incentives (path C)	Buy from an agency, or make a small core. Flip to make when the brand holds a strong proof function and the agency sells the media it plans.
Negotiated buying (upfronts, private marketplaces, retail media deals)	Scale, relationships, appetite for inventory risk	Fees or spreads; inventory risk; payment terms; audit	An undisclosed role; a cheap package the brand cannot assess	Buy, with the role disclosed per transaction. Principal buying is acceptable for a defined scope with an audited comparison ^[8] . Flip to make for a retailer or platform-native brand with its own deal desk.
Auction buying inside one platform	Platform setup, feed quality, creative variants	Setup labour; feeds; platform fees	Harvesting and channel shift; hidden model state	Partner with the platform (path C) for execution; keep allocation and proof outside ^[15] .
Agent orchestration layer (workflow, permissions, logs)	Systems integration, security, product management	Integration; maintenance; security; decommissioning; brand labour	Brittle links; staff fixing interfaces; operator dependence	Partner (path B) unless the brand has a product team. Flip to make when an exit drill shows the managed route cannot return usable state.
Measurement, incrementality and brand tracking	Experiment design, modelling, scepticism	Proof team; test budgets; data access; null results kept	Proof captured by the party that places the spend	Make the mandate and the budget; buy the method. The reviewer must be able to test media the supplier sells.
Reporting and reconciliation	Data engineering, finance	Tooling; reconciliation labour; definitions	Attribution blended with causal claims	Buy the pipes; make the definitions.

TASK ANALYSIS	SKILL NEEDED ANALYSIS	FULL COST TO COUNT ANALYSIS	DECIDING RISK ANALYSIS	DEFAULT ROUTE, AND THE CONDITION THAT FLIPS IT ANALYSIS
Contracting, audit and exit	Procurement, legal, technology	Negotiation; audit fees; exit drill cost	Rights assumed, not granted	Make. Run the exit drill before signing.

This paper recommends that three tasks stay in-house whatever the brand's size: setting the goal and its limits; holding the measurement mandate and its budget; and owning the contract, audit and exit. Everything else is a comparison. The surveys support coexistence; this allocation of rights is the paper's recommendation, not a finding. Most respondents in the ANA and WFA samples hold some skill inside and still buy from agencies ^{[9][10]}. In-housing does not mean doing more tasks, and keeping accountability does not require employing everyone who performs the work. It means keeping the decisions that define the business and buying the execution that suits the brand's scale. ^{[9][10]} · **INFERENCE**

A064 A066 A067

Count full cost honestly. A fee comparison that stops at the invoice misses retained staff, tool licences, review time, security, integration and the cost of leaving. An in-house build that saves fees but needs three engineers to keep the links alive is not cheaper. A managed service that bundles tools at a premium is not dearer if it removes those three roles and the exit drill works. Price the exit before entry on every route. · **ANALYSIS**

Watch the deciding risk. For creative direction it is a weak idea produced cheaply at volume. For negotiated buying it is an undisclosed role. For orchestration it is a brittle set of links the brand owns and cannot fix. For measurement it is proof captured by the party that spends. The route that keeps the deciding risk lowest at an acceptable full cost wins, even when it is not the cheapest on the invoice. · **INFERENCE** A047

Openness and client control are two questions, not one

Two questions must stay separate. Openness: can a technology provider be connected and replaced? Client control: can the brand exercise its rights and move useful operating state to another operator? The matrix below has no scores and plots no firm. Openness, ownership and neutrality are distinct. An open ecosystem can have weak client exit. A closed one can have a sound export. Neither settles who gains from the media chosen, and neither says whether the work is good; the loop in section 19 does that. · **ANALYSIS**

Four cells, described, not scored. No named firm is placed in any cell.

Column classifications: · **ANALYSIS**

ANALYSIS	WEAK CLIENT CONTROL ANALYSIS	STRONG CLIENT CONTROL ANALYSIS
Open: providers can be connected and replaced	Many vendors, one operator the client depends on. Vendor swaps are easy; operator exit is not. Open entry for vendors does not by itself give a client an easy exit.	Broad choice and strong exit rights. Can be costly to run, because the brand pays for the joins and the drills; this review holds no comparable cost cases, so no cell is ranked by cost.
Restricted: providers are hard to connect or replace	Limited choice and weak control. Can still deliver good value on a bounded task at a good price, if the brand knows what it is giving up.	A tightly managed tool set with strong records and exit rights. Closed does not mean captured.

Use the matrix per task, not per supplier. A brand can sit in the top-right cell for its decision core, the bottom-left cell for one platform's auction buying, and the top-left cell for a managed service it has chosen to depend on for now. The paths in sections 21 to 23 coexist inside one brand. This table decides where each applies. · ANALYSIS

SECTION 25

A commercial design

A proposal, not a real private contract: scope by task, fees in illustrative units, conflicts, test funding, proof rights, exit terms, and who owns which kind of learning.

IN SHORT

No term here is drawn from any reviewed agreement and none is attributed to any named group or to Horizon. The design does not make an agency neutral; it makes the conflict visible and testable. It does not guarantee client value; it makes value measurable and the cost of failure assignable.

A PROPOSAL, NOT A REAL PRIVATE CONTRACT

No term below is drawn from any reviewed agreement. None is attributed to any named group or to Horizon. Numbers are illustrative units for arithmetic, not an estimate of market prices.

Scope, by task

Scope is written by task, not by discipline, using the table in section 24. For each task the contract lists who decides, who acts, what data is used, which authority level applies (section 16), and what the brand can take on exit. Four scope classes

do the work. Reserved to the brand: goal, margin floor, caps, exclusions, consent rules, measurement mandate, final supplier approval, stop right. Delegated within a recorded grant: planning, creative production, platform setup, bounded optimisation, reporting. Shared: test design (brand approves, agency proposes) and creative direction (agency leads, the brand's creative lead can overrule). Excluded: anything that moves material budget across platforms without approval; any use of another client's confidential data; any principal transaction not disclosed per line. · **ASSUMPTION**

Fees

Fee components in illustrative units, base fee 100. Shares are assumptions for arithmetic, not market rates.^[8]^[51]

Column classifications: · **ILLUSTRATIVE** · **ANALYSIS**

COMPONENT ILLUSTRATIVE	BASIS ILLUSTRATIVE	ILLUSTRATIVE SHARE OF 100 ILLUSTRATIVE	WHY IT IS SEPARATE ANALYSIS
Base service fee	Named roles and service levels, not hours	45	Funds essential staff and the test programme; protects craft from a race to volume
Output prices	Per accepted asset class, with quality and revision rules agreed first	20	Prevents cheap variants from inflating output
Technology and licence charges	Itemised; pass-through at cost or a disclosed margin	15	Stops overlapping licences and hidden bundle premiums
Media route charge	Agency route: a disclosed fee. Principal route: the role disclosed per line and an audited comparison right	10	Keeps the commercial role visible ^[8]
Bounded outcome pool	Paid only on agreed incremental profit over a matching window, with caps, refunds and a check by a party without a sales interest	10	Rewards added profit, not attribution ^[51]

Conflicts

- The agency discloses, per line, whether it acts as agent or principal, and whether it or an affiliate earns from any recommended product, whether media, data or software.
- Pay for media selection is separated from any income from the media selected where feasible. Where not feasible, the brand holds a comparison right on equal inventory and terms.
- Partner subsidies, commissions or revenue shares that could sway tool selection are disclosed before a pilot is approved.
- The agency may reuse general methods across clients. It may not pool this client's data or transfer learned state derived from it without written consent.

Test funding

- A ring-fenced learning budget (illustratively 5% of the annual fee plus an agreed media test allowance) is set before launch.
- Failed tests are paid from that budget. A failed test is logged as learning and does not trigger a fee penalty.
- The agency may share pilot risk for a stated premium. If it does, the record of the test still belongs to the brand.
- Spend committed before proof arrives is bounded by exposure limits per authority level. The brand carries that spend unless the contract moves part of it to the agency. A cap is not a liability ceiling, because booked commitments and a failed stop can exceed it.

Proof rights

- The brand names a proof function, internal, external or both. It has access to itemised fees, spend, logs, raw permitted data, model and tool versions, and test designs.
- The proof function may test media the agency sells. It may publish adverse results to the brand's leadership. Its findings change allocation.
- Success criteria are frozen before each test; changes are logged. Null results are kept.
- Attribution reports and causal estimates are delivered as separate documents and never blended.

Exit terms

- Export rights cover accounts; audience definitions; creative source files and usage rights; briefs; decision logs; approval history; test designs and results; settings and rules; and any transferable agent state, each in a named format.
- Third-party licences and agency property are listed with their post-exit status and price. Portable model weights are not promised.
- An exit drill, an export plus a rebuild of one workflow elsewhere, is run before the contract passes its first year and repeated each year. Its cost, downtime and lost state are documented.
- Exit assistance is a named transition period, staffed and priced in advance. Retention and deletion schedules are fixed.

Ownership and liability

A proposed ownership-and-liability table. Explicitly not Horizon's terms and not any reviewed contract.

Column classifications: · **ANALYSIS**

PARTY ANALYSIS	OWNS OR DECIDES ANALYSIS	DUTY ANALYSIS	PROOF AND EXIT TEST ANALYSIS
Brand	Business goals, reserved choices, lawful customer context, stop rights	Funds spend, retained staff and tests unless the contract shifts them	The brand can revoke authority; all active agents stop or fall back safely
Operator (agency)	Service quality, integrations, delegated execution	Incident duties, repair of partial actions, logs, exit assistance	A completed vendor substitution and an operator exit drill with documented loss
Vendors	Their own software and licensed assets	Service under their actual agreements; no transfer of model weights promised	Licence terms priced for post-exit use
Proof function	Method, test calendar, null results	Access to costs, logs and tests; freedom to challenge the operator and the media sellers	Safeguards shown to permit an adverse result; any adverse result found is acted on

Four kinds of learning

Four kinds of learning and the terms each needs. Owning the first does not grant the last.

Column classifications: · ANALYSIS

KIND ANALYSIS	EXAMPLE ANALYSIS	DEFAULT OWNER UNDER THIS PROPOSAL ANALYSIS	EXPORT, LICENCE, RETENTION, DELETION AND REBUILD TERMS ANALYSIS
Client facts and decision history	Goals, caps, exclusions, approvals, actions taken	Brand	Named export format; retained for the agreed period; deleted on schedule; sufficient to rebuild the decisions
Experimental designs and results	Holdout designs, outcomes, null results	Brand, with the agency able to cite anonymised method	Exportable; the record of a failed pilot belongs to the brand
Reusable agency methods	Selection rubrics, playbooks, integration patterns	Agency	Licensed, not transferred; no client data inside
Proprietary vendor model state	Weights, learned parameters, identity links	Vendor	Not portable by default; priced if available; listed as left behind on exit

Reuse of general methods does not permit pooling client data or transferring confidential learned state. Other clients' confidential information is protected on exit and in reuse. The rights needed to rebuild a decision are the first two kinds plus the settings and rules that produced it. A brand that holds only the first kind has files, not decisions. · ANALYSIS

FIGURE 19 ILLUSTRATIVE

What leaves with the brand

Four kinds of asset leave on four kinds of terms. Owning the records does not grant the methods or the models. Every right here is a contract condition.

What leaves with the brand

proposal, not a reviewed contract

Current operator: who holds each asset class today under the proposal, and on what terms.
Every right here is a contract condition. Owning the records does not grant the methods or the models.

	Brand records	Client-specific rules and state	Agency methods and IP	Platform and vendor model state
	Owner: Brand	Owner: Brand, where the contract says so	Owner: Agency	Owner: Vendor or platform
Access	Direct: the proof function reads the logs, not the operator's dashboard	Settings, rules and tool versions, through the proof function	No access right by default	Under the vendor's terms; client rights depend on them
Export	Named formats for accounts, audiences, creative, briefs, logs, approvals and tests	Settings, rules and transferable agent state in named formats; the rest is left behind	Not exported	Not portable by default; portable weights are not promised
Reuse	The brand reuses freely; the agency may cite anonymised method	Successor reuses the rules; this client's learned state moves only with written consent	The agency reuses general methods across clients; no client data pooled inside	Not reusable by the brand; the platform keeps its learned state
Licence	Not needed: brand-owned under this proposal	None for the rules; non-transferable learned state is priced or left behind	Licensed, not transferred; listed with post-exit status and price	Priced if available; terms priced for post-exit use
Retention and deletion	Agreed retention period; deleted on schedule; a failed pilot's record stays with the brand	Schedules fixed in the contract	The agency's own	The vendor's own schedule
Rebuild	Rebuilds the decisions only together with the settings and rules	Records plus rules are what a rebuild needs; files alone are not decisions	The successor rebuilds its own methods or licenses the incumbent's	Left behind on exit; a successor relearns from the brand's records and rules

A file export alone is not a successful exit; the settings and rules must move with the records.
An exit drill, an export plus a rebuild of one workflow elsewhere, is run before the contract passes its first year and repeated each year. Its cost, downtime and lost state are documented.
A proposal, not a reviewed agreement; no term is attributed to any named group. No exit drill result exists in the record.

What leaves with the brand **ILLUSTRATIVE** Four asset classes (brand records; client-specific rules and state; agency methods; platform model state) against six exit conditions: access, export, reuse, licence, retention and deletion, rebuild. A successor view shows what moves, what moves if contracted, what must be licensed or rebuilt, and what is left behind. Every right is a contract condition; a file export is not an exit. Units: none (a grid of stated conditions). Scope: A proposal from sections 25 and 27; no term is drawn from any reviewed agreement or attributed to any named group or to Horizon; rights are contract-dependent and no exit drill result exists in the record. A051 Text equivalent: appendix "Figure text equivalents".

The economics questions, without invented numbers

For any agency-led model, including the Horizon case, the contract and the proof function must answer six questions. No reviewed release answers any of them for Horizon. Which income is strategy, setup, managed operation, analytics, licence pass-through, media income and outcome component? Each needs its own scope and price test. Who funds a pilot that fails, and who owns the resulting test record? Who captures the saving when an integration is reused: the agency as margin, or the client as a lower price? Does a partner subsidy, commission or bundle distort which vendor is selected? Who carries committed spend while proof is delayed? How much retained brand skill is needed to direct and challenge the operator, and is that cost in the comparison? Any model built on these questions is a proposal, not a finding about any firm's contracts. • **INFERENCE**

What this design does not do. It does not make the agency neutral; it makes the conflict visible and testable. It does not guarantee client value; it makes value measurable and the cost of failure assignable. It does not prove the path in which it sits. It is a condition for A or B to deliver what they promise. • **ANALYSIS**

SECTION 26

One retailer, three agents

An illustrative design. A media agent, an offer agent and a retention agent pursue conflicting goals. A shared profit goal, hard limits, clear precedence and a human stop resolve them.

IN SHORT

The retailer, the agents, the numbers and the outcomes are invented for the argument. It is not a Horizon client result and not an observed case. It shows how conflicts can be resolved without instant causal proof. It does not show that any vendor's agents work this way or that the mandate would be profitable.

ILLUSTRATIVE DESIGN ONLY

The retailer, the agents, the numbers and the outcomes are invented. This is not a Horizon client result and not an observed case. No universal profit floor is implied and no instant causal measurement is assumed.

A retailer wants more new customers without harming profit. Three agents act in its system. A media agent seeks reach among likely new customers. An offer agent seeks conversion through deeper discounts. A retention agent favours current buyers because they respond cheaply. Left alone, each wins its own metric and loses the business goal. Reach buys attention that does not convert. Discounts convert at a loss. Retention harvests demand that was already there. · **ILLUSTRATIVE**

The mandate. Illustrative settings for one invented retailer; a floor is a setting for this retailer's economics, not a rule for all.

Column classifications: · **ILLUSTRATIVE**

ELEMENT ILLUSTRATIVE	ILLUSTRATIVE SETTING ILLUSTRATIVE	WHO SETS IT ILLUSTRATIVE
Common goal	Incremental contribution profit from new customers over 13 weeks, after discounts, returns, fulfilment and media; the brand-health measure must not fall below its 12-week baseline	CMO with the CFO
Time horizon	Fast loop: daily delivery checks. Learning loop: a 13-week holdout test, with pre-specified interim reads at the end of weeks 4 and 8 that release budget tranches but do not change the mandate. Planning loop: 12-month cohort value and brand tracking	CMO; proof lead
Margin floor	No offer may take first-order contribution below 8% of net sales for the targeted segment	CFO
Budget and exposure limits	\$400,000 media over 13 weeks as a contingent maximum, released in three tranches: \$90,000 for weeks 1 to 4; up to \$150,000 for weeks 5 to 8; up to \$160,000 for weeks 9 to 13. Each later tranche is unlocked only by the pre-specified interim read for the prior tranche, signed by the budget owner on the proof lead's reading. If a tranche is spent before its read, spending pauses until the read. No agent can release a tranche. No agent may move more than 10% of the weekly budget between channels without approval, counted on a rolling seven-day window across all agents and including pending commitments. Discount budget capped at \$120,000 face value. Maximum spend before the first interim read: \$90,000	Budget owner

ELEMENT ILLUSTRATIVE	ILLUSTRATIVE SETTING ILLUSTRATIVE	WHO SETS IT ILLUSTRATIVE
Interim evidence gates	Week 4 read: added new-customer contribution on the first three weeks' holdout is positive on the pre-registered rule, and no exclusion breach has occurred. Week 8 read: the same on weeks 1 to 7, plus the brand-health measure holding. A passed read releases the next tranche. At a failed interim gate, pause new commitments and queued actions that can still be stopped. Hold the next tranche and return the plan to the CMO and budget owner. Unspent budget does not carry execution authority forward automatically. Resume only under a logged human decision consistent with the mandate. Keep existing non-cancellable commitments visible. Neither read changes the mandate; only the 13-week and 12-month reads can. Interim reads are gates for exposure, not proof	Budget owner decides; the proof lead supplies the read
Delayed or inconclusive reads	A delayed read is not a failed result, but it does not extend spending authority. The same holds for an inconclusive read. Any permitted continuation is specified in advance, with a time limit and an exposure cap. Actions already taken stay logged as executed. Illustrative, week 4: up to 7 days after the read was due, from what is left of tranche 1, capped at \$22,500 (one week's even share); the next tranche stays held. Week 8: none; new execution stays paused until the read arrives or the CMO and budget owner approve a logged exception with its own spending cap, expiry and rationale; tranche 3 stays held	Budget owner
Customer exclusions	Customers whose last purchase was within the 12 months before the send date are excluded from acquisition offers; under-18s excluded from targeting; suppressed and opted-out records excluded everywhere; local rules differ by market and are checked by legal	Data steward; legal
Permitted data use	First-party purchase history for exclusion only; platform signals within platform terms; no sensitive categories; no pooled data from other clients; retention 13 months; deletion on exit	Data steward

Which decision wins

Precedence is fixed before launch and logged. First, constraints beat objectives: exclusions, consent rules and the margin floor are hard, and an agent that cannot satisfy them does not act. Second, the common goal beats any agent's metric: a proposed action is scored on expected added contribution from new customers, not on reach, conversion or response. Third, ties go to the cheaper reversible

action: a media shift that can be undone tomorrow beats a discount that trains customers to wait. Fourth, uncertainty raises the approval level: an action whose expected effect rests on an unproven signal is capped at a smaller exposure until the holdout test (a group held back from the campaign for comparison) reads. · **ILLUSTRATIVE**

A worked clash, invented, with the unit economics stated so the decisions can be reproduced. The retailer's targeted segment has an average order of \$100 at list price, cost of goods of \$60, fulfilment of \$10, and expected returns of 5% of orders, each costing \$6 to handle. First-order contribution per order, before media, is $(1 - r) \times (P \times (1 - d) - C - F) - r \times H$, where P is the list price, d the discount, C the cost of goods, F fulfilment, r the return rate and H the handling cost of a return; net sales are $(1 - r) \times P \times (1 - d)$. On a returned order the formula treats the goods and the fulfilment cost as recovered, so only the handling cost H remains; that is a simplification of this illustration, not a rule. The floor is 8% of net sales. With no discount the ratio is 29.7%. In week three the offer agent proposes a 25% code to lift conversion: contribution falls to \$4.45 on net sales of \$71.25, a 6.2% ratio, below the floor, so the code is rejected and a 12% code is proposed instead, at \$16.80 on \$83.60, a 20.1% ratio, within the floor. Media cost sits in the common goal, not in the floor. The media agent proposes shifting 15% of the weekly budget into a broad-reach video buy. The 10% limit applies to the sum of all shifts between the same channels by all agents in the past seven days, pending commitments included. If nothing else has moved, 10% executes and the remaining 5% goes to the budget owner for approval as one request; if any part of the limit is already used, the whole proposal is held for the budget owner, so no agent can split a move to pass under the limit. The retention agent proposes a lapsed-buyer email. Its list holds customers whose last purchase was between six and twelve months before the send date. Every one of them bought within the past twelve months, so all are inside the exclusion; the email is blocked for this programme and routed to the retention programme's own budget. A customer whose last purchase was more than twelve months ago is outside the exclusion; the mandate counts such a customer as a reactivation, reported separately from new customers at the holdout read. · **ILLUSTRATIVE**

Discount arithmetic for the worked clash. Illustrative inputs for one invented retailer; the formula is the point, not the values. A returned order is treated as recovering its goods and fulfilment cost; only the \$6 handling cost remains.

Column classifications: · **ILLUSTRATIVE**

ELEMENT ILLUSTRATIVE	ILLUSTRATIVE SETTING OR RESULT ILLUSTRATIVE
List price per order (P)	\$100
Cost of goods (C)	\$60
Fulfilment (F)	\$10
Expected returns (r) and handling cost (H)	5% of orders; \$6 per returned order
Contribution per order before media	$(1 - r) \times (P \times (1 - d) - C - F) - r \times H$; net sales = $(1 - r) \times P \times (1 - d)$
Floor	8% of net sales for the targeted segment

ELEMENT ILLUSTRATIVE	ILLUSTRATIVE SETTING OR RESULT ILLUSTRATIVE
No discount	\$28.20 on \$95.00 net sales: 29.7%
12% code	\$16.80 on \$83.60: 20.1%, within the floor
25% code	\$4.45 on \$71.25: 6.2%, below the floor; rejected

FIGURE 20 ILLUSTRATIVE


One retailer, three agents

Three agents chase three metrics. A shared profit goal, hard limits, fixed precedence and a human stop decide what runs. Invented numbers; a design to test, not an outcome.

Three proposals, one policy gate

Invented numbers: a design to test, not an outcome

Offer agent proposes: a 25% code to lift conversion, in week three.



Outcome: rejected · Constraints beat objectives: the margin floor is hard, so the agent does not act. Decided by: The rule itself (the floor is set by the CFO); no human decision is needed to reject.

Nothing executes; the refusal is logged.

Available authority now: tranche 1: \$90,000, released at start; no usage figure is stated, the walkthrough starts at zero used.

Budget tranches: released only by a signed read; no agent can release one



A failed read pauses new commitments and queued actions that can still be stopped, even with budget left: the next tranche is held and the plan returns to the CMO and budget owner until a logged human decision. A delayed or inconclusive week 4 read allows up to 7 days from what is left of tranche 1, capped at \$22,500 (illustrative); week 8 has none. A spent tranche pauses spending until its read.

contribution per order = $(1 - r) \times (P \times (1 - d) - C - F) - r \times H \cdot \text{net sales} = (1 - r) \times P \times (1 - d) \cdot \text{floor } 8\% \text{ of net sales}$

P \$100 · C \$60 · F \$10 · r 5% · H \$6 → no discount \$28.20 on \$95.00 (29.7%) · 12% code \$16.80 on \$83.60 (20.1%) · 25% code \$4.45 on \$71.25 (6.2%, below the floor)

Assumed, as the paper states (6 October 2026): a returned order adds nothing before handling; goods and fulfilment are treated as recovered, only the handling cost remains. Media cost sits in the common goal, not in the floor. A floor is a setting for this retailer, not a rule for all.

Invented numbers. No sales or vendor outcome is shown. Reactivations are reported separately from new customers.

One retailer, three agents **ILLUSTRATIVE** An invented retailer’s mandate, walked through: three proposals enter a policy gate and leave eligible, rejected or escalated. Eligible is not permitted: an offer needs sign-off and an available tranche. A failed interim read pauses new execution even with budget left. A 25% code fails the 8% floor; a 12% code waits; a 15% shift exceeds the 10% limit. Illustrative. Units: USD and shares, illustrative; weeks. Scope: An invented mandate and worked clash from section 26; a design to test, not an observed case, a vendor result or a universal rule. Text equivalent: appendix “Figure text equivalents”.

Who can override, and when.

Column classifications: · **ILLUSTRATIVE**

OVERRIDE ILLUSTRATIVE	WHO ILLUSTRATIVE	WINDOW ILLUSTRATIVE
Stop all agents	CMO, budget owner, or the proof lead on a data-use breach	Immediate; all connected tools, queues and booked commitments checked
Lift a cap for one action	Budget owner, with finance sign-off above an illustrative \$50,000	Before execution
Release the next budget tranche	Budget owner, on the proof lead's signed interim read; never an agent	At the week 4 and week 8 gates only
Change the margin floor	CFO only	Only at a planned review, never mid-test
Reject creative or an offer on brand grounds	Creative lead	Any time, before and during the run
Change exclusions or data use	Data steward with legal	Before execution; logged

No agent overrides a human. No human overrides the exclusions or consent rules by exception. · **ILLUSTRATIVE**

How logs are kept

Each action record holds the input signals and their source; the model or tool version; the rule or score that chose the action; the proposed exposure; the approver, human or grant id; the time; the result; and any later correction. Logs are exportable in a named format and held for the retention period. The proof function reads logs directly, not through the agency's dashboard. A log that shows an action taken on a provisional signal must keep that fact, so later proof can reverse the plan. · **ILLUSTRATIVE**

When later evidence changes the mandate

At the week 4 and week 8 interim reads: a passed read releases the next tranche; a failed read pauses new commitments and queued actions that can still be stopped, even if budget remains, holds the next tranche and returns the plan to the CMO and budget owner until a logged human decision; a delayed or inconclusive read allows only the continuation set in advance in the mandate. A delayed read is not a failed result, but it does not extend spending authority. There is no continuation after week 8: new execution stays paused until the read arrives or the CMO and

budget owner approve a logged exception with its own spending cap, expiry and rationale, and tranche 3 stays held. None of these changes the mandate. At the 13-week holdout read: if added new-customer contribution is positive and the brand-health measure holds, the budget owner may raise the weekly shift limit. If the gain came from discounts to customers who would have bought anyway, the offer agent's authority falls one level and the discount cap is cut. At the 12-month cohort read: if customers acquired with deep offers show low repeat value, the margin floor review moves the floor up, and the offer agent's objective changes from first-order conversion to cohort contribution. At any time: a data-use breach stops the programme; a creative-quality failure returns the brief to the creative lead. The fast loop never changes the mandate; it changes execution within it. Only the learning and planning loops change the mandate, and only through a logged approval. · **ILLUSTRATIVE**

Creative quality and brand effects

The creative lead approves every acquisition asset before scale. The lead can reject an offer that cheapens the brand even when it clears the margin floor. The brand-health measure sits in the common goal, so a programme that acquires customers while the measure falls has failed its own mandate. Deep offers that recruit customers who wait for the next discount show up at the 12-month read, not the 13-week read. That is why the planning loop exists. · **ILLUSTRATIVE**

What this example shows and does not show. It shows how a shared profit goal, hard constraints, exposure limits and clear precedence can resolve conflicts between agents without instant causal proof. It does not show that any vendor's agents work this way, that the numbers are typical, or that the mandate would be profitable. It is a design to test, not an outcome. A simulated 13-week schedule spends at most the released tranches: \$90,000 by the end of week 4, \$240,000 by the end of week 8 and \$400,000 by the end of week 13, and never more than the authorised exposure on any date. · **ANALYSIS**

SECTION 27

The retained team

Roles, decision rights, escalation and stop rights, plus the charter of the proof

function.

IN SHORT

A proposal. Staffing depth depends on brand size. Roles can be shared, part-time or bought, but each must have a named owner. The brand needs fewer operators and more people who can direct, test and challenge a partner. One possible small-brand arrangement is three accountable people: one for goal and budget, one for data and proof, one for creative; execution can be shared, fractional or bought, as long as each right has a named owner.

Roles, what each owns, what it never delegates, what it delegates, and its stop right. A proposal.

Column classifications: · ANALYSIS

ROLE ANALYSIS	OWNS ANALYSIS	KEEPS, NEVER DELEGATED ANALYSIS	DELEGATES ANALYSIS	STOP RIGHT ANALYSIS
CMO	The goal, the operating model, the partner mandate	Goal, margin floor, reserved choices, renewal and exit decisions	Execution by task per section 24	Yes: stop any programme
Budget and risk owner (finance partner)	Caps, exposure limits, the cash-to-proof gap	Approval of material budget moves; liability allocation	Within-grant moves to operators	Yes: stop spend above a cap
Product or operations lead	The joins between systems, tools and partners	Integration standards, change control, the exit drill	Build and maintenance to vendors or the agency	Yes: stop a failing integration
Data steward	Lawful use, provenance, consent, exports	Permitted-use rules, retention, deletion, export formats	Data engineering to vendors	Yes: stop on a data-use breach
Creative lead	Brand meaning, distinctiveness, fit, rights	Approval of what the brand says; reopening the brief	Production craft to agencies	Yes: reject creative or an offer on brand grounds
Proof lead	Measurement mandate, test calendar, null results	Success criteria, access rights, the right to publish adverse results	Methods and tools to specialists	Yes: stop on a proof breach, such as criteria changed after results
Procurement and contract owner	Scope, fees, conflicts, audit, exit terms	Disclosure requirements; comparison rights	Negotiation support to advisers	No stop right; escalates
Security and legal	Permissions, identity, revocation, jurisdiction	Grant issuance and revocation policy	Implementation to vendors	Yes: revoke any grant

Decision rights by authority level (section 16). Level 6 is not granted under this design until level 5 has passed two review cycles and a causal read by a party without a sales interest.

Column classifications: · ANALYSIS

AUTHORITY LEVEL ANALYSIS	WHO APPROVES ANALYSIS	WHO CAN RUN IT ANALYSIS	WHO REVIEWS ANALYSIS
1. Recommend or make content	Creative lead for content; any owner for recommendations	Agency or tool	Creative lead; proof lead samples
2. Human-approved workflow	The task's owner	Agency or tool	Proof lead samples; logs retained
3. Execute preset rules	Budget owner sets the rules; data steward sets the data limits	Operator or agent under a grant	Weekly exception review
4. Adapt within bounds	Budget owner and product lead jointly	Agent within bounds	Proof lead at each read; exception review
5. Coordinate across systems	CMO, budget owner, product lead and security	Operator under a tested grant	Monthly review; exit drill annually
6. Move material budget alone	CMO and CFO jointly, with a loss budget, and only after the conditions above	Agent under a tested grant	Proof by a party without a sales interest; board sponsor informed

Escalation

1. Exception in the fast loop (pacing breach, delivery fault, blocked action): the operator fixes it within the grant; logged; reviewed weekly.
2. Cap or limit reached: the action pauses; the budget owner decides within a window set by exposure, reversibility and measurement lag, for example one business day for a reversible media shift (illustrative); logged.
3. Conflict between agents or teams not resolved by precedence rules: the product lead convenes the task owners; the CMO decides if unresolved within the window set for that task by its exposure and reversibility; two business days is the illustrative default.
4. Data-use, consent or jurisdiction question: the data steward and legal decide; the programme pauses on a breach.
5. Proof dispute (the agency and the proof lead disagree on a result): both reports go to the CMO unchanged; the proof lead's criteria stand unless changed at a planned review.
6. Partner failure (missed service level, failed stop, failed export): procurement triggers the remedy; the exit drill result decides whether to escalate to renewal or exit.
7. Timing rule for every window above. Three things set it: the exposure at stake, how far the action can be reversed, and how long the evidence takes to arrive. A reversible, low-exposure action can wait for the weekly review. An irreversible, material one escalates the same day.

The proof function's charter

The proof function is defined by decision rights, not by job title. Neither an in-house analyst nor an outside reviewer is separate from the spender just by being called so. The charter: an agreed test calendar; success criteria frozen before each test, with changes logged so an optimiser cannot redefine success; null results

kept; direct access to costs, logs, raw permitted data and tool versions; a client-appointed reviewer with authority to reject a result; the right to test media the supplier sells; and the right to publish an adverse result to the brand's leadership. A fast signal may trigger a small bounded change within a grant. It never counts as proof. · **INFERENCE**

What the brand keeps on exit

The decision history, approval history, test designs and results, exclusion and consent rules, creative source files with rights, account access and the exit drill record. It does not automatically keep agency methods or vendor model state. Those are priced and licensed, or left behind knowingly. · **ANALYSIS**

SECTION 28

The next 90 days

One bounded decision pilot plus an exit drill, with an accountable owner and a pre-agreed rule for expanding or stopping.

IN SHORT


Days 1 to 30 reconcile the money, freeze the rules and test the stop, a refused request, a retry and the exposure cap before any live spend. Days 31 to 60 run the pilot at authority level 3 or 4, rehearse the stop in flight and test one export. Days 61 to 90 challenge the result, run the exit drill and decide against the frozen rule. A failed pilot that yields usable evidence is a success of the method.

The practical CMO action is one bounded decision pilot plus an exit drill. It has an accountable owner and a rule for expansion or stopping agreed before the pilot starts. The plan below gives each step a week, an owner, a baseline, a cost line, a stop criterion and a review. It is a proposal; the costs are categories, not estimates. · **ANALYSIS**

FIGURE 21 ILLUSTRATIVE

The next 90 days

One bounded decision pilot plus an exit drill, in four phases over about thirteen weeks. No live spend before the signed page and four control tests pass.



WEEK ANALYSIS	ACTION ANALYSIS	OWNER ANALYSIS	BASELINE IT NEEDS ANALYSIS	COST LINE ANALYSIS	STOP CRITERION ANALYSIS	REVIEW ANALYSIS
1	Choose one business decision to pilot, for example new-customer acquisition in one market or one channel's allocation; write the goal, horizon, margin floor, caps, exclusions and data rules in the section 26 format	CMO with CFO	The current goal and how it is measured	Senior time	No single decision can be named with a measurable goal	Week 2 sign-off
1 to 3	Reconcile the money: one brand's media, fees, data, production, internal labour and measurement cost; who pays whom; where any supplier acts as principal	Finance and procurement	Last four quarters of invoices against contracted services	Analyst time; audit support if needed	Money cannot be reconciled to contracted services	Week 3
2 to 3	Map decisions and permissions: every supplier with authority to spend or transform data; accounts; grants; dependencies	Product lead; data steward; security	Account and grant inventory	Internal time	No one can name who may spend	Week 3
3 to 4	Record the baseline: current quality, time to market, business metrics, internal staffing; unknowns marked, not zeroed	Proof lead	Cost per task; time to market; quality judgment; attributed and, where available, incremental results; exit time per task	Internal time	Baseline cannot be written for the chosen decision	Week 4
4	Freeze the success, stop and expansion rules on one signed page: what result expands the pilot, what stops it, what the maximum exposure is	CMO, CFO, proof lead	The baseline above	None	The three parties cannot sign one page	Week 4 gate: the pilot does not start without it

WEEK ANALYSIS	ACTION ANALYSIS	OWNER ANALYSIS	BASELINE IT NEEDS ANALYSIS	COST LINE ANALYSIS	STOP CRITERION ANALYSIS	REVIEW ANALYSIS
5 to 8	Run the decision pilot at authority level 3 or 4, never higher, with logged actions, exceptions and review effort	Delivery lead, agency or internal	Frozen rules; exposure limit	Pilot media, a sum small enough to stop in a day; tool cost; review hours	An exclusion or consent breach; exposure above the cap; a quality failure the creative lead will not accept	Exception review at a cadence set by exposure, reversibility and measurement lag; weekly is the illustrative default
5 to 8	Run a comparator where feasible: a holdout or matched comparison, pre-registered	Proof lead	Pre-registered design	Test budget from the learning fund	Design cannot be pre-registered before launch	Week 8 read
4	Before any live spend: test one stop, one refused out-of-bounds request, one forced retry and the exposure cap, on a sandbox or a zero-budget grant	Security; product lead	Grant inventory	Internal time; vendor support	The stop does not reach every connected tool and booked commitment; an out-of-bounds request is not refused; a retry duplicates a booking; the cap can be exceeded	Week 4 gate: no live authority until all four pass
6	In-flight rehearsal: repeat the stop with live spend running, and test one export	Security; product lead	The week 4 results	Internal time; vendor support	The live stop misses a tool or a booked commitment; an export cannot be read	Same day; a failed stop halts the pilot
6 to 8	Pilot AI on one bounded task with total review cost counted: approved useful work per total hour, before and after	Operations	Current hours per approved deliverable	Tool cost; review time	Quality falls or review cost exceeds the saving	Week 8
5 to 8	Keep creative and brand assessment beside delivery metrics	Creative lead	Current brand-health read	Research cost if tracked	A brand-grounds rejection stands	Weekly
5 to 8	Measure the relationship beside the controls: brief quality rated by both sides, approval delay in days, rework share of hours, alongside supplier-conflict checks	Brand lead with the agency lead	Current brief, approval and rework record	Internal time	None; a measure, not a gate	Week 8 read

WEEK ANALYSIS	ACTION ANALYSIS	OWNER ANALYSIS	BASELINE IT NEEDS ANALYSIS	COST LINE ANALYSIS	STOP CRITERION ANALYSIS	REVIEW ANALYSIS
9 to 10	Challenge the result and the counterfactual: a written challenge by the proof lead or an outside reviewer with access and freedom to disagree; both reports go to the CMO unchanged	Proof lead or outside reviewer	The pre-registered design and the frozen criteria	Reviewer fee	Criteria were changed after results	Week 10
10 to 12	Exit drill: export plus a rebuild of a small part of the workflow outside the incumbent; lost state, cost, downtime and licences documented	Product lead; procurement	Export formats named in the contract	Internal time; vendor transition support	Usable state cannot be rebuilt on the promised terms	Week 12
11 to 12	Full-cost reconciliation: fees, tools, review, retained staff, failed tests and exposure before proof, against the claimed benefit	Finance	The week 3 reconciliation	Analyst time	Benefit cannot be stated after full cost	Week 12
13	Decide against the frozen rule: expand, revise or stop; what is retained, delegated or bought next; a pre-agreed expand-or-stop rule applied in one meeting	CMO as accountable owner; board sponsor informed	The signed page from week 4	None	The rule is not applied as written	The single review meeting

No agent receives live authority before the week 4 control tests pass. Any one stop criterion in weeks 5 to 8 stops the pilot the same day. The review in week 13 is a single meeting with the frozen rule on the table. A failed pilot that yields usable evidence is a success of the method; an untested global rollout is not. The decision may favour a group, an internal core or a mix. The research supports a standard of proof, not an answer. · **ANALYSIS**

Baselines the pilot must not run without: current cost per task; current time to market; current quality judgment; current attributed and, where available, added results; current exit time per task; current brief quality, approval delay and rework share. Unknowns are marked as unknown, never forced to zero. · **ANALYSIS**

Ten research questions and update triggers

Ten questions this paper cannot settle from its register, each with the observation sought, the trigger that would update the paper's reading, and a window. Where an expectation is kept, its basis and its threshold are stated now.

IN SHORT

None is a forecast with a probability. Each runs from the 5 October 2026 cutoff. A trigger that has not fired by its window leaves the paper's reading unchanged, not confirmed. A disclosure milestone is not proof of economic success or failure.

Earlier editions gave verdicts. This one gives questions a reader can check. Absence from this paper's limited register is not a basis for predicting future absence, so most rows ask what observation would settle the question rather than predicting that none will appear. Where an expectation is kept, the row states its analytical basis and defines its threshold now, and the trigger is the logical negation of the expectation. Compound questions are split. All windows run from the 5 October 2026 cutoff. · ANALYSIS

Ten research questions and their update triggers. Thresholds are set now; none is left to be chosen later.^{[8][27][7][20][21][23][9][10][11][37][31][41][17][18][13][14]}

Column classifications: · ANALYSIS

QUESTION ANALYSIS	OBSERVATION SOUGHT, AND ANY EXPECTATION KEPT ANALYSIS	UPDATE TRIGGER, WITH THE THRESHOLD DEFINED NOW ANALYSIS	WINDOW ANALYSIS	WHERE TO LOOK ANALYSIS
1. Will buying scale keep its commercial role in negotiated inventory?	Disclosed principal or negotiated media income at the large groups, and the next buyer survey. Expectation kept: it stays material. Basis: ANA member use rose from 47% to 58% ^[8] ; negotiated inventory is not auctioned inside one platform.	The reading changes if two of the six groups disclose principal or negotiated media income falling as a share of media income for two consecutive annual reports, or if a comparable ANA wave reports principal use below the prior wave's 47%. Threshold: a fall in share over two reports, or an incidence below 47%.	To the end of 2028	Group accounting policies on principal contracts; the next ANA principal-media wave

QUESTION ANALYSIS	OBSERVATION SOUGHT, AND ANY EXPECTATION KEPT ANALYSIS	UPDATE TRIGGER, WITH THE THRESHOLD DEFINED NOW ANALYSIS	WINDOW ANALYSIS	WHERE TO LOOK ANALYSIS
2. Will any group disclose a recurring technology-fee line with retention?	A separately reported recurring technology or platform fee line, with a client retention or renewal figure, in audited accounts. Basis: the register holds only \$16 million committed and mixed at one group [27]. No expectation of absence is kept.	Any of the six groups reports such a line at or above 5% of its net revenue, a threshold this paper sets now as a choice, with retention disclosed. That would move claim (b1) in section 15 for that group from unestablished to supported for the revenue model; it would say nothing about delivery or client value.	To FY2027 reporting, in the first half of 2028	Revenue notes in annual reports; segment disclosures
3. Will a matched-scope client-profit study of an integrated group service be published by a party without a sales interest?	A study with a credible counterfactual, a full cost boundary on the nine lines of section 20 and a named scope. Basis: the only case in the register is a vendor-published before-and-after without a control [7]. No expectation of absence is kept.	One such study, whatever its result: a net gain moves claim (b3) for that scope; a null or adverse result strengthens the counterfactual. A study that fails the counterfactual or cost test does not fire the trigger.	To the end of 2027	Trade-body and academic publications; client-side disclosures
4. Will public evidence reach level 6 of the authority ladder?	A named deployment with stated spend scope, time window, approval limits, a loss budget and an incremental-profit result read by a party without a sales interest. Basis: the register's highest documented level is 5, described and untested (section 16).	One such deployment, documented to that standard, moves the level 6 row from no reviewed evidence to named live use at the third gate. A product page that states spend scope and stop rights without a result moves it only to documented implementation.	To the end of 2027	Product documentation that states spend scope and stop rights; client case studies with controls
5. Will a pinned-version, cross-vendor compatibility result be public for AAMP or AdCP?	A published run naming both implementations and versions, with the surviving record on a failed request, a retry, a revoked grant and a duplicate booking. Basis: proposals, SDK descriptions and a pinned repository, and the FAQ's own caution about transport and meaning [20] [21] [23].	One such run moves the protocol's row in section 17 to compatibility test for that pair and version only. A conformance badge without the four cases does not fire the trigger.	To the end of 2027	Conformance programme publications that name both sides; buyer-run tests

QUESTION ANALYSIS	OBSERVATION SOUGHT, AND ANY EXPECTATION KEPT ANALYSIS	UPDATE TRIGGER, WITH THE THRESHOLD DEFINED NOW ANALYSIS	WINDOW ANALYSIS	WHERE TO LOOK ANALYSIS
6. Will in-housing continue to coexist with outside agency use?	The next comparable survey waves. Expectation kept: coexistence at near-universal incidence. Basis: 92% of ANA respondents still used outside agencies; all WFA/Observatory respondents did ^[9] ^[10] .	The reading changes if a comparable wave reports fewer than 80% of respondents with an in-house agency still using outside agencies. Threshold: 80%, set now.	The next ANA and WFA waves, to the end of 2028	Survey waves with stated denominators
7. Will outcome fees stay a minority of agency compensation in most disciplines?	The next WFA remuneration wave. Expectation kept: a minority. Basis: 63% of 69 WFA respondents expect more such fees, which remain a small share today ^[11] .	The reading changes if a comparable wave reports outcome or performance fees above 50% of compensation in a majority of disciplines, the negation of a minority. Growth to 20% or 30% would be consistent with the expectation, not against it.	The next WFA remuneration wave, to the end of 2028	WFA remuneration studies
8. Will a further material impairment or disposal accompany the integration story at one of the six groups?	Goodwill notes and disposal announcements. Expectation kept: at least one. Basis: Publicis' FY2016 charge, Dentsu's FY2025 impairment and WPP's Kantar sale ^[37] ^[31] ^[41] .	The reading changes if no group books a goodwill impairment at or above 1% of its net revenue for the year, and no group disposes of an acquired capability business, across FY2026 and FY2027. Threshold: 1% of net revenue, set now.	To FY2027 reporting	Goodwill notes; disposal announcements
9. What is LiveRamp's status?	A dated completion or termination record. This is an update rule, not a forecast. Basis: agreement announced 17 May 2026; the 16 September 2026 bond release still described approvals as pending ^[17] ^[18] .	Either record ends the pending status; the paper then carries a dated update note and never infers ownership from silence.	Until a completion or termination record dated after 16 September 2026 is opened	Issuer releases and filings
10a. Does Horizon's public design stay closest to Scenario B?	Releases or contract disclosures that show where decision rights and operating state sit. Basis: the June and October 2026 releases describe a managed open ecosystem ^[13] ^[14] .	The classification moves toward A if a disclosure shows brand-held decision rights and portable state, and toward C if it shows dependence on platform model state the operator cannot export. A classification change says nothing about client value.	To the end of 2027	Horizon releases; client disclosures

QUESTION ANALYSIS	OBSERVATION SOUGHT, AND ANY EXPECTATION KEPT ANALYSIS	UPDATE TRIGGER, WITH THE THRESHOLD DEFINED NOW ANALYSIS	WINDOW ANALYSIS	WHERE TO LOOK ANALYSIS
10b. Will pilot conversion be published with a denominator?	Sustained live deployments divided by completed eligible pilots for a dated cohort, with stops and pending cases reported. Basis: the releases give participation counts, not adoption [14].	One such publication moves the Horizon case from programme activity to named live use at the third gate. It is a disclosure milestone, not proof of economic success or failure; a tested client exit would separately move the portability signpost.	To the end of 2027	Horizon releases; client disclosures

Rule for reading the table: a trigger that has not fired by its window leaves the paper's reading unchanged; it does not confirm it. The paper will record each outcome in a dated update note. · ANALYSIS

SECTION 30

Five questions for the conversation

The public crosswalk: where the paper answers each of the host's five questions, the limit of each answer, and one follow-up.

IN SHORT

The host's position to test is that agencies may earn more value through strategy, creative judgment, analytics and coordination, if the work creates value for the brand. The paper supports the standard of proof, not the conclusion. A guest agreeing with the thesis is not separate proof of it.

The author plans a short conversation with a guest about these questions. No transcript exists at the cutoff. Nothing here infers the guest's identity, role, views or answers. The host's position to test is that agencies may earn more value through strategy, creative judgment, analytics and coordination, if the work creates value for the brand. The table maps each question to the sections that answer it, the limit of that answer, and one follow-up worth asking. · FACT

Five questions, where the paper answers, the limit, and one follow-up.[39][7][9][10][8][11][21][49][50]

Column classifications: · ANALYSIS · ATTRIBUTED

QUESTION ANALYSIS	WHERE THE PAPER ANSWERS ANALYSIS	LIMIT ATTRIBUTED	ONE FOLLOW-UP ANALYSIS
As AI takes on more work, what should a brand still pay an agency to do? When does a group add value beyond a few specialist partners?	Sections 24 (make, buy or partner), 15 (claims b, c, k) and 14 (value created and captured)	No matched client study; no group premium measured. Coordination predates AI [39]; integration milestones exist but client value is not established [7].	Name one join a group made for a client that two specialists could not have made, and what it cost in full.
Does in-housing mean doing more tasks, or retaining control of goals, data, decisions and learning? What can the brand take on exit?	Sections 27 (retained team), 25 (four kinds of learning; exit terms), 15 (claims d, j) and 12	Incidence is not share of work [9] [10]; brand control of learning is a scenario condition, not a verified right; no exit drill result is in the public record.	What did the last client that left take with it, and what could it not use?
How do fees, principal margins and outcome targets shape recommendations? How should an agency earn when the client spends less but earns more?	Sections 18 (fee routes; task model), 25 (commercial design) and 15 (claims h, i)	No fee file or spread reviewed; principal use rose from 47% to 58% in ANA's samples [8]; 63% of WFA respondents expect more performance fees, still a small share [11]; the task model is illustrative.	What would you charge for a recommendation to spend less?
Who resolves conflicts across agents, grants authority and stops the system? Who carries spending risk while reliable results are pending?	Sections 26 (the retailer), 16 (authority ladder), 17 (controls) and 22 (liability in path B)	Approval before booking is described in an SDK release [21]; no authority log was inspected; no live stop test was run; the \$70,000 illustration is arithmetic, not a loss.	Has a stop ever been pulled mid-run, and what did it miss?
What skills and decision rights does the CMO retain, and what should change in the next 90 days?	Sections 27 (retained team), 28 (90 days) and 19 (two horizons)	Two disclosed pay plans [49][50] are cases, not a norm; the 90-day plan is a proposal.	Which one decision would you take back from your agency this quarter?

The one question to ask

If an agency builds your marketing system, have you gained control, or moved your dependence to the system that makes the decisions? The follow-up: what can the client actually take when it leaves? A strong answer names rights, usable learning and a tested transition, not just ownership of source data. · **INFERENCE**

The guest's eventual answer enters the paper as attributed testimony. It is kept distinct from the host's view, with any disagreement preserved. It does not verify the thesis. If a transcript arrives, a separate timestamped comparison of host, guest and paper will classify each statement as experience, factual claim, judgment or scenario. Facts will be checked before they enter the paper. Only material cleared for publication will be used. No quotes appear in this edition. · **FACT**

For the conversation itself: one concrete case, the retailer in section 26. One serious challenge: the record shows groups bought the parts and reported margins but not matched client profit, so why should a CMO believe the premium exists? One first step: the bounded decision pilot plus exit drill in section 28. Protocol names, accounting terms and survey denominators belong in the notes, not on the microphone. · **ANALYSIS**

Method, coverage and release boundary

What was checked, how, with what denominators, and what this edition does not claim.

IN SHORT

A targeted audit of 72 selected claims and a separate 20-claim Horizon supplement, 49 source passages, 32 inherited recomputations within 104 numerical and comparability checks and 362 original paragraphs. No source was re-opened in this edition or in its 6 October 2026 update; an external review of that date opened 13 and its checks are attributed, not adopted. Every record carried forward unread is listed here once, with its effect on the conclusion. Three passes by one assistant are three processes, not outside review. The browser, PDF and printed-type checks ran on the final build of 7 October 2026 and are in the published test record. Readability was measured on 7 October: 69.4, against a target of 85.

Coverage, with denominators

The prior audit inspected 49 source passages. It audited 72 selected atomic claims from the original paper, plus a separate 20-claim Horizon supplement. The two sets are kept apart, and the supplement is not a full audit of Horizon or its partners. The original paper held 362 chapter and executive paragraphs, a 232-claim register and 329 financial records. Not all of those were audited; the unaudited remainder stays in the private archive. The 32 inherited calculations were recomputed; 31 reproduce and one cannot be recomputed from the stored records. The calculation ledger of this edition holds 104 numerical and comparability checks: the 32 inherited recomputations, 48 numbers checked from the revised edition against their stored inputs, 15 task-model edge cases, 7 peer-growth observations and 2 precision rules. Only the 32 are recomputations in the strict sense; 35 rows carry no formula, and the two precision rules are rules, not computations. Of the 104, 97 reconcile or reconcile with a label fix and 7 cannot be recomputed from the stored records. The public claim ledger below carries 210 rows. Each row carries a status. Of the 210 rows, 160 are active claims. Their displayed assertion is the wording this edition prints, and the verdict grades that wording. Ten are withdrawn; each keeps only the wording that was withdrawn, which is not printed. Forty are editorial rows: statements about this paper's own editions, audit and method. Where a printed wording changed from the audited original, the row carries both wordings, and the verdict grades the current one.

The source register carries every record this edition cites, with its stored passage and the search-level status recorded on 5 October 2026. · **FACT**

No source was re-opened in this edition. Outbound page fetches to every source host were blocked in the session that prepared it. Each inherited record was rechecked by web search for its URL, title, date, corrections and supersession. A search result is a discovery aid, never support. Two repositories were cloned by git on 5 October 2026 and pinned by commit and file hash: AdCP at commit 6adfebe2, whose release notes date version 3.2.1 to 30 September 2026 ^{[22][25]}, and OpenProposal at commit 66ce40a8, with tag v3.0-draft-1 present ^{[24][20]}. Every other factual statement rests on the stored passage from the 5 October 2026 audit, carried forward unread and labelled as such in the register. Candidate primary sources found by search but not opened are listed in a separate candidate-sources file, not in the source register. They are not used as support anywhere in this edition. Of the 51 register records, 49 passages were inspected in the 5 October 2026 audit and two repositories were cloned by git; none was re-opened in this edition or in its 6 October 2026 update; the rest of the check was search-level. An external critical review dated 6 October 2026 reports that it opened and inspected passages in 13 of these sources and reproduced all 47 arithmetic rows. This update names that review where it relies on one of its checks: the LiveRamp offer price (sections 6 and 13), the WPP client-loss attribution (section 7) and the WFA ranking of briefing, trust and feedback (section 19). Those are attributed external checks. They are not this paper's own verification. The re-opening of the 49 inherited sources remains open. ^{[22][25][24][20]} · **FACT**

Records carried forward unread, listed once

The group chapters no longer repeat their caveats. Every material record from the original edition that this edition carries forward unread is listed here once. Each row names the check that would settle it and its effect on the chapter's conclusion. None is used as fact. One access limit sits behind this list: outbound fetches were denied in the sessions that prepared this edition. That limit shapes the record. No result in this register is a narrower statement than no useful evidence in the market. · **ANALYSIS**

Material records carried forward unread, the check that would settle each, and its effect on the conclusion.

Column classifications: · **ANALYSIS**

CHAPTER ANALYSIS	RECORD CARRIED FORWARD UNREAD ANALYSIS	CHECK THAT WOULD SETTLE IT ANALYSIS	EFFECT ON THE CONCLUSION ANALYSIS
Section 4, the timeline	About sixty events of the original timeline not re-checked; seven trading-desk lineages (Xaxis, Accuen, Cadreon and Matterkind, Amnet and Amplifi, Affiperf, Varick); two abandoned deals (the Publicis and Omnicom merger plan ended in 2014; the Dentsu international sale ended in February 2026)	Dated issuer releases for each event	None on the thesis; the table stays a floor, not a series. The abandoned deals would strengthen the no-straight-line reading.
Section 6, Publicis	The Digitas closing (January 2007 in the original edition); the Epsilon integration plan; the LiveRamp shareholder vote of 17 August 2026 and the related filing	The completion release; the integration announcement; LiveRamp's filings	None: the chain runs on agreement and completion records already held, and the vote would not change the pending status.
Section 7, WPP	InfoSum terms; 24/7 Real Media, TNS, AKQA, Satalia and Goat; the three-layer stack, platform usage counts, production site, self-serve route and agent hub; the 2020 AKQA and Grey combination separated in 2025 with £641 million of goodwill impairments	Deal releases; the 2025 annual report's goodwill note	The impairment would enlarge the countercase and is the largest reversal in the original record; the usage records would add company-reported use, not client results. Neither changes the chapter's conclusion.
Section 8, Omnicom and IPG	Annalect, Credera, IPG's 2018 Acxiom purchase and the deal prices; Omni's 2018 launch and the January 2026 release naming Acxiom and Flywheel; IPG's sale of Huge and R/GA; the 2024 impairment on IPG's digital specialists; the 2014 abandoned Publicis merger; the US regulatory conditions	Issuer releases; IPG and Omnicom filings; the regulatory order	The countercases would enlarge the chapter's countercase. The Omni releases carry the chapter's central shift from partner-neutral to owned, which stays analysis of the original edition's account until they are read.
Section 9, Dentsu	The 2016 Merkle rationale; the 2023 one-management-team structure; the 2026 audience product; the H1 2026 regional split; the February 2026 decision not to sell the international business; the March 2026 chief executive change	Dentsu releases and the H1 2026 report	The regional split and the 2026 decisions would sharpen the cost-base question; none would assign the impairment to one deal.
Section 10, Havas	Family-linked holdings as of September 2026; FullSIX, Havas CX, Uncommon, Channel Bakers and the 2026 deals; released tools; the H1 2026 figures; the prospectus health-client loss and the €400 million 2024 to 2027 programme	Shareholder disclosures; deal releases; the H1 2026 report; the listing prospectus	The control position is unknown at the cutoff; the prospectus items would count against breadth. The mixed-model question stays open either way.

CHAPTER ANALYSIS	RECORD CARRIED FORWARD UNREAD ANALYSIS	CHECK THAT WOULD SETTLE IT ANALYSIS	EFFECT ON THE CONCLUSION ANALYSIS
Section 11, Stagwell and S4 Capital	72andSunny, Anomaly, the Harris assets, Brand New Galaxy and ADK GLOBAL; the three 2026 product releases (a January workflow layer, a June media operating system, a July preferred-supply marketplace); one named-client cost claim; S4 Capital's prior-year margin and debt values	Stagwell releases; S4 Capital's 2024 results	The releases would document the design the chapter analyses; the client cost claim would be a company-reported result, not a matched test. The conclusion does not depend on them.
Section 13, the numbers	Publicis H1 2026 organic growth; Omnicom Q2 2026 core organic growth; Dentsu H1 2026 organic growth; Havas H1 2026 organic growth; S4 Capital H1 2026 like-for-like change	The five primary results documents	None: the rows would return without a shared axis, and no growth rate shows a client result.
Sections 12 and 17, platforms	Primary pages for platform agent products described in search results as beta or pilot	The product pages, dated	None: a stage would be recorded; no gate above documented implementation would change.
Sections 1 and 4, the categories	Purchases of creators, production and AI tooling (Influential, Captiv8, Tag, WPP Production and CoreAI in the original edition)	Deal releases	None: the layered reading holds on the audited categories; the three would add layers, not change the pattern.

What this edition did and did not do

Three passes were run by one assistant: evidence and numbers; adversarial reasoning and commercial design; editorial, citation and build consistency. They are three distinct processes, not three sources of evidence. Their agreement is not outside fact-checking. The release checks ran on the final build of 7 October 2026: a real-browser check of 305 assertions at 375, 1024, 1366, 1440 and 1920 pixel widths, which runs every control of every interactive figure in the page and in the enlarged view; an automated type and overflow check of all 170 PDF pages; and a visual review of the figure pages. The results are in the published test record (figure-test-record.md). A screen reader, a physical phone, real browser zoom and other browsers were not used. Readability was measured on this text with a stated tool on 7 October 2026: the narrative scores 69.4 on the Flesch reading-ease scale (69.5 after the 6 October update, 71.1 before it), against a target of 85, which this edition does not meet; the plain edition scores 87.3, above its own gate of 85. Link integrity was not probed in this review, because outbound fetches were denied, and no source was re-opened. This edition ran no live media spending, no customer-data experiment, no compatibility test and no stop or revocation test. Every control in section 17 and every test in section 28 is proposed for a buyer, not run by this review. No interview was conducted and no private fee file was read. The task model and the retailer case are illustrative. The commercial design, retained team and 90-day plan are proposals. • **FACT**

Update-note policy. Facts that change after the 5 October 2026 cutoff are recorded as dated update notes at the end of this paper. Examples: a LiveRamp completion or termination record, a protocol version beyond the pinned commit, or a product moving from pilot to general availability. The body text is not silently rewritten. Corrections to the body are logged in the corrections ledger with the date, the location and the change. · **FACT**

Instruments

Four instruments follow. The version map shows what moved between the three editions and why. The claim ledger lists every audited atomic assertion with its wording, classification, sources, verdict and correction. The source register lists every cited source with its stored passage, dates and search-level status. The corrections ledger lists every correction made to this paper, with the date and the change. · **FACT**

The 48-row version map is in the web edition and in the download data/version-comparison.csv; it is not printed here.

The 210-row claim ledger is in the web edition and in the download data/atomic-claim-ledger-v3.csv; it is not printed here.

Numbered in order of first citation. Every entry carries the passage the 5 October 2026 audit recorded for it; this edition re-opened none of them, and pinned two repositories by git clone. The locator and the limits of each stored passage are recorded with it.	
[1]	Digitas acquisition agreement. Publicis Groupe, 2006-12-20 (filing). PUB-S001 · L4–19,47–78 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2006 agreement
[2]	Sapient acquisition completion. Publicis Groupe, 2015-02-06 (filing). PUB-S003 · L17; completion release · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2015-02-06
[3]	Epsilon completion and integration plan. Epsilon / Publicis Groupe, 2019-07-02 (official_release). PUB-S006 · L10,16–43 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · July 2019
[4]	Flywheel closing. Omnicom, 2024-01-02 (official_release). OMC-S020 · L469–479 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2024-01-02
[5]	Omnicom completes IPG acquisition. Omnicom, 2025-11-26 (official_release). OMC-S024 · L5,8 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2025-11-26
[6]	Stagwell and MDC complete combination. Stagwell, 2021-08-02 (official_release). ALT-S006 · L76–90 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2021-08-02

- [7] Acxiom Digital Recognition/Guitar Center case. Omnicom, 2026-07-14 (official_release).
OMC-S033 · L15–23 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2026 case publication
-
- [8] Governance around principal media failing to keep pace. Association of National Advertisers, 2026-03-19 (research).
ECO-S002 · L128–146; Methodology · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2026 survey; prior wave comparison
-
- [9] In-house agencies here to stay. Association of National Advertisers, 2023-05-02 (research).
ECO-S003 · L129–136 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2023 member survey
-
- [10] In-housing set for rapid and continued growth at multinationals. The Observatory International / WFA study, 2023-12-20 (research).
ECO-S004 · L65–86 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2023 study
-
- [11] Global brands shifting to paying for outputs not time. World Federation of Advertisers / Agency Mania Solutions study, 2026-09-29 (research).
ECO-S006 · L33–44,55–75 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2026 publication
-
- [12] Generative AI at Work. Oxford University Press / Quarterly Journal of Economics; Brynjolfsson, Li and Raymond, 2025-02-04 (research).
ECO-S022 · Abstract L50–87 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Published 2025-02-04
-
- [13] HorizonOS Blu agentic integration announcement. Horizon Media via PR Newswire, 2026-06-18 (official_release).
HOR-S002 · L476–494; Proof Points and Client Impact, L490–494 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2026-06-18 announcement; results period not supplied
-
- [14] HorizonOS Labs third cohort. Horizon Media via PR Newswire, 2026-10-01 (official_release).
HOR-S003 · L478–486; partner descriptions L491–505; pilot next steps L506; corporate note L510–512 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2026-10-01 announcement; results period not supplied
-
- [15] About Performance Max campaigns. Google, undated living documentation (other).
PLT-S001 · L19–49 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Observed 2026-10-05
-
- [16] Performance+ hits general availability and Brand+ enters open beta. Amazon Ads, 2025-03-21 (official_release).
PLT-S003 · Title L3, date L7, body L35–50 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Dated March 21, 2025; body observed cutoff
-
- [17] LiveRamp agreement and proposed future organization. Publicis Groupe, 2026-05-17 (official_release).
PUB-S009 · L23,30,45–48,95–108 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Agreement May 2026
-
- [18] Publicis bond financing and pending approvals. Publicis Groupe, 2026-09-16 (official_release).
PUB-S012 · L23,26–31 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2026-09-16
-
- [19] Retail media integrated products launched. Publicis Groupe, 2023-06-16 (official_release).
PUB-S023 · L29–36 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2023 product launch
-

- [20] IAB Tech Lab introduces AAMP 3.0 with OpenProposal. IAB Tech Lab, 2026-09-22 (regulator_or_standard).
ROOT-S002 · L15–21 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2026-09-22
-
- [21] AAMP 2.0 buyer and seller agent SDKs. IAB Tech Lab, 2026 (regulator_or_standard).
ROOT-S003 · L15–21,28–38,62–71 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Observed 2026-10-05
-
- [22] Advertising Context Protocol repository. Ad Context Protocol project; GitHub repository, Not shown (regulator_or_standard).
ROOT-S005 · README and repository structure, L230–277 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Observed 2026-10-05
-
- [23] AdCP FAQ. Ad Context Protocol project; GitHub documentation, Not shown (regulator_or_standard).
ROOT-S006 · FAQ L266–308 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Observed 2026-10-05
-
- [24] IABTechLab/OpenProposal repository — AAMP 3.0 OpenProposal draft (pinned). GitHub (IABTechLab), 2026-09-22 (draft 1 published for comment) (regulator_or_standard).
ROOT-S301 · see document_or_version_id · accessed 2026-10-05 (git clone) · Draft specification
-
- [25] AdCP repository release records at the pinned commit (CHANGELOG.md, docs/reference/release-notes.mdx, docs/reference/roadmap.mdx, docs/reference/versioning.mdx). GitHub (adcontextprotocol), 2026-10-05 (pinned snapshot) (regulator_or_standard).
ROOT-S302 · see document_or_version_id · accessed 2026-10-05 (git clone) · AdCP 3.x line
-
- [26] WPP acquires InfoSum in major investment in its AI-driven data offer. WPP, 2025-04-03 (official_release).
WPP-S004 · L11–23 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2025-04-03
-
- [27] Q2 2026 earnings presentation. Stagwell, 2026-07-30 (filing).
ALT-S018 · Slides9,13–17; L84,115–143 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · H1 2026
-
- [28] Stagwell 2025 Form 10-K. Stagwell, 2026-03-13 (filing).
ALT-S023 · MD&A p30; L988–1018 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · FY2025; method from Q3 2025
-
- [29] History: Dentsu establishes Dentsu Aegis Network. Dentsu Group, Not shown (company_report).
DH-S003 · L140–153 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2012–2013
-
- [30] Dentsu confirms 100% ownership of Merkle. Dentsu Group, 2020-04-15 (official_release).
DH-S008 · L9–12 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2020-04-15
-
- [31] Financial Report FY2024–2025. Dentsu Group, 2026 (company_report).
DH-S016 · PDF pp82,90–92; note 14 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · FY2025
-
- [32] Successful Euronext Amsterdam listing. Havas, 2024-12-16 (official_release).
DH-S029 · L9–15 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2024-12-16
-
- [33] Horizon Global formation. Havas, 2025-09-29 (official_release).

DH-S038 · L14–19 and scope description · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2025-09-29

[34] HorizonOS launch. Horizon Media via PR Newswire, 2025-12-18 (official_release).

HOR-S001 · L485–503; About Horizon Media Holdings, L507–509 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2025-12-18 announcement; results period not supplied

[35] Omnicom redistributes specialist staff into agency brands. AdExchanger article reprinted by Omnicom Media, Not re-established; inherited record gives 2019 (reporting).

OMC-S018 · L4–15 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Historical article; date not re-established

[36] Hearts United launch. Omnicom, 2026-08-28 (official_release).

OMC-S034 · L7,10,17 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2026-08-28

[37] FY2016 results and impairment. Publicis Groupe, 2017-02-09 (company_report).

PUB-S004 · L94–101 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · FY2016

[38] Publicis H12026 financial report. Publicis Groupe, July 2026; exact publication day not established (company_report).

PUB-S015 · Printed pp4,11,25; PDF pp6,13,27 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · H1 2026; Lotame Apr2025

[39] Power of One transformation plan. Publicis Groupe, 2015-12-02 (official_release).

PUB-S035 · L66–87 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2015 plan; 2016 implementation date

[40] WPP Media launches as fully integrated, AI-powered media company. WPP, 2025-05-28 (official_release).

WPP-S001 · L11–23 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2025-05-29

[41] Completion of Kantar transaction. WPP, 2019-12-05 (official_release).

WPP-S007 · L11–18 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2019-12-05

[42] GroupM Launches Xaxis: audience buying solution. GroupM announcement distributed by Business Wire, 2011-06-28 (official_release).

WPP-S012 · L15–24 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2011-06-28

[43] 2025 URD income statement. Publicis Groupe, Not shown on inspected page; inherited record gives 2026-04-30 (filing).

PUB-S039 · URD printed p260; income statement L18–31 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · FY2025

[44] S4 Capital audited FY2025 results. S4 Capital, 2026-03 (company_report).

ALT-S028 · PDF pp1–4, financial summary and notes · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · FY2025

[45] Omnicom Q2 2026 earnings release. Omnicom, 2026-07-28 (filing).

OMC-S036 · Q2 Core Operations table; L25–60 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · Q2 2026

[46] H12026 guidance page. Publicis Groupe, 2026-07-16 (company_report).

PUB-S041 · L23–36,51–52 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · H1/Q2 2026

- [47] 2026 Interim Results. WPP, 2026-08-06 (company_report).
WPP-S003 · Key figures L13–56; operating review L67–80 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · H1 2026
- [48] Full-year2025 results. Havas, 2026-02-17 (company_report).
DH-S041 · Key results L13–15,63–65 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · FY2025
- [49] Coca-Cola 2026 proxy. The Coca-Cola Company; proxy filing, 2026-03-16 (filing).
ECO-S016 · PDF pp52,56–61; printed pp48,52–57 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2025 pay; 2026 proxy
- [50] Kohl's 2026 proxy. Kohl's Corporation; proxy filing, 2026-04-10 (filing).
ECO-S017 · PDF p91; printed p84; also pp93 onward · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · FY2025 pay
- [51] A comparison of approaches to advertising measurement. INFORMS / Marketing Science; Gordon, Zettelmeyer, Bhargava and Chapsky, 2019-04-04 (research).
ECO-S020 · Abstract L213–219 · accessed 2026-10-05 (stored audit open); 2026-10-05 (this session, search-level only) · 2019 paper

Download the register: [CSV](#) · [JSON](#).

The 255-row corrections ledger is in the web edition and in the download data/corrections-ledger-v3.csv; it is not printed here.

METHOD, DISCLOSURE AND DOWNLOADS

How this paper was made

Public-source research with every claim, source, numerical check and correction published beside the text.

IN SHORT

Research cutoff 5 October 2026. 210 atomic claims (160 active, 10 withdrawn, 40 editorial), 51 sources (49 passages inspected in the prior audit, 2 repositories pinned by commit, 49 rechecked by search only, none re-opened in this edition), 104 numerical and comparability checks and 255 corrections are downloadable below.

Public-source study. Research cutoff 5 October 2026; history from January 2006. The prior audit inspected 49 source passages and audited 72 selected claims from the original paper plus a separate 20-claim Horizon supplement; 32 inherited

calculations were recomputed and 362 original paragraphs inventoried. The 232-claim register and 329 financial records of the original paper were not all audited and stay in the private archive. No source was re-opened in this edition or in its 6 October 2026 update; checks were search-level, and two protocol repositories were pinned by commit on 5 October 2026. Three passes by one assistant are three processes, not outside review. Browser, PDF and printed-type checks ran on the final build of 7 October 2026 and are in the published test record; readability was measured on this text on 7 October 2026 (narrative 69.4, against a target of 85, not met; 69.5 on 6 October and 71.1 before that); link integrity was not probed, because outbound fetches were denied. No live media spending, customer-data experiment, compatibility test or stop test was run. The commercial design, retained team and 90-day plan are proposals; the task model and the retailer case are illustrative. An external critical review dated 6 October 2026 reports that it opened 13 of the cited sources and reproduced the 47 arithmetic rows; where this edition relies on one of its checks the text names the review, and those checks are attributed, not this paper’s own verification. The calculation ledger holds 104 numerical and comparability checks, of which 32 are recomputations. Facts that change after the cutoff are recorded as dated update notes, not silently rewritten.

INTERESTS AND DISCLOSURE

The author has worked at Samba TV since May 2025. Samba TV is a TV measurement and ad data company. Earlier he worked at Hearts & Science, which Omnicom combined into Hearts United in August 2026, and at Lotame, which Publicis bought in 2025; both are named in this paper, and each is treated on the public record alone. He co-leads the Signals & Measurement working group of AgenticAdvertising.org, which publishes AdCP, a protocol discussed in section 17; Samba TV is listed among AdCP's members. No Fluff Advisory, the author's own practice, is an AgenticAdvertising.org member. Membership, working-group roles and announcements earn no protocol stage in this paper, and the author's own work on the protocol is not cited. No company named in this paper funded, reviewed or saw it before publication.

Downloads

Print edition (PDF, A4) holding-companies-2026.pdf	Reading edition (Markdown) holding-companies-2026.md	Plain-language edition plain/index.html
atomic-claim-ledger-v3 (CSV) data/atomic-claim-ledger-v3.csv	atomic-claim-ledger-v3 (JSON) data/atomic-claim-ledger-v3.json	source-register-v3 (CSV) data/source-register-v3.csv

source-register-v3 (JSON) data/source-register-v3.json	calculations-v3 (CSV) data/calculations-v3.csv	calculations-v3 (JSON) data/calculations-v3.json
corrections-ledger-v3 (CSV) data/corrections-ledger-v3.csv	corrections-ledger-v3 (JSON) data/corrections-ledger-v3.json	version-comparison (CSV) data/version-comparison.csv
version-comparison (JSON) data/version-comparison.json	Statistics the page quotes (JSON) data/stats.json	Figure test record: what was tested, how, and what was not (Markdown) data/figure-test-record.md
Dated events, 2006– 2026 (editable SVG) figures/timeline.svg	The present-day loop and the three paths, in one layout (editable SVG) figures/baseline.svg	From PAM to the governed loop (editable SVG) figures/governed-loop.svg
A. Brand-controlled, open orchestration (editable SVG) figures/scenario-a.svg	B. Agency-led, managed orchestration (editable SVG) figures/scenario-b.svg	C. Platform-dominated, fragmented autonomy (editable SVG) figures/scenario-c.svg
Two linked loops: client and capability (editable SVG) figures/horizon.svg	Who keeps a task saving? (editable SVG) figures/economics.svg	Five control questions (editable SVG) figures/control-questions.svg
The task profile (editable SVG) figures/task-profile.svg	The next 90 days (editable SVG) figures/ninety-days.svg	Value created and value captured (editable SVG) figures/value-capture.svg
Seven fee routes (editable SVG) figures/fee-routes.svg	Openness and client control (editable SVG) figures/openness-control.svg	How to read a deal (editable SVG) figures/acquisition-chain.svg
Who owns what, since when (editable SVG) figures/entity-capability.svg	From a claim to evidence (editable SVG) figures/evidence-explorer.svg	Spending to proof: two clocks (editable SVG) figures/spending-to-proof.svg
One retailer, three agents (editable SVG) figures/retailer-agents.svg	What leaves with the brand (editable SVG) figures/learning-exit.svg	How judgment creates value (editable SVG) figures/value-judgment.svg

UPDATE HISTORY

6 October 2026 Update of 6 October 2026, after an external critical review dated the same day. This update changed: the edition label, which no longer says 'Final'; the thesis in section 1 and summary point 5; the LiveRamp offer price and the WPP client-loss attribution, now printed as statements the external review confirmed in the cited sources and this edition did not re-open; proposition (b) in section 15, split into three graded claims; section 29, now research questions and update triggers; one brand-side cost boundary across sections 20 to 24; the retailer's budget tranches, discount arithmetic, exclusion dates and aggregation rule in section 26; the baseline label and the pre-live control tests in sections 20 and 28; bounded wording in place of rankings and absolutes; a value-production loop in section 19; a task profile beside the authority ladder in section 16; the partnership definitions in section 3, the Digitas status in section 6 and the top evidence gate in section 2; and the claim ledger's wording versions and status field, with 'numerical and comparability checks' in place of 'recomputations'. Every record this edition carries forward unread is now listed once in section 31 with its effect on the conclusion. The re-opening of the 49 inherited sources remains open: no source was opened in this update either, and the external review's source checks are attributed to it, not adopted as this paper's own verification. The calculator reset (review finding F14) and the print and figure typography repairs (F15) were made at build and figure level later on 6 October 2026, after this manuscript edit: the reset now restores the five defaults (100, 60, 20, 25, 5) and recomputes the chart, table and sentence together; the figure text equivalents print once in an appendix and each scenario prints once; the figures were then rebuilt with larger labels, as the next note records. These are build and figure changes, not manuscript changes; the corrections ledger logs them (ER-F14, ER-F15).

6 October 2026 Figure update, later on 6 October 2026, after the review update above. This is a rendering and interaction pass, not a new independent audit. No source was re-opened. No fact after the 5 October 2026 cutoff was added. No research fact, source or date changed. Every figure now sits in one shared shell: a number, a status badge, a one-line takeaway, controls, an Enlarge button that opens a viewer (fit width, fit all, 100%, zoom, pan, Escape to close), a caption, a text equivalent, a sources-and-assumptions note and an editable SVG download. On phones each figure shows a compact layout in place of a drawing that scrolls sideways. The fourteen existing figures were redrawn from data files that carry claim and source ids. The five control questions (section 1) are five cards, not a cycle. The timeline (section 4) runs on one linear axis with filters, a dated event list and a lens view. How to read a deal (section 5) is a numbered checklist with twelve worked cases. Value capture (section 14) and the seven fee routes (section 18) are one commercial-model explorer with two entry views. The task profile (section 16) shows stated conditions for two tasks at a time. The task-saving figure and the calculator (section 18) share one calculation and four presets; the equal split at a fee of 95 is new and illustrative. The baseline and paths A, B and C

(sections 20 to 23) are one comparison component with compact chapter views and a six-question table. The Horizon loops (section 22) open on the client loop, with a step-through. Openness and control (section 24) is a described two-by-two with no vendor count. The 90-day plan (section 28) is four phase cards with visible gates. Six figures were added on 6 October 2026: from a claim to evidence (section 2, with a protocol view linked from section 17); who owns what, since when (section 3, with chapter views in sections 6 to 11); spending to proof (section 18); how judgment creates value (section 19); what leaves with the brand (section 25); one retailer, three agents (section 26). The paper now has 20 figures and 20 SVG downloads. Nine points wait for the author and are marked open in the corrections ledger: the \$90,000 cap default and the commitments input in the spending figure; the returns wording and the weekly-budget dollar example in the retailer walkthrough; the row labels of the section 24 table; the two proposed measures in the judgment figure; whether to print the AdCP 3.2.2 patch date; the capability columns and the S4 Capital row in the ownership figure; the proposed roles in the control questions; the wording for records the deal figure leaves blank; and a compact protocol table for section 17. One assistant checked the figures in three passes, in a browser at seven widths, in light and dark themes, and in the PDF. A screen reader and a physical phone were not used.

6 October 2026 Figure repair, later on 6 October 2026, after the author's live validation of the figure update (a browser and representation review of the published page by one assistant, received the same day). The validation found the figures clearer but not ready for sign-off, and listed ten findings. All ten are addressed here, and the open points of the note above are settled. The enlarged view now moves the whole figure into the dialog, so every switch inside it works and stays in step with its list and detail panel; the zoom readout is measured from the drawing after every change, so it matches what the reader sees. The spending figure (section 18) now holds spend to the cap: each day spends at most what the cap leaves. A slow stop is a labelled assumption, off by default, and any spend past the cap is drawn as its own line, not hidden. The read schedule is stated: after the first read, each day's spend is read a set number of days later. The \$90,000 cap stays, marked as borrowed from section 26 and not a recommended limit. The retailer figure (section 26) now keeps three steps apart: a proposal can be eligible, then wait for the creative lead, then run only if a tranche has money left. With tranche 1 spent before its read, nothing runs. The returns assumption is stated in the text, the table and the figure, with an alternative offered as a figure option only; the weekly budget behind the 10% shift rule is defined. The ownership figure (sections 3 to 12) now keeps apart what an entity is, how it relates to the group, the deal's status, and the dates. Digitas reads as owned on the record but not checked again here; Lotame and LiveRamp are businesses, LiveRamp still pending; the \$16 million Stagwell line is a reported figure, not a product; Horizon is the independent case with its own private parent; SharkNinja is a named client. Each group has its own download, and the all-groups file is a labelled reading sheet. Exports wrap their text for the fallback font. The deal figure names the worked case it shows, and two labels replace one: not assessed in this review,

and not shown in the reviewed record. The baseline and path figures (sections 20 to 23) no longer print a yes-or-no “brand decides” line; each stage names who runs it, whose sign-off it needs and what is delegated. The section 24 table uses the figure’s definition: open means providers can be connected and replaced. The AdCP row keeps 3.2.1 as the version analysed, with the v3.2.2 patch tag of 4 October 2026 recorded beside it, not substituted. Section 17 carries a short protocol table from the same data. The judgment figure keeps its two proposals as questions, not measures. The control questions say a title grants no legal authority. No research fact, source or date changed, and nothing was fetched. Still open: on the two tallest views, fitting the whole drawing to the screen makes labels about 14 pixels (fit to width keeps them at 18 or more); the Horizon single-loop views are scaled copies of the overview; the ownership print takes five half-empty sheets; and one author decision, whether execution pauses when the week 4 read fails while tranche 1 still has money. Verification record: one assistant ran four repair passes and one merge, not independent reviews. On the merged build, the browser check ran every control of every interactive figure both in the page and in the enlarged view (272 transitions, each in both places) and checked the readout against the drawing’s real scale. The author’s headline cases were reproduced and pass. A screen reader, a physical phone, real browser zoom and other browsers were not used. The test record is published with the data downloads as figure-test-record.md.

7 October 2026 Update of 7 October 2026, after the author’s live check of the figure repair. The check confirmed the repair on the live page and left two items, both now done. First, the author decided what happens when an evidence gate fails. Section 26 now says: at a failed interim gate, pause new commitments and queued actions that can still be stopped; hold the next tranche and return the plan to the CMO and budget owner; unspent budget does not carry execution authority forward; resume only under a logged human decision consistent with the mandate; keep non-cancellable commitments visible. A delayed or unclear read is not a failure. Any continuation must be set in advance, with a time limit and a spend cap. In this example, after a delayed week 4 read, work may go on for up to 7 days from what is left of tranche 1, capped at \$22,500; there is no continuation after week 8. Actions already taken stay logged. The retailer figure follows the same rule. Second, the ownership figure was too tall to read whole. Each row now shows the name, type, relationship and status; dates, checks and source passages sit in the panel when a row is selected. The capability mapping is a separate view. The viewer’s buttons now say Fit width (reading) and Fit all (overview), and on screens 900 to 1,199 pixels wide it opens with its side panel hidden. A correction: the note above said that fitting the tallest views to the screen gave labels of about 14 pixels. That was not true of the Horizon ownership view, which gave about 10. After this update, at 1366 by 900 pixels every ownership view keeps labels at about 18 pixels or more, whether fitted to width or whole. No research fact or source changed. One assistant made and checked these changes. A screen reader, a physical phone, real browser zoom and other browsers were not used. The measured sizes are in the test record. Later the

same day, the figures were set to one reading width: on any desktop a drawing is at most 880 pixels wide, so its labels read at about 18 pixels beside 17-pixel text, and the Horizon single-loop view is no longer enlarged as far.

7 October 2026 Editorial cleanup of 7 October 2026, after the author's final review. The page now opens with a short current-status note; these dated notes moved here, after the method section. The edition label now reads 7 October 2026. Section 31 and claim V3-029 now report the release checks that ran on the final build, with their dates and results, instead of describing them as pending, and they keep the list of what was not tested. Readability was measured again on the final text: 69.4 for the narrative, against a target of 85 that this edition does not meet, and 87.3 for the plain edition. The hardest passages are the enumerations of the nine cost lines and the loop descriptions in parts II and III; they were left as written, because shortening them would drop terms the cost boundary needs. Section 26 now states the delayed-read rule in one sentence, as the author decided: a delayed read is not a failed result, but it does not extend spending authority. There is no continuation after week 8; new execution stays paused until the read arrives or the CMO and budget owner approve a logged exception with its own spending cap, expiry and rationale, and tranche 3 stays held. The retailer figure's table says the same; the spending-to-proof figure is unchanged. No finding or source changed (FG-61 to FG-63).

7 October 2026 Figure addition of 7 October 2026, at the author's request. Section 20 gains one paragraph and one figure, "From PAM to the governed loop": the present-day loop, here called PAM for plan, activate and measure, beside a possible future of three loops on three clocks around the five rights. It restates the section 20 baseline and the loop rules of section 26; no finding, source, date or claim changed, no source was re-opened and nothing after the 5 October 2026 cutoff was added. It is the paper's only figure with motion: dots run on its rings at relative speeds, not to scale. The motion stops under a reduced-motion setting, in print and in the download, and a control on the figure pauses it.

Figure text equivalents

The text equivalent of each figure, printed once here rather than beside the drawing. Each figure’s caption, status badge, units, scope and sources are printed with the figure itself.

Figure 1: Five control questions

ILLUSTRATIVE

Five rights drawn as five cards in two rows (three, then two, with a note on the roles in the sixth place), a set with no arrows between them: the paper lists them in one sentence and draws no order or cause. Each card shows the question, the answer in a word, the accountable role (a proposal from section 27) and the evidence to request. 1 Set the goal. Who sets the goal? Owner. Accountable: the CMO, with the CFO. Ask to see one signed page with the goal, the margin floor, the caps and the exclusions. 2 Grant authority. Who may act? Grant. Accountable: the budget owner issues the grant; security and legal can revoke it. Ask to see a grant that names tasks, accounts, sums and time windows, and nothing else. 3 Stop spending. Who can stop the spend? Stop right. Accountable: the CMO or the budget owner; the proof lead on a data-use breach. Ask to see a rehearsed stop that reaches every connected tool, queue and booked commitment. Its detail names the technical revocation authority (security and legal can revoke any grant; the brand can revoke authority and every active agent stops or falls back safely) and the contractual cancellation limits (a cap is not a liability ceiling, because booked commitments and a failed stop can exceed it; spend stops within the agreed window and partial actions are listed for repair). 4 Challenge evidence. Who can challenge the proof? Separate proof. Accountable: the proof lead, or an outside reviewer with access and freedom to disagree. Ask to see direct log access, success criteria frozen before each test, null results kept, and attribution and causal estimates as separate documents. 5 Retain usable learning. Who keeps the learning when the deal ends? Tested exit. Accountable: the product lead and procurement run the exit drill; the data steward owns export formats. Ask to see an exit drill: an export plus a rebuild of one workflow elsewhere, with lost state, cost, downtime and licences documented. Selecting a card reveals its contract clause, a log example, its failure mode and the chapters to read; the table below holds the same. Roles are the section 27 proposal: a title grants no legal authority, the contract and the grant do; one person may hold several roles and no role is a separate hire. No source in the record shows a brand holding all five rights. Illustrative framework, not an observation of any product.

Reader question: What are the five rights that decide what a brand keeps, and what grants each one?

Figure 2: From a claim to evidence

OBSERVED

Default view: five cards in one row, joined by thin reading-order arrows. Gate 1, announcement, shows a dated claim about what a firm can do, or about a deal; it does not show that the thing was built, used or closed. Gate 2, documented implementation, shows named systems, tasks, inputs, outputs and limits; not that any client uses it. Gate 3, named live use, shows a named deployment with role, market, date and approval scope; not that it works at scale or pays. Gate 4, meaningful scale, shows active clients or spend, repeat use, duration and failure rate; not that the client gained. Gate 5, independently demonstrated value, shows a credible test with full cost, quality and added effect; it supports value within the tested setting, and durability, transferability and performance at larger scale still need evidence. Selecting a gate lists the paper’s examples at that gate with their claim ids and sources; no example reaches gate four or five for a named group account, a gap in the record, not a score. Not every gate is needed for every useful study; empty evidence is not a zero score and not proof of no capability. Second view, protocol stages: a grid of six items (IAB Tech Lab AAMP 3.0 with OpenProposal; AAMP 2.0; AdCP; platform agent products; Google automated bidding and Amazon automated buying products; Horizon Blu APIs and MCP connections) against the paper’s six stages (proposal, specification, reference implementation, compatibility test, deployment, measured adoption), each cell in a few words with the full section 17 wording in the table, and a column for the record and version verified (two repositories cloned by git and pinned by commit; stored passages for the rest; none opened for the platform agent products). The AdCP row carries one

version record: the version analysed is 3.2.1, generally available from 30 September 2026, at commit 6adfebe2; a later patch tag, v3.2.2 of 4 October 2026, inside the cutoff, is recorded beside it with what its changelog entry changes (documentation), and is not substituted for the version analysed. Not established is a gap in inspection, not a zero. Shared transport is not shared meaning, rights or state. Tests run by this review: none.

Reader question: For any claim in this paper, which gate has the record reached, what is its boundary, and what would move it?

Figure 3: Who owns what, since when

OBSERVED

Observed record, selected not exhaustive: six groups and the independent agency-group case (Horizon Media), from Tier 1 and Tier 2 records to the 5 October 2026 cutoff, each shown in two views drawn from one data file. The ownership view lists the parent and its entities; each row gives the entity's name and one line with its type, its relation to the parent and its transaction status. A selected row's detail, the group tables and this text keep five separate fields, plus the passages in this paper each row rests on: the entity type (parent, operating group, agency brand, business, product, trading desk, joint venture, programme, named client, reported figure), the relation to the parent (owned, stake kept, joint venture, programme, named client use, proposed purchase, announced, structure or plan, part of the group, built or launched, reported figure, counterweight), the transaction status with its effective date and the date of the record it rests on, and a verification line that keeps "owned on the record" apart from "ownership verified by this review" (a completion record held). The capability-mapping view, kept apart from the ownership structure, shows six capability columns (strategy or consulting, creative or production, data or identity, activation, commerce, measurement) are analysis: a filled mark means the paper's words for the asset fall in that column on a completed record; an outlined mark means the record is announced, proposed or partly sold; an empty cell means not assessed in this review, not that the capability is absent. A count of marks is not a score. Two labels stand for empty evidence throughout: "Not assessed in this review" and "Not shown in the reviewed record"; neither means the capability does not exist. Publicis: Publicis Groupe (Parent; Parent, the owner of the rows below; dated events: no dated event drawn; capability columns: none mapped). Chief Client Officers across disciplines (plan of 2 December 2015) (Operating group · Structure or plan; In place Jan 2016 · record 2 Dec 2015; A plan record; not today's roster; dated events: reorganised 2 Dec 2015; capability columns: none mapped). Digitas (Business · Owned; Purchase announced 20 Dec 2006; Owned on the record; not verified by this review: closing not checked; dated events: agreement announced 20 Dec 2006; capability columns: Creative / production (announced or proposed), Activation (announced or proposed)). Sapient (Publicis.Sapient) (Business · Owned; Completed 6 Feb 2015; Ownership verified by this review: completion record held; dated events: completed 6 Feb 2015, impairment booked 9 Feb 2017; capability columns: Strategy / consulting). Epsilon (Business · Owned; Completed 1 Jul 2019 · record 2 Jul 2019; Ownership verified by this review: completion record held; dated events: completed 2 Jul 2019; capability columns: Data / identity). Lotame (Business · Owned; Completed Apr 2025 · record H1 2026 report, note 3.2; Owned on the company's own note; to be re-read; dated events: completed Apr 2025; capability columns: Data / identity). LiveRamp (Business · Proposed purchase; Agreed 17 May 2026 · pending, status 16 Sep 2026; Not owned: pending at the cutoff; never shown as owned; dated events: agreement announced 17 May 2026, still pending, status checked 16 Sep 2026; capability columns: Data / identity (announced or proposed)). Retail media products (June 2023) (Product · Built or launched; Launched 16 Jun 2023; Issuer's own description; no deal record; dated events: launched 16 Jun 2023; capability columns: Activation, Commerce). WPP: WPP (Parent; Parent, the owner of the rows below; dated events: no dated event drawn; capability columns: none mapped). GroupM, renamed WPP Media (28 May 2025) (Operating group · Part of the group; Renamed 28 May 2025; Issuer's own description; no deal record; dated events: renamed 28 May 2025; capability columns: Activation). Kantar (60% sold December 2019; 40% retained at that date) (Business · Stake kept; Part sold 5 Dec 2019; The stake at that date; not checked since; dated events: sold in part 5 Dec 2019; capability columns: Measurement (announced or proposed)). InfoSum (Business · Announced; Purchase announced 3 Apr 2025; Not drawn as owned: completion not in the record held; dated events: agreement announced 3 Apr 2025; capability columns: Data / identity (announced or proposed)). Mindshare (Agency brand · Part of the group; Named in the record 28 May 2025; Issuer's own description; no deal record; dated events: no

dated event drawn; capability columns: none mapped). Wavemaker (Agency brand · Part of the group; Named in the record 28 May 2025; Issuer's own description; no deal record; dated events: no dated event drawn; capability columns: none mapped). EssenceMediacom (Agency brand · Part of the group; Named in the record 28 May 2025; Issuer's own description; no deal record; dated events: no dated event drawn; capability columns: none mapped). Xaxis (trading desk, launched June 2011) (Trading desk · Built or launched; Launched Jun 2011; Issuer's own description; no deal record; dated events: launched Jun 2011; capability columns: Data / identity, Activation). Omnicom and IPG: Omnicom (Parent; Parent, the owner of the rows below; dated events: no dated event drawn; capability columns: none mapped). Omnicom Media (Operating group · Part of the group; In place on the record 28 Aug 2026; Issuer's own description; no deal record; dated events: reorganised 2019; capability columns: Activation). Interpublic (IPG), a lineage inside Omnicom from 26 November 2025 (Business · Owned; Completed 26 Nov 2025; Ownership verified by this review: completion record held; dated events: completed 26 Nov 2025; capability columns: Data / identity, Activation). Flywheel (Business · Owned; Completed 2 Jan 2024; Ownership verified by this review: completion record held; dated events: completed 2 Jan 2024; capability columns: Commerce). Acxiom (IPG lineage) (Business · Owned; Named in the record Jul 2026; Owned on the record; not verified by this review: the record is unread; dated events: named use reported Jul 2026; capability columns: none mapped). Hearts United (Hearts & Science with Mediahub, 28 August 2026) (Agency brand · Part of the group; Renamed 28 Aug 2026; Issuer's own description; no deal record; dated events: renamed 28 Aug 2026; capability columns: none mapped). OMD, PHD, UM, Initiative (siblings named in the Hearts United release) (Agency brand · Part of the group; Named in the record 28 Aug 2026; Issuer's own description; no deal record; dated events: no dated event drawn; capability columns: none mapped). Omni (the group's shared platform) (Product · Part of the group; Named in the record 2 Jan 2024; Owned on the record; not verified by this review: the record is unread; dated events: no dated event drawn; capability columns: none mapped). Dentsu: Dentsu (Parent; Parent · Impairment booked FY2025; dated events: impairment booked FY2025; capability columns: none mapped). Aegis (Dentsu Aegis Network from 26 March 2013) (Business · Owned; Completed 26 Mar 2013; Ownership verified by this review: completion record held; dated events: agreement announced Jul 2012, completed 26 Mar 2013; capability columns: Activation). Merkle (100% from 15 April 2020) (Business · Owned; Completed 15 Apr 2020; Ownership verified by this review: completion record held; dated events: completed 15 Apr 2020; capability columns: Data / identity). Havas: Havas (listed on Euronext Amsterdam 16 December 2024) (Parent; Parent · Listed 16 Dec 2024; dated events: listed 16 Dec 2024; capability columns: none mapped). Havas Media Network (Operating group · Part of the group; In place on the record 29 Sep 2025; Issuer's own description; no deal record; dated events: no dated event drawn; capability columns: none mapped). Horizon Global (joint venture with Horizon Media, 50/50) (Joint venture · not a merger; Announced 29 Sep 2025; Issuer's own description; no deal record; dated events: announced 29 Sep 2025; capability columns: Activation (announced or proposed)). Stagwell: Stagwell (combined with MDC 2 August 2021) (Parent; Parent · Recorded Q3 2025; dated events: completed 2 Aug 2021, recorded Q3 2025; capability columns: Creative / production, Activation, Measurement). Enterprise-software revenue label (\$16 million committed, H1 2026) (Reported figure · not an entity; Reported in a filing of 30 Jul 2026; Not an entity: no product entity in the record held; dated events: recorded 30 Jul 2026; capability columns: none mapped). S4 Capital (counterweight, not part of Stagwell) (Parent · Counterweight; Not part of the group · results Mar 2026; dated events: no dated event drawn; capability columns: none mapped). Horizon Media (the independent agency-group case: a private parent and its agency; independence means outside the six groups, not the absence of a holding structure): Horizon Media Holdings (private parent) (Parent; Parent, the owner of the rows below; dated events: no dated event drawn; capability columns: none mapped). Horizon Media (the agency issuing the releases) (Business · Part of the group; Named in the record 18 Dec 2025; Issuer's own description; no deal record; dated events: no dated event drawn; capability columns: none mapped). HorizonOS (launched 18 December 2025) (Product · Built or launched; Launched 18 Dec 2025; Issuer's own description; no deal record; dated events: launched 18 Dec 2025; capability columns: Strategy / consulting, Activation). Blu (announced 18 June 2026) (Product · Announced; Announced 18 Jun 2026; Issuer's own description; no deal record; dated events: announced 18 Jun 2026; capability columns: Data / identity (announced or proposed), Activation (announced or proposed)). HorizonOS Labs (partner programme) (Programme · access, not ownership; Cohort announced 1 Oct 2026; Issuer's own description; no deal record; dated events: cohort announced 1 Oct 2026; capability columns: none mapped). SharkNinja (named client; audience collaboration with retailers) (Named client · Named client use; Use reported 18 Jun 2026; Issuer's own description; no

deal record; dated events: named use reported 18 Jun 2026; capability columns: none mapped). Horizon Global (joint venture with Havas Media Network, 50/50) (Joint venture · not a merger; Announced 29 Sep 2025; Issuer's own description; no deal record; dated events: Announced 29 Sep 2025 (drawn on the Havas tab); capability columns: none mapped). Twenty deal cards give what was acquired, the intended connection (an inference), the evidence of implementation and of client benefit, and the gate the record reaches (announcement for most; documented implementation for the Publicis retail media products; named live use for IPG's Acxiom result and Horizon's Blu with SharkNinja). Client benefit is not shown in the reviewed record for any group deal; the one figure held is a vendor-reported recognition rate with no control group. LiveRamp is proposed (17 May 2026) and still pending at the 16 September 2026 check; it is never shown as owned. The Stagwell \$16 million row is a revenue label in a filing, not a product entity. Gaps are shown as gaps: the Digitas closing, the Omni releases, IPG's 2018 Acxiom purchase, IPG's disposals and impairment, AKQA and Grey, the Stagwell 2026 products, the Havas purchases, the trading-desk lineages and two abandoned deals were carried forward unread. Downloads: per group, the ownership view and, as a second file, its capability mapping; and one reading sheet of the seven ownership views.

Reader question: Who owns what, since when, through which kind of relationship, and what does the record show the deal did?

Figure 4: Dated events, 2006–2026

OBSERVED

Observed record, selected not exhaustive: 26 dated events from Tier 1 and Tier 2 records, one swimlane per group plus Horizon Media as an independent, on one linear axis from 2006 to the 5 October 2026 cutoff, with a second view of 2024 to the cutoff on its own axis. Marker shape gives the entity level (parent deal, operating group, agency brand, product, partnership, financial record); fill gives the status (closed or launched; announced only; pending at the cutoff). Nothing ranks the groups. Publicis: Digitas agreed, 20 Dec 2006 (parent; announced). Publicis: Sapient completed, 6 Feb 2015 (parent; closed). Publicis: Power of One plan, 2 Dec 2015 (group; closed). Publicis: €1.44bn impairment, 9 Feb 2017 (record; closed). Publicis: Epsilon completed, 2 Jul 2019 (parent; closed). Publicis: Retail media launch, 16 Jun 2023 (product; closed). Publicis: Lotame bought, Apr 2025 (parent; closed). Publicis: LiveRamp agreed, pending, 17 May 2026 (parent; pending). WPP: Xaxis launched, Jun 2011 (product; closed). WPP: Kantar 60% sold, 5 Dec 2019 (parent; closed). WPP: InfoSum announced, 3 Apr 2025 (parent; announced). WPP: GroupM → WPP Media, May 2025 (group; closed). Omnicom: Specialists into brands, 2019 (group; closed). Omnicom: Flywheel completed, 2 Jan 2024 (parent; closed). Omnicom: IPG completed, 26 Nov 2025 (parent; closed). Omnicom: Hearts United formed, 28 Aug 2026 (brand; closed). Dentsu: Aegis completed, 26 Mar 2013 (parent; closed). Dentsu: Merkle 100%, 15 Apr 2020 (parent; closed). Dentsu: ¥396bn impairment, FY2025 (record; closed). Havas: Euronext listing, 16 Dec 2024 (parent; closed). Havas: Horizon Global JV, 29 Sep 2025 (partnership; announced). Stagwell: MDC combination, 2 Aug 2021 (parent; closed). Stagwell: Growth method changed, Q3 2025 (record; closed). Horizon: HorizonOS launched, 18 Dec 2025 (product; closed). Horizon: Blu announced, 18 Jun 2026 (product; announced). Horizon: Labs cohort 3, 11 of 600+, 1 Oct 2026 (partnership; announced).

Reader question: When did each group widen or narrow its scope, at which entity level, and which of today's links are still only announced or pending?

Figure 5: How to read a deal

ILLUSTRATIVE

A template from section 5 with a worked-case selector. Six numbered links in reading order, numbers and a hairline only, no arrows: 1 prior gap (what was missing before the deal; filled by the buyer's dated statement of purpose); 2 asset or skill (what was bought, when and for how much where stated; agreement and completion records kept separate); 3 intended connection (how the asset was meant to join other work; a dated integration plan); 4 use on record (did a named system, team or product connect; a product description, a named deployment or a reorganisation record); 5 client effect (did a client do better after all costs; a matched test with full cost and added effect); 6 group economics (what the group's

own accounts show; revenue, net revenue, margin, impairment with definitions). A separate panel holds two checks: the countercase (what went wrong, was sold or was reversed: impairments, disposals, separations, lost scope) and the rival reading (what else explains the same record: cross-selling, cost removal, investor narrative, the cycle). A third panel holds three value pools kept apart: the acquired team may gain distribution; clients may gain a better outcome, a lower total cost or less coordination work; shareholders may gain cost removal; none shows the others. Under every link a status word: observed (a dated record), attributed (the source's own words), inference (the paper's reading, never a deal fact), not applicable, or one of two labels for empty evidence, "Not assessed in this review" (no record opened or held for the link) and "Not shown in the reviewed record" (the paper looked and found none); neither means the capability does not exist. With a case chosen the badge reads "Worked case: name · status" and a one-line summary (the asset link's own words, with announcement, pending or disposal conspicuous) sits above the checklist; the template badge returns only after Clear case. Twelve worked cases from the group chapters, by group and date: Publicis: Digitas (agreed 20 December 2006; closing not checked here), Sapient (closed 6 February 2015; the 2016 impairment is the countercase), Epsilon (closed 2 July 2019), Lotame (April 2025 per company note), LiveRamp (agreed 17 May 2026; pending at the cutoff; never owned). Omnicom: Flywheel (closed 2 January 2024), Interpublic (closed 26 November 2025). Dentsu: Aegis (agreed July 2012; closed 26 March 2013), Merkle (100% from 15 April 2020). WPP: Kantar (60% sold 5 December 2019; a disposal), InfoSum (announced 3 April 2025). Stagwell: Stagwell and MDC (combined 2 August 2021). Client effect is not shown in the reviewed record for any case; the acquired-team and shareholder pools are not assessed in this review for any case; no case ranks one group above another.

Reader question: How do I read any holding-company deal link by link, and which value pools can diverge?

Figure 6: Value created and value captured

ILLUSTRATIVE

Analysis of the section 14 table, not a ranking. Nine mechanisms sit under the party that gains as volume rises; the brand pays where a payer is stated. Agency or group: labour or service fees; spend-linked fees; principal media spreads; outcome or performance fees. Either the agency or the provider or platform, as provider, supplier or incumbent: technology subscriptions and platform fees; data services; switching costs. Provider or platform: platform spend and in-platform automation. Either side, depending on dates: working capital and float. For each, whether the client can gain (every answer is conditional; switching costs only indirectly) and the evidence that would show it. Labour fees: if the fee is below the brand's own cost; evidence: approved useful output per total hour, quality held, rework falling. Spend-linked fees: only if the extra spend is profitable at the margin; evidence: marginal client profit per extra unit of spend. Media spreads: if the bundled price for comparable inventory is below the agency route; evidence: the role disclosed per transaction and a matched comparison. Subscriptions: if the tool replaces a costlier process and exit charges are bounded; evidence: true usage, overlapping licences removed, exit cost priced. Data services: if match quality and lawful use improve targeting or measurement and the data can move; evidence: provenance, match rate, use rights, portability. Switching costs: indirectly, if continuity lowers rework and the client accepted the cost knowingly; evidence: an exit drill priced before signing. Platform spend: on bounded tasks with reliable local signals; evidence: an outside cross-platform added-demand test. Outcome fees: if the metric is added profit over a matching window; evidence: the fee tracks added profit after substitution, discounting and brand effects. Working capital: either side, by payment dates; evidence: payment and due dates. Selecting a mechanism draws the payer, the earner and the cost bearer, with money and evidence as separate relationships, and states who earns as spend rises, who funds failed tests (the brand, unless the contract moves it) and who checks results (the brand's proof function). The same explorer and data serve section 18. This review does not compare the mechanisms' scalability or client returns, and no client-side test for a named group account is on record.

Reader question: For each way a group, agency or platform captures value as volume rises, can the client gain too, and what evidence would show it?

Figure 7: The task profile

ILLUSTRATIVE

Illustrative profile; every cell is a design definition from section 16, no product was tested and nothing is measured. Six dimensions, each recorded as a stated condition, never as a point on a scale: recommendation or execution; adaptive discretion (preset rules, or adapts within a stated scope); system scope (one system or many); spend and time limits (amount and period); approval (who, how often); reversibility (what a stop undoes). Three example tasks. A content tool: recommends; proposes only; no execution surface; no spend, per request; a human edits or ignores every output; nothing to undo, but cancellation and data-use controls are still needed. A human-approved cross-system workflow: executes after approval; preset rules; many systems; no spend until approved, per campaign; a named approver at every step; reversible until booked. A single-platform budget optimiser: executes; adapts within a stated scope; one system; material spend within the platform cap, for the campaign; approval per campaign, not per bid; partly reversible, because a stop halts new actions and booked inventory stays committed. The point: autonomy, integration breadth, spend and reversibility are separate questions; a cross-system workflow can be low autonomy and broad integration, a single-platform tool high discretion with material spend. The page compares at most two tasks at a time; this table holds all three.

Reader question: Why is one ordered autonomy ladder misleading, and which separate questions should a grant answer for each task?

Figure 8: Seven fee routes

ILLUSTRATIVE

Analysis of the section 18 table and the working terms in section 3. Agency buying: a supplier acts for the client under a mandate. Money: the brand pays the agency a disclosed fee, and pays the net media cost to the publisher through the agency under the mandate, counted once at the price the brand pays. Advice: the agency's plan and recommendation, acting for the client. Evidence: delivery logs and reports, tested by the brand's proof function. Principal resale: a supplier buys media on its own account and resells it; the contract sets price, risk and rights. Money: the brand pays the agency one bundled price; the reseller's spread sits inside that price and is disclosed, never added again as a second line; the reseller buys from the publisher on its own account. Advice: a recommendation from a seller with inventory of its own. Evidence: the role disclosed per line and an audited comparison on equal inventory and terms. The two routes may coexist inside one mandate, stated line by line. Use of principal media rose to 58% of 114 qualified client-side ANA members in the past year, from 47% in the prior wave: a member sample, not a share of spend; no spread is disclosed by any group reviewed. No arrow or box carries an amount. Seven fee routes, the brand paying on each: labour or service fees (earns the agency or group when billable scope or hours rise; the agency carries delivery cost and the brand carries rework it cannot see); spend-linked fees (earns the agency or group when billable spend rises; the brand carries unprofitable spend); principal media spread (earns the reseller when the bundled price exceeds its cost; the reseller carries inventory risk it chose and the brand carries the price it cannot compare); technology subscriptions and platform fees (earns the provider or platform from seats, usage and renewals; the brand carries overlapping licences and exit cost); data services (earns the provider or platform from permitted data work; the brand carries unlawful use and lost provenance); strategic analytics and advice (earns the agency or group; the brand carries advice shaped by a sales interest); performance or outcome fees (earns the agency or group when the agreed metric clears the hurdle; the brand carries fees paid on harvested or shifted sales). Every client saving is conditional on the route's test; the table lists the condition, the conflict and the evidence needed, and the selected route shows who earns as spend rises, who funds failed tests (the brand, unless the contract moves it) and who checks results (the brand's proof function). Fees are a design choice, not neutrality.

Reader question: On each fee route, who earns as volume rises, who can save, and who carries the risk?

Figure 9: Spending to proof: two clocks

ILLUSTRATIVE

An illustration with assumed inputs, not a forecast and not a profit model. Inputs: proposed daily media spend \$10,000, divisible; the first read after 7 days; an authorised cap of \$90,000, an explicitly borrowed illustration from the section 26

retailer mandate because section 18 states no cap, not a recommended brand limit; outstanding commitments \$0, an editable assumption; the first read lands on time; the stop bites the same day (no slow-stop overshoot assumed). The cap is checked before each day's spend: allowed spend each day = $\max(0, \min(\text{proposed daily spend}, \text{cap} - \text{spent} - \text{commitments}))$, so a cap below one day's spend clips the first day, a cap that is not a multiple of the daily spend clips the last day, and commitments at or above the cap leave nothing to spend. Read schedule: the first read lands on day 7; after it each day's spend is read 7 days later, a rolling lag, not one assessment of all earlier spend. Lines, each dollar on one line only: spent and awaiting evidence (money out whose read is not yet due: every dollar before the first read, then the last seven days of spend); spent with a read received; outstanding commitments (booked but not yet spent, which a stop may not undo; never also counted as spent); remaining authority (the cap less spent less commitments, never below zero); spent past the stop (excess from a slow stop, an assumption that is off by default); and exposure beyond the cap (spent plus commitments less the cap, when positive), shown as its own number rather than hidden by clipping remaining authority to zero. Day by day at the defaults: day 0, nothing spent, \$90,000 of authority; day 1, \$10,000 spent and awaiting, \$80,000 left; the same each day to day 7, the first read due, with \$70,000 spent, all of it awaiting evidence, and \$20,000 left; day 8, \$80,000 spent, \$70,000 awaiting, \$10,000 left; day 9, \$90,000 spent, the cap reached with that day's allowed spend \$10,000, further spend needs an approval or a pause; day 10, nothing more spent, \$60,000 still awaiting a read; spent past the stop and exposure beyond the cap are \$0 on every day. The stop point is the first day spent plus commitments reach the cap (day 0 when commitments already reach it), or the day a scheduled read is due and has not landed, whichever comes first; from that day nothing more is spent unless the slow-stop assumption is set, and then the excess is its own line. A read is a read, not guaranteed causal proof; a fast proxy never counts as a causal result. Spend awaiting evidence is not a loss; it is spend whose return is not yet known, carried by the party that funds it unless the deal moves it. A cap is a brand-set rule about exposure, not a liability ceiling; booked commitments and a failed stop can exceed it, and the figure then shows the excess as a number. Media expenditure only; agency fees sit on a separate line and are not in this illustration. Cash exposure is not a financing expense: the section 20 boundary carries financing and commitments as its own line. The live view lets a reader change the six inputs; the static view and print use the labelled defaults.

Reader question: How much money is out before the first delayed read, what is still booked, and where does more spend need a sign-off?

Figure 10: Who keeps a task saving?

ILLUSTRATIVE

Illustrative model; every input is an assumption; arbitrary units with a base fee of 100 that stays 100 in every case. Formula: $\text{cost} = \text{labour} \times (1 - \text{saving_pct} \div 100) + \text{other} + \text{new cost}$; $\text{profit} = \text{fee} - \text{cost}$; $\text{client fee saving} = 100 - \text{fee}$. Assumed: fee 100, labour 60, other 20, labour saving 25%, new tool and review cost 5. Base: no saving, no new cost; labour 60, other 20, cost 80, profit 20. Agency keeps the net saving: fee 100; labour 45, other 20, new cost 5, cost 70, profit 30, client saving 0. Client keeps the net saving: fee 90; the same cost 70, profit 20, client saving 10. Equal split, a new illustrative preset: fee 95; cost 70, profit 25, client saving 5. Added cost absorbs the saving: new cost 15 at fee 100; cost 80, profit 20, client saving 0. The live calculator uses the same formula and presets, keeps the five inputs, and draws a cost above the fee as a shortfall against the fee line, never as a negative bar. One task, output and quality held fixed, arbitrary units. Nothing here measures added sales.

Reader question: If automation cuts the labour in one task, where does the money go under a held fee and under a cut fee?

Figure 11: How judgment creates value

ILLUSTRATIVE

The paper's proposed operating logic, not a measured causal growth loop; the arrows mean next, not cause. Six stages in reading order, each with one measure or diagnostic question, read one at a time by selecting a stage (the drawing carries the stage names and a prompt; the table lists all six): 1 Brief: The business problem, the audience, the constraints and the

measure of success, written before work starts. Measure from the paper: Brief quality rated by both sides. 2 Shared customer understanding: One picture of the customer that both sides can see and challenge. Proposed diagnostic question, not a measure (the figure's data file; the paper names no measure for this stage): Can both sides state the same customer picture, and can either side disagree with it on the record? 3 Creative choices: Made and defended by a named lead, with the right to reject volume. Measure from the paper: A brand-grounds rejection stands; rejected volume is counted as a cost, never per prompt. 4 Timely client decisions: Approval windows agreed in advance and measured. Measure from the paper: Approval delay in days. 5 Useful feedback: Specific, reasoned and recorded, so the next round improves. Measure from the paper: Rework share of hours. 6 Learning kept: What was tried, what worked and why, in a form the next brief can use. Proposed diagnostic question, not a measure (the figure's data file; the paper names no measure for this stage): Does the next brief cite what the last round tried, what worked, what failed or showed nothing, and why? Learning kept leads back to the next brief. Two horizons side by side. Near-term delivery: Approved useful work per total hour; Brief quality, approval delay, rework share; Attributed results, as credit not cause. Long-term brand value: The brand-health read, kept beside delivery metrics; Harvest and build run as a portfolio, each with its own budget, horizon and measure; No number assigned to a brand effect the paper has not measured. Four ways a gain can be false: harvesting; channel shifting; discounting profit away; weakening brand health. Support of a bounded kind: WFA respondents ranked briefing, trust and feedback above financial incentives; rankings, not causal effects, and the source was not re-opened. No measured loop: the paper assigns no number to a brand effect and holds no causal study of briefing or feedback.

Reader question: What makes the work better, and what one measure or question tells me whether each step is working?

Figure 12: The present-day loop and the three paths, in one layout

ILLUSTRATIVE

Simplified present-day map, not a measured sequence; shown by default. Seven stages run as a closed loop: 1 Brief (brand) → 2 Planning (agency planner; brand approves) → 3 Activation (agency or platform) → 4 Measurement (agency) → 5 Attribution and reporting (agency or platform) → 6 Optimisation (operator) → 7 Revised plan (brand with agency). Eight tasks are shown as seven stages: attribution and reporting are split from measurement so that credit and cause stay distinct. Arrows: 1 → 2 information (brief); 2 → 3 approval (approved plan and budget); 3 → 4 execution (work runs); 4 → 5 information (credit by each platform's rule); 5 → 6 information (reports); 6 → 3 fast feedback (bounded changes within rules: the short loop on proxies); 6 → 7 information (what changed); 7 → 1 delayed feedback (revised goals); 7 → 2 approval (a budget change beyond the approved limits); 4 → 7 delayed feedback (causal evidence, later, if tests are run: the long loop, drawn as the outer path). Attribution assigns credit; it does not prove incremental growth. Each stage box names who runs the stage and, on an approval line, whose sign-off its output needs: the brand sets the brief (1) and decides the revised plan (7); the agency planner develops the plan and the brand approves the plan and budget (2 → 3); activation, measurement, attribution and optimisation run within the approved plan, with bounded changes inside the rules and any budget change beyond the approved limits returning to the brand (7 → 2). "Within the approved plan" or "within the grant" on an approval line means no new sign-off inside limits already granted. The detail panel and the tables give three fields per stage, who runs it, whose sign-off its output needs and what is delegated, beside the paper's own "Owner of decisions" cell, which is quoted unchanged. Each stage carries a layer tag (authority, execution or evidence), this figure's reading of the owner and measurement columns. The same drawing shows A, B and C in the same positions, with their own stages (each chapter's text equivalent lists them), their band (A: records, proof and learning brand-held by contract; B: the agency-run service; C: inside each platform, two or three lanes) and their conditions. Layers can be shown or hidden: execution (solid dark), authority (thick orange arrows and the approval line on each stage), evidence (dashed blue information and dotted green feedback, with fast checks and delayed evidence as two clocks). Previous step, Next step and Reset walk the seven stages; stepping never produces evidence. Switching path keeps the selected stage where the paper maps a like stage (brief 1; planning 2, which is stages 2 and 3 in A; activation 3 to stage 4; measurement 4 and attribution 5 to stage 5; revised plan 7). Optimisation 6 has no stage in A, B or C (it is the fast 5-to-4 loop), and test and review 6 has no baseline stage (causal evidence is the delayed 4-to-7 edge); neither is forced into equivalence. The comparison table beneath the drawing states, for the baseline and each path: who sets goals

(brand marketing lead; CMO with finance and data leads; brand goal owner with the lead agency designing the service; brand with finance or an adviser); who runs the work (agency or in-house operator and platform systems; a named specialist within the grant; a named agency team or agent under contract; each platform within its own system); who approves more spend (the agreed brand approval beyond limits; the brand budget owner, with any wider grant returning to stage 3; the brand reserving large moves; the brand or a delegated agency owner per platform); who tests growth (later, if at all, with attribution assigning credit; a brand-funded proof team; a proof function with its own mandate; a brand-funded team separate from media sellers); who retains learning (decision history may not be kept; brand decision history and tests with vendor model state behind; brand facts and history, agency methods, vendor model state; the portfolio decision history with hidden state behind); who carries costs (the party funding the spend carries the spend-to-proof gap; in A the brand pays retained staff, integration, data, measurement and exit directly and funds failed tests; in B services and technology are paid through the partner at itemised prices while retained staff, proof and exit stay with the brand; in C platform spend counts once as media, platform fees as technology, and the brand funds most failed demand tests). Open is not good and closed is not bad; a scenario is not a forecast. No arrow encodes timing, probability or effect size.

Reader question: What are the stages between a brief and a revised plan, who owns each today, which loop is fast (proxies) and which is slow (causal evidence), and how does each of the three paths change who decides, runs, tests and keeps?

Figure 13: From PAM to the governed loop

CONDITIONAL

Two parts. Above, the present-day loop of section 20 as one ring, PAM for plan, activate and measure: 1 Brief (plan) → 2 Planning (plan) → 3 Activation (activate) → 4 Measurement (measure) → 5 Attribution and reporting (measure) → 6 Optimisation (activate) → 7 Revised plan (plan) → back to the brief. Beside it, where PAM breaks: it closes on credit, not cause; one wheel hides two clocks; the learning may leave. Below, a possible future: three concentric loops around a hub of the five rights of section 1 (set the goal, grant who may act, stop the spend, challenge the proof, keep the learning), held by contract in any path. Inner, the fast loop, hours to days: act → signal → adjust within the grant. It changes execution, never the mandate, and a fast signal is never proof. Middle, the learning loop, weeks: 1 Mandate (goal, floor, caps) → 2 Grant (task, sum, window) → 3 Execute (work runs in the grant) → 4 Record (a log for each action) → 5 Prove (pre-registered test) → 6 Permit (authority up or down) → 7 Learn (decision history kept). Evidence changes what agents may do next. Outer, the planning loop, months: it revises the mandate itself. Only the learning and planning loops change the mandate, and only through a logged approval. A stop line runs from the hub: the stop works at any time and reaches every loop, tool, queue and booking; it does not cancel commitments already booked. On screen a dot runs on each ring and the stage under it lights up. Adjacent rings turn in opposite directions only to keep them apart; the speeds are relative, not to scale. The motion stops under a reduced-motion setting, in print and in the download, and a control pauses it. A conditional design, not a forecast.

Reader question: What changes in the marketing loop when software takes on more of the work, and what stays the same?

Figure 14: A. Brand-controlled, open orchestration

CONDITIONAL

Chapter view of path A: the seven stages in the shared layout with their layer tags (authority, execution, evidence), the brand-held band and the two clocks; the full loop with every arrow is the editable download and the comparison figure in section 20. A conditional model that holds only if its assumptions hold; not a forecast. Seven stages run as a closed loop: 1 Set goals and rights (brand) → 2 Plan and create (brand and agency) → 3 Approve the plan (brand) → 4 Run bounded work (specialist) → 5 Collect signals (brand) → 6 Test and review (proof team) → 7 Keep learning; replan (brand). Arrows: 1 → 2 information (brief); 2 → 3 information (proposal); 3 → 4 approval (the grant: scope, cap, stop right); 4 → 5 execution (logged work); 5 → 4 fast feedback (checks within the grant); 5 → 6 information (records); 6 → 7 delayed feedback (causal evidence, later); 7 → 1 feedback (revised goals); 7 → 3 approval (any wider grant). Who runs each stage

and whose sign-off its output needs: the CMO sets goals and rights with finance and data leads (1); the brand lead plans with the agency's creative lead for craft, and the plan goes to the budget owner as a proposal (2); the brand budget owner issues the grant (3); the named specialist runs bounded work within the grant, with no new sign-off inside it (4); the brand data steward holds the records (5); the brand-funded proof team tests and reviews (6); the brand plan owner replans, and any wider grant returns to stage 3 (7). Stages 5, 6 and 7 sit in a brand-held band: records, proof and learning, held by contract. No arrow encodes timing, probability or effect size.

Reader question: If the brand runs the core, who sits at the gate, where do the records, the proof and the learning live, and where do the two loops close?

Figure 15: B. Agency-led, managed orchestration

CONDITIONAL

Chapter view of path B: the seven stages in the shared layout with their layer tags, the agency-run band and the two clocks; the full loop with every arrow is the editable download and the comparison figure in section 20. A conditional model that holds only if its assumptions hold; not a forecast. Seven stages run as a closed loop: 1 Agree the mandate (brand and agency) → 2 Plan and create (agency lead; brand approves strategy) → 3 Grant authority (brand) → 4 Run managed work (agency) → 5 Report full cost (agency reports, brand reconciles) → 6 Challenge results (separate proof function) → 7 Review and renew (brand). Arrows: 1 → 2 information (mandate); 2 → 3 information (proposal); 3 → 4 approval (grant of scope; large moves reserved); 4 → 5 execution (logged work); 5 → 4 fast feedback (service checks within the mandate); 5 → 6 information (full-cost report); 6 → 7 delayed feedback (proof, later); 7 → 1 feedback (renewal); 7 → 3 approval (changes beyond the grant). Who runs each stage and whose sign-off its output needs: the lead agency designs the service and the brand goal owner sets the goal (1); the agency lead plans and creates and the brand approves the strategy (2); the brand grants the scope and reserves large moves (3); the named agency team runs managed work within the grant, with no new sign-off inside it (4); the agency reports full cost and the brand reconciles it (5); the proof function challenges the results (6); the brand budget owner reviews and renews, and changes beyond the grant return to stage 3 (7). Stages 4 and 5 sit in an agency-run band: the managed service, whose full-cost report is a condition to prove. No arrow encodes timing, probability or effect size.

Reader question: If an agency runs the system, which rights does the brand reserve, where does the service loop close, and who can challenge the result?

Figure 16: Two linked loops: client and capability

CONDITIONAL

Conditional, interpretive design inspired by the strategy, not Horizon's implemented architecture. Client loop, six stages: 1 set the goal; 2 form choices; 3 approve scope (approval gate); 4 run the work; 5 record signals; 6 test and revise, then revisit the goal. Fast checks from 5 to 4 stay inside the existing grant; delayed evidence from 6 to 1 and 2 can change the mandate. Capability loop, six stages: 1 define the gap; 2 shortlist candidates; 3 approve a pilot (approval gate); 4 run the test; 5 judge evidence; 6 adopt or stop, feeding the next gap. Cross-links: client evidence (stage 6) and needs (stage 2) define the next gap (capability 1); capability 6 returns options only to client stage 2; using one needs that account's stage 3 approval. Arrow kinds: dashed information arrows join stages 1 to 2, 2 to 3 and 5 to 6 in each loop and carry the three cross-links; a thick approval arrow runs from each stage 3 to stage 4; a solid execution arrow from each stage 4 to 5; dotted feedback arrows run from client 5 to 4, 6 to 1 and 6 to 2, and from capability 6 to 1. No arrow encodes timing, probability or a measured effect; the second loop is a capability loop, not proof of compounding value.

Reader question: How could an agency's capability search and a client's results reinforce each other, where are the approval gates, and what would show that they do?

Figure 17: C. Platform-dominated, fragmented autonomy

CONDITIONAL

Chapter view of path C: the seven stages in the shared layout with their layer tags, the platform band (two or three lanes, drawn as stacked cards; local runs, siloed signals and model state stay inside) and the two clocks; the full loop with every arrow is the editable download and the comparison figure in section 20. A conditional model that holds only if its assumptions hold; not a forecast. Seven stages run as a closed loop: 1 Set portfolio goals (brand) → 2 Build assets and rules (creative team and platforms) → 3 Approve each scope (brand, per platform) → 4 Optimise locally (each platform) → 5 Read siloed signals (platform reports; brand checks) → 6 Test across channels (brand-funded proof) → 7 Rebalance; learn (brand). Arrows: 1 → 2 information (goals); 2 → 3 information (assets, rules); 3 → 4 approval (a bounded budget per platform); 4 → 5 execution (local runs); 5 → 4 fast feedback (local checks within preset rules); 5 → 6 information (siloed reports); 6 → 7 delayed feedback (tests, later); 7 → 1 feedback (revised goals); 7 → 3 approval (rebalancing). Stages 4 and 5 are drawn as stacked boxes inside a platform band: several platforms, each its own silo. Who runs each stage and whose sign-off its output needs: the brand sets portfolio goals with finance or an adviser (1); the creative lead and the platform setup teams build assets and rules, the creative lead approving what the brand says (2); the brand, or an agency owner it delegates to, approves each platform's bounded budget and goal (3); each platform optimises within that scope, with no new sign-off inside it (4); the platforms report and the brand checks (5); the brand-funded team tests across channels (6); the brand budget owner rebalances, and rebalancing returns to stage 3 (7). No instant causal loop is assumed. No arrow encodes timing, probability or effect size.

Reader question: If platforms run execution, what must stay outside them, where does each loop close, and who reads across the silos?

Figure 18: Openness and client control

ILLUSTRATIVE

A two-by-two matrix from section 24. Columns: client control, weak (rights exist on paper or not at all; state does not move) or strong (usable decision, access, stop and exit rights; state moves). Rows: ecosystem openness, open (providers can be connected and replaced) or restricted (providers are hard to connect or replace); openness is the ability to connect and replace providers, not how many vendors a marketplace holds. Open and weak control, easy vendor swap, hard exit: many vendors, one operator the client depends on; for example a managed service the brand has chosen to depend on for now; vendor swaps are easy, operator exit is not. Open and strong control, broad choice, strong exit: broad choice and strong exit rights; for example the brand's decision core; can be costly to run, because the brand pays for the joins and the drills; no cell is ranked by cost. Restricted and weak control, limited choice, bounded task: for example one platform's auction buying; can still deliver good value on a bounded task at a good price, if the brand knows what it is giving up. Restricted and strong control, managed tools, strong exit: a tightly managed tool set with strong records and exit rights; for example a closed stack with export rights; closed does not mean captured. Selecting a cell reveals its failure mode and contract test; a file export alone is not a strong exit, because a brand that holds only files has files, not decisions. No scores; no firm is placed; no cell is good or bad by itself. Use per task, not per supplier.

Reader question: Are openness and client control the same question, and what does each cell of the matrix look like in practice?

Figure 19: What leaves with the brand

ILLUSTRATIVE

A proposal, not a reviewed contract; every right here is a contract condition, and no term is attributed to any named group. A grid of four asset classes against six exit conditions, with a second view for a successor operator. Brand records (default owner: Brand; on exit: moves): access — Direct: the proof function reads logs directly, not through the operator's dashboard.; export — Named export format for each item: accounts, audience definitions, creative source files and usage rights, briefs, decision logs, approval history, test designs and results.; reuse — The brand reuses freely; the agency may cite anonymised method from the test record.; licence — Not needed: owned by the brand under this proposal.; retention and deletion — Retained for the agreed period; deleted on schedule; the record of a failed pilot belongs to the brand.; rebuild — Sufficient to rebuild the decisions only together with the settings and rules that

produced them.. Successor: The successor receives the records in the named formats and the exit drill record. Client-specific rules and state (default owner: Brand, where the contract says so; on exit: moves if contracted); access — Access to settings, rules and tool versions through the proof function.; export — Settings and rules in a named format; any transferable agent state in a named format. Non-transferable state is listed as left behind.; reuse — The successor reuses the rules; the operator may not transfer learned state derived from this client's data elsewhere without written consent.; licence — None for the rules; learned state that is not transferable is priced or left behind.; retention and deletion — Retention and deletion schedules fixed in the contract.; rebuild — The first two kinds plus the settings and rules are what a rebuild needs. Files alone are not decisions.. Successor: The successor rebuilds the decision from records plus rules; portable client state is a contract condition, not a default. Agency methods and IP (default owner: Agency; on exit: licence or rebuild): access — No access right by default.; export — Not exported.; reuse — The agency may reuse general methods across clients. It may not pool this client's data inside them.; licence — Licensed, not transferred; no client data inside. Agency property is listed with its post-exit status and price.; retention and deletion — The agency's own.; rebuild — The successor rebuilds its own methods or licenses the incumbent's.. Successor: The successor licenses what it needs at a listed price, or rebuilds. Platform and vendor model state (default owner: Vendor or platform; on exit: left behind): access — Under the vendor's own terms; client rights depend on terms.; export — Not portable by default. Portable model weights are not promised.; reuse — Not reusable by the brand; the platform keeps its learned state.; licence — Priced if available; licence terms priced for post-exit use.; retention and deletion — The vendor's own schedule.; rebuild — Listed as left behind on exit; a successor relearns from the brand's records and rules.. Successor: The successor starts without the incumbent's model state unless it is licensed; owning the source data does not grant it. Successor view legend: moves = Exported in a named format and usable by the successor; moves if contracted = Moves only where the contract grants it; licence or rebuild = Licensed at a listed price, or rebuilt by the successor; left behind = Not portable; listed as left behind; relearned from the brand's records. An exit drill, an export plus a rebuild of one workflow elsewhere, is run before the contract passes its first year and repeated each year. Its cost, downtime and lost state are documented. A file export alone is not a successful exit; the settings and rules must move with the records. Owning the records does not grant the methods or the models. No exit drill result exists in the record for any brand or operator. A successful move to a new operator is a signpost to watch, not something seen in the record. On phones the grid gives way to four asset cards, each opening to its six conditions and the successor action.

Reader question: On exit, what can the brand take, what must it license, and what must a successor rebuild?

Figure 20: One retailer, three agents

ILLUSTRATIVE

An invented illustration from section 26: the retailer, the agents, the numbers and the outcomes are invented; no sales or vendor outcome is shown. States: proposed, policy gate, eligible, escalated, rejected, then sign-off, authority, executed, logged and later evidence. An agent proposes; the gate checks, in this order, exclusions and consent rules, the margin floor, the discount budget cap, the aggregate weekly shift limit, tranche and authority, the uncertainty cap and creative-lead sign-off. Any failed hard constraint rejects; a limit that needs a human escalates to the budget owner; otherwise the action is financially eligible, which is not yet permitted. After the gate come three stages: awaiting creative approval (an offer or acquisition asset needs the creative lead's sign-off before scale, and the lead can reject an offer that cheapens the brand even when it clears the floor; a media shift or an email needs none); executable within the available tranche (tranche 1 at start; tranche 2 only after the passed week 4 read, signed; tranche 3 only after the passed week 8 read, signed; a tranche spent before its read holds the action until the next read passes); executed and logged (signals and source, model or tool version, rule or score, proposed exposure, approver, time, result, later correction), meeting later evidence at the week 4, week 8, 13-week and 12-month reads. Nothing executes until every required approval and authority check has passed. The offer agent's 25% code: contribution \$4.45 on net sales of \$71.25, 6.2%, below the 8% floor; rejected by the rule itself, the floor being set by the CFO, with no human decision needed. Its 12% code instead: \$16.80 on \$83.60, 20.1%, within the floor; eligible at the gate, counted against the \$120,000 discount budget, and waiting for the creative lead's sign-off; the live view can set the sign-off, labelled as assumed for the walkthrough, and the code is then permitted and executes within the available tranche. The media agent's 15% shift into a broad-reach

video buy exceeds the 10% limit: if nothing else has moved, 10% is eligible and executes within the available tranche while the remaining 5% goes to the budget owner as one request; if any part of the allowance is already used, the whole proposal is held, so no agent can split a move. The shift rule: the limit is 10% of the weekly budget; the denominator is the weekly budget, frozen for the week (the figure's stated reading; the mandate states the share, the window and the aggregation); it counts the sum of all shifts by all agents, pending commitments and pending moves included, in a rolling seven-day window. Illustration only, an even split of tranche 1 across its four weeks: $\$90,000 / 4 = \$22,500$ a week; 15% = $\$3,375$; 10% = $\$2,250$; 5% = $\$1,125$; no weekly budget in dollars is stated. The retention agent's lapsed-buyer email to customers whose last purchase was six to twelve months ago: every one is inside the twelve-month exclusion; rejected, routed to the retention programme's own budget; a customer whose last purchase was more than twelve months ago counts as a reactivation, reported separately from new customers, never as a new customer. Tranches: $\$90,000$ for weeks 1 to 4, released at start; up to $\$150,000$ for weeks 5 to 8, unlocked only by the week 4 read; up to $\$160,000$ for weeks 9 to 13, unlocked only by the week 8 read; $\$400,000$ as a contingent maximum. A passed read, signed by the budget owner on the proof lead's reading, releases the next tranche; neither read changes the mandate; if a tranche is spent before its read, spending pauses until the read and the authority stage is held; no agent can release a tranche. At a failed interim gate, new commitments and queued actions that can still be stopped are paused, even with money left in the current tranche: the walkthrough shows authority none, the authority box paused (failed gate; unspent budget does not carry authority), no executed path, the next tranche held and the plan returned to the CMO and budget owner; it resumes only under a logged human decision consistent with the mandate; existing non-cancellable commitments stay visible (none in this walkthrough); actions executed before the read stay logged as executed. This applies to tranche 1 after a failed week 4 read and to tranche 2 after a passed week 4 read and a failed week 8 read. A delayed or inconclusive read is not a failure, but it does not extend spending authority. For the week 4 read the illustrative mandate sets the continuation in advance: execution may continue for at most 7 days after the read was due, within what is left of tranche 1 and capped at $\$22,500$ (one week's even share of tranche 1); tranche 2 stays held; a control in the live view says whether the continuation is still open or used up. Outside the window or the cap, or with tranche 1 spent, new execution pauses until the read lands. No continuation is set for the week 8 read, so a delayed or inconclusive week 8 read pauses new execution until it lands or the CMO and budget owner approve a logged, capped exception. For the 12% code, signed off (assumed): week 4 not yet due, permitted within tranche 1; week 4 failed, paused with tranche 1 unspent; week 4 delayed or inconclusive, permitted within the continuation, then paused once it is used up; week 4 passed and week 8 failed, paused with tranche 2 unspent; week 4 passed and week 8 delayed or inconclusive, paused until the read lands; both passed, permitted within tranche 1, with tranches 2 and 3 also released. Unit economics: contribution per order before media = $(1 - r) \times (P \times (1 - d) - C - F) - r \times H$; net sales = $(1 - r) \times P \times (1 - d)$; P $\$100$, C $\$60$, F $\$10$, r 5%, H $\$6$; floor 8% of net sales; no discount $\$28.20$ on $\$95.00$, 29.7%. As the paper states, a returned order adds nothing before handling: goods and fulfilment are treated as recovered and only the $\$6$ handling cost remains. The calculator also offers an alternative figure assumption, fulfilment incurred on all orders: contribution = $(1 - r) \times (P \times (1 - d) - C) - F - r \times H$, giving $\$27.70$ on $\$95.00$, 29.2%; $\$16.30$ on $\$83.60$, 19.5%; $\$3.95$ on $\$71.25$, 5.5%; pass and fail at the 8% floor unchanged; an assumption of the figure, not a claim about the retailer, and the default stays the paper's. Media cost sits in the common goal, not in the floor. A floor is a setting for this retailer, not a rule for all. The live view recomputes the formula for any input and any chosen discount; the static view and print use the printed inputs and the 25% code.

Reader question: When three agents ask for things that clash, which one runs, which is refused, which goes to a person, and by what rule?

Figure 21: The next 90 days

ILLUSTRATIVE

A proposal, not an observed plan: four phase cards joined by gates, about thirteen weeks; card widths are not durations. Between the cards: the week 4 hard gate, the week 8 read, the week 12 read and, after the last card, the week 13 decision. Weeks 1 to 4, freeze the rules: pick one decision (CMO, CFO); reconcile the money (finance); map grants and write the baseline (product lead, proof lead); freeze the success, stop and expansion rules (CMO, CFO, proof lead); then, before any live spend, the control tests on a sandbox or a zero-budget grant: one stop, one refused out-of-bounds request, one

forced retry and the exposure cap (security, product lead). Week 4 gate: one signed page and all four tests pass; no live authority before it; the arrow out of the gate is the first live spend. Weeks 5 to 8, run the pilot: the live pilot at authority level 3 or 4, never higher (delivery lead); the week 6 in-flight stop rehearsal with one export test (security, product lead), apart from the week 4 tests; the holdout read (proof lead). Week 8 read: the holdout result, approved useful work per total hour and the relationship measures. Weeks 9 to 12, challenge and exit: challenge the result (proof lead or an outside reviewer); the exit drill (product lead, procurement); the full-cost bridge (finance). Week 12 read: the exit drill record and the full-cost bridge. Week 13, decide: expand, revise or stop against the frozen rule, in one meeting (CMO as accountable owner). Any stop criterion stops the pilot the same day. A failed pilot that yields usable evidence is a success of the method. Selecting a phase lists its owners, evidence and tasks with stop criteria; an optional checklist stores the reader's own ticks in the browser and verifies nothing.

Reader question: What does the CMO do in the next 90 days, in what order, who owns each step, and what has to pass before any live spend?

The Evolution of Holding Companies: Control, client value and three paths for marketing work. © 2026 Evgeny Popov, No Fluff Advisory. Research cutoff 5 October 2026; published 6 October 2026, updated 7 October 2026. The data behind every figure and count is in the downloads.