International Journal of Computer Information Systems and Industrial Management Applications
ISSN 2150-7988 Volume 18(22s) 2026 pp. 822-843
https://doi.org/10.70917/ijcisim-2026-5479
Article

Integrating Marketing Automation and Commerce Platforms into a Unified Customer Engagement Architecture

Ashok Apati

Independent Researcher, India
Email: gangarapukeerthi88@gmail.com

Abstract: Marketing automation and digital commerce are usually described as adjacent systems that need connecting. They are better understood as two systems that keep time differently. Commerce runs at the speed of a customer who is browsing, authenticating, ordering and paying, and needs the current state of identity, entitlement and account access to answer the request in front of it. Marketing automation runs on segments, journeys, campaign audiences, consent rules and scheduled engagement, and is designed to reason about populations over intervals rather than about one customer in one moment. The two also carry different data models: marketing thinks in contact keys and communication permissions, commerce in accounts, roles, pricing and ordering rights. This paper argues that most integration failure between them lives in that gap, not in the field mappings usually blamed for it. It contributes four artefacts: a source of truth ownership matrix assigning each of nine customer data domains to one owning platform; a treatment of the common customer key, the governed identity relationship across platforms, as the precondition without which marketing and commerce cannot stay aligned at all; a classification of which flows across the boundary must be real time and which may be scheduled, with the consequence of misclassifying each; and a divergence detection and reconciliation loop governed by the rule that every important data flow must have one defined source of truth, because without it teams repair the symptom in one system while the cause sits in another. The paper is explicit about being part reference architecture and method proposal, part practitioner account of one estate at a large North American consumer brand enterprise running both business to business and direct to consumer commerce. That estate supports more than 150,000 customers, more than 100,000 annual customer orders and approximately 110,000 customer registrations, with approximately 35 enterprise application programming interfaces, approximately 15 real time integrations and approximately 10 scheduled integration processes. Registration processing improved to approximately 800 milliseconds and page load performance to around one second. No campaign, conversion, email engagement or marketing lift measurement is available for this work, and none is claimed anywhere in the paper.
Keywords: marketing automation, digital commerce, customer engagement architecture, source of truth, identity resolution, consent management, master data management, eventual consistency, integration timing, data governance

1. INTRODUCTION

A customer who registers on a brand's site, receives a welcome message, logs in a week later, places an order and then updates their communication preferences has interacted with one brand and five or six systems. The brand experiences those systems as a portfolio. The customer experiences them as one relationship, and judges the brand on whether that relationship behaves consistently. That asymmetry is the standing problem in customer engagement architecture, and the customer journey literature describes it well by treating the journey, rather than the touchpoint or the channel, as the unit of analysis [1].

The version of the problem addressed here is narrower and, in the author's experience, more tractable: the reconciliation of marketing automation with digital commerce. These are the two systems most often expected to behave as one, because personalised engagement is supposed to reflect what the customer just did, and they are the

two whose internal assumptions agree least. The omnichannel literature has begun to name this as an informational problem rather than a channel problem, pointing at data integration, attribution and identity as the binding constraints [2]. Marketing automation research documents that the technology delivers value only where content, data and process are already organised to feed it, and that firms routinely underestimate that precondition [3].

Practitioners tend to meet the problem as a field mapping exercise. A contact record needs to reach the marketing platform. An order needs to reach a journey trigger. A consent flag needs to reach whatever will honour it. Framed that way the work looks like plumbing and gets estimated like plumbing. It then overruns, and the overrun is attributed to data quality.

That framing is wrong in a specific way. Marketing automation and commerce do not primarily differ in their fields. They differ in two respects that field mapping cannot reach.

They differ in timescale. Commerce often requires real time or near real time information because customers are actively browsing, logging in, placing orders, making payments or checking account details. Marketing automation works through segmentation, journeys, campaign audiences, consent rules and scheduled engagement. One system answers a question now; the other assembles an audience for later. Any integration that treats those as the same clock will be wrong in one direction or the other, and both directions have costs.

They differ in data model. Marketing systems care about contact keys and communication permissions. Commerce systems care about account access and ordering rights. Identity systems care about authentication. Customer relationship management systems care about profile and service history. Each is a defensible model for its own purpose. None is a model of the customer, and the customer is what has to be shared.

FROM THOSE TWO OBSERVATIONS THE PAPER DEVELOPS FOUR CONTRIBUTIONS.

The first is a source of truth ownership matrix. The proposition is not that one platform should own everything, nor that ownership should be negotiated case by case. It is that each of a small number of customer data domains should have exactly one owning platform, stated in advance, with every other platform treated as a consumer of that domain. Nine such domains are proposed in Section 5.

The second is the common customer key. The paper's central technical claim is that building a common customer key, an explicit identity relationship across platforms, is the precondition for everything else. Without it, marketing and commerce will always struggle to stay aligned, because there is no shared referent for the entity both are reasoning about. Segmentation, personalisation, consent enforcement and cross channel engagement all reduce to statements about a customer, and a statement about a customer is meaningless if the systems disagree about which customer it is.

The third is a timing classification scoped to the marketing and commerce boundary, separating the flows whose latency the business genuinely constrains from those that may synchronise on a schedule, and naming what goes wrong in each direction when the classification is set by habit rather than by requirement.

The fourth is a divergence detection and reconciliation loop, on the position that divergence between platforms is not an incident class to be eliminated but a steady state condition to be observed, bounded and repaired, and that the precondition for repair is a defined source of truth per flow.

Section 2 states the evidentiary basis and separates the paper's prescriptive register from its descriptive one, which matters more here than in most practitioner papers and is therefore done before any argument is made. Section 3 positions the work. Section 4 develops the two clocks and two data models problem. Sections 5 through 9 present the four contributions. Section 10 covers organisational alignment, Section 11 a sequenced adoption path, Section 12 the observed scale of the estate, Section 13 limitations and Section 14 conclusions.

2. SCOPE AND EVIDENTIARY BASIS

This section exists because the paper mixes two kinds of claim, and conflating them would misrepresent the evidence.

2.1 What kind of paper this is

The paper is part reference architecture and method proposal, and part practitioner account of a specific estate. Those are different claim types with different warrants.

The reference architecture, the ownership matrix, the timing classification, the reconciliation loop and the sequenced adoption path are proposals. They are informed by the author's practice, but they are stated as what an organisation should do, and their warrant is argument and transferability, not measurement. They are written in the present tense with explicit modal verbs, and they say "should", "must" or "is recommended".

The practitioner account describes one estate the author worked in: a large North American consumer brand enterprise operating both business to business commerce for professional B2B accounts and retail partners, and direct to consumer digital experiences. The employer is not named. Claims of this type are written in the past tense and are introduced with a phrase locating them in that estate. Where the author's own workstreams are named, they are named exactly and not elaborated: customer identity modernisation, registration redesign, and a customer key alignment workstream referred to internally as SubKey alignment.

2.2 What is measured and what is not

The estate's scale and two performance figures are available. No marketing effectiveness measurement is available.

Stated plainly, and this is the constraint that shaped the whole paper: there are no campaign conversion percentages, no email engagement figures, no marketing lift numbers, no campaign launch timings and no synchronisation error rates for this work. The author does not hold them, and declined to estimate them rather than overstate numbers without confirmed data. The paper honours that. No before and after marketing comparison appears anywhere in it, because no baseline exists to compare against. Readers looking for evidence that a unified engagement architecture improves campaign performance will not find it here, and should not infer it from the operational figures in Section 12, which are platform scale and latency figures and are evidence of neither.

One further exclusion is deliberate. The estate realised operational cost savings from cloud modernisation, a separate programme. That figure is not reported in this paper, because it is not attributable to the marketing and commerce integration work and reporting it alongside this work would invite exactly that inference.

2.3 The register map

TABLE 1 MAPS EACH SUBSTANTIVE SECTION TO ITS REGISTER AND ITS BASIS. IT IS PROVIDED SO THAT A READER CAN TELL, WITHOUT RECONSTRUCTING IT, WHICH PARTS OF THE PAPER ARE CLAIMS ABOUT THE WORLD AND WHICH ARE RECOMMENDATIONS ABOUT PRACTICE.

TABLE 1. REGISTER AND EVIDENTIARY BASIS BY SECTION

Section	Register	Basis
4. Two clocks and two data models	Analytical, with observed instances	Author's characterisation of the conflict; the platform pairs named are from the estate described
5. Reference architecture and ownership matrix	Prescriptive proposal	Proposed assignment of nine domains; the platform categories

		reflect the estate, the assignment rule is a recommendation
6. The common customer key	Prescriptive proposal, with one observed workstream	The precondition claim is the author's position; the identity and key alignment work is described as performed
7. Identity and consent as the hard integration	Descriptive, with analytical explanation	Reported as harder than expected in the estate; the explanation of why is analytical
8. Integration timing	Prescriptive classification	Flows and their required timing as stated by the author; misclassification consequences are analytical, not measured
9. Detecting and resolving divergence	Prescriptive method	Practices recommended and used in the estate; no incident counts, error rates or reconciliation volumes are available
10. Organisational alignment	Descriptive	Reported approach and the functions coordinated with
11. Sequenced adoption path	Prescriptive recommendation	Stated as a recommendation for an organisation planning this work, not as a sequence executed and evaluated
12. Observed scale and outcomes	Descriptive, measured where stated	Platform scale and two latency figures; no marketing measurement

3. RELATED WORK

The literature this paper draws on sits in five clusters that rarely cite each other. That separation is itself part of the problem, since the marketing side and the distributed data side of this question are studied by different communities.

3.1 Journeys, omnichannel information and marketing automation

Lemon and Verhoef consolidate customer experience research around the journey as the unit of analysis and argue that experience is produced across touchpoints the firm only partly controls [1]. That framing is adopted here with one addition specific to architecture: the touchpoints a firm does control are owned by different platforms, so a journey level commitment is immediately a data ownership commitment.

Cui and colleagues survey omnichannel marketing and identify the binding constraints as informational rather than channel structural, naming data integration across channels, identity linkage and measurement as the recurring difficulties [2]. This paper supplies the architectural half of that diagnosis: the mechanisms by which identity linkage and governed data flow are actually established across an enterprise platform estate.

Järvinen and Taiminen study marketing automation adoption in a business to business setting and find that automation delivers value only where the surrounding content and process capability exists, and that the technology is frequently acquired ahead of that capability [3]. The sequencing recommendation in Section 11 is a stronger version of the same finding, placing not only content but identity, consent and data ownership ahead of automation.

3.2 Consistency, timing and the limits of shared state

The distributed systems literature supplies the formal vocabulary the marketing and commerce timing problem lacks. Gilbert and Lynch's proof of Brewer's conjecture establishes that a system cannot simultaneously guarantee consistency, availability and partition tolerance [4]. Brewer's later reassessment is the more useful text for practitioners, because it reframes the result as a design space rather than a prohibition, arguing that the choice is made per operation and per invariant rather than once for a whole system [5]. That is precisely the position this paper takes about marketing and commerce flows: the classification is per flow, and there is no global answer.

Vogels' account of eventual consistency describes systems that accept a window of divergence in exchange for availability, and names the window itself as the thing to be designed [6]. Pritchett's BASE formulation, basically available, soft state, eventually consistent, is the closest existing description of how a marketing platform actually relates to a commerce platform in production: it holds soft state about the customer that is expected to converge, not state that is transactionally correct at every instant [7].

Where a business process spans systems that cannot share a transaction, Garcia Molina and Salem's saga construction remains the canonical alternative, decomposing a long lived operation into local transactions with compensating actions [8]. Onboarding across identity, consent, CRM and commerce access is a saga in this sense, and Section 7 treats it as one.

Laigner and colleagues survey data management in microservice architectures and report that practitioners struggle with exactly the questions this paper is about: which service owns which data, how to keep replicated state coherent, and how to reason about consistency once a single database has been given up [9]. Their finding that ownership is decided informally and then defended technically is what the ownership matrix in Section 5 is designed to prevent.

3.3 Entity resolution and the customer key

The formal treatment of the common customer key problem is record linkage. Fellegi and Sunter established the probabilistic foundation, formalising the decision of whether two records refer to the same entity as a likelihood ratio test with explicit error rates [10]. That framing earns its place here for one reason above all: it makes explicit that matching decisions have two error types and that reducing one increases the other. A customer key strategy that does not state which error it prefers has not been designed.

Elmagarmid, Ipeirotis and Verykios survey duplicate record detection and catalogue the field matching, blocking and decision techniques available [11]. Christen's treatment of data matching extends this to the practical apparatus, from preprocessing and blocking through comparison, classification and evaluation [12]. The relevance is direct: the duplicate customer record condition that Section 4 describes as a symptom is an entity resolution failure, and enterprises routinely attempt to fix it by reconciling systems pairwise rather than by establishing a resolved key once.

3.4 Master data, ownership and data quality

Otto's case study of master data architecture design establishes that the architecture decision is the assignment of ownership, and that the technical topology follows it rather than the reverse [13]. That is the position Section 5 takes, and it is why the matrix, not the integration diagram, is presented as the primary artefact.

Khatri and Brown decompose data governance into five decision domains, principles, quality, metadata, access and lifecycle, and stress that governance is about locating decision rights rather than about tooling [14]. Abraham, Schneider and vom Brocke's structured review confirms that governance research is dominated by structural questions of ownership and accountability, with enforcement mechanisms comparatively underdeveloped [15]. Both bear on Section 9, where the governing rule is stated as a decision right: every important data flow has one system designated authoritative, and that designation is what makes a discrepancy resolvable rather than merely visible.

Batini and colleagues survey data quality assessment and improvement methodologies, including the comparison techniques used to establish whether two representations of the same data agree [16]. The reconciliation reports in Section 9 are an instance of that method class applied at the integration boundary rather than within a single repository.

3.5 Consent, preference and the privacy constraint

Martin and Murphy review the role of data privacy in marketing and establish consent and preference as substantive constraints on what marketing may do with customer data, rather than as a compliance annex to it [17]. This paper takes the architectural position that follows: if consent constrains what may be sent, then consent state is an input to campaign execution and must be as current and as authoritative as any other input.

Empirical work shows how far implementations fall short of that. Utz and colleagues study consent notices in the field and document that design choices in the consent interface materially determine the recorded outcome [18]. Nouwens and colleagues analyse consent interfaces after the European General Data Protection Regulation and find widespread patterns that steer the recorded preference away from the user's actual one [19]. Both findings matter architecturally, not only ethically. If recorded consent is partly an artefact of interface design, then the consent record is a claim about the customer that downstream systems will enforce, and the integrity of that claim is an architectural responsibility of whichever platform is designated authoritative for it.

On the regulatory side, the California Consumer Privacy Act establishes consumer rights to know, to delete and to opt out of the sale or sharing of personal information, with definitions of the personal information those rights attach to [20]. The architectural implication is the one Section 5 encodes: a right to opt out is enforceable only if there is a single place where the opt out is recorded and a governed path by which every consuming system learns of it.

3.6 Gap addressed

Each cluster is strong internally. The marketing literature establishes that omnichannel integration is an informational problem [1], [2], [3] without specifying the architecture. The distributed data literature supplies precise vocabulary for divergence and ownership [4] through [9] but is written about services rather than commercial platform estates. The entity resolution literature solves the matching problem [10], [11], [12] largely as an analytics or data cleaning task, not as a live operational key that authentication, ordering and campaign execution all depend on. Governance and master data research locates ownership as the decisive question [13], [14], [15] without treating the marketing and commerce boundary specifically. The consent literature establishes the constraint [17], [18], [19], [20] without addressing how consent state reaches the systems that must honour it.

What is missing is an account that puts these together where practitioners actually meet them: a specific, assignable ownership scheme across the platforms a large consumer brand runs, a stated position on the customer key, a per flow timing classification, and a reconciliation discipline. That is what this paper contributes.

4. THE PROBLEM: TWO CLOCKS AND TWO DATA MODELS

4.1 Two clocks

Commerce and marketing automation are both correct about time, for their own purposes, and incompatible with each other.

Commerce is interactive and its latency budget is set by a waiting human. A customer authenticating needs the current identity state. A customer viewing an account needs current entitlement and access. A customer ordering needs current pricing and current ordering rights. If any of that is stale, the platform does not degrade gracefully; it gives a wrong answer to a customer who is present, and the wrongness is immediately visible.

Marketing automation is population oriented and its natural unit is the interval. Segments are evaluated. Journeys advance on entry criteria and waits. Campaign audiences are assembled, reviewed and sent. Consent rules

are applied at the point of send. None of that requires the current state of one customer at one instant, and much of it is actively better served by a stable snapshot, because an audience that changes while it is being sent to is harder to reason about, harder to reproduce and harder to explain to the person accountable for what went out.

The consequence is that a naive shared data layer serves neither. Design it for commerce and marketing inherits volatility it did not ask for and cannot audit. Design it for marketing and commerce inherits staleness at exactly the moments a customer is making a decision.

Applying the vocabulary from Section 3.2: the consistency choice belongs to the operation and the invariant it must preserve, not to the architecture as a whole [5]; where divergence is accepted, the window is a design parameter and must be stated [6]; and what a marketing platform holds about a customer is soft state expected to converge, not state that is transactionally correct at every instant [7]. Naming it that way is not pedantry, because it changes what teams argue about. A team that believes its marketing platform holds the truth about a customer treats every discrepancy as a defect. A team that understands it holds convergent soft state asks instead how large the window is, whether it is bounded, and whether anything irreversible happens inside it.

4.2 Two data models

The second half of the conflict is that the two sides are not modelling the same object. Marketing systems are organised around a contact and its permissions, and must answer whether this contact may be sent this kind of message through this channel. Commerce systems are organised around an account and its rights, and must answer whether this authenticated user may act on this account, at what prices, with what ordering rights. In a business to business context that is not a variant of the marketing question but a different question about a different entity, because the person and the account are distinct and the relationship between them is itself data. Identity systems are organised around a credential and an authentication outcome, and customer relationship management systems around a profile and a service history [1].

Each model is fit for purpose and none is a model of the customer. That is the structural reason field level integration cannot resolve the problem: there is no field to map, because the entities do not correspond. Establishing the correspondence is the work, and it is the subject of Section 6.

4.3 How the failure presents

In the estate described in this paper, before the integration work, the customer facing ecosystem was not fully unified. Commerce, marketing automation, customer identity, CRM, registration, consent management, enterprise resource planning, order management and payment systems each supported an important part of the customer journey without operating as one connected engagement architecture.

The following consequences of that condition are stated analytically, as the classes of friction such a disconnection can produce, and not as an itemised incident record, because no incident record is being reported: duplicate or incomplete customer records; inconsistent consent or communication preferences; slower onboarding; delays in customer data synchronisation; and additional manual effort for internal teams. For business to business customers, including professional B2B accounts and retail partners, small delays or ambiguity in account access, ordering or communication have a disproportionate effect, because those customers are transacting as part of their own work rather than as a discretionary consumer activity.

This condition is common in large consumer brands for a structural reason rather than a local one. Estates accumulate over years. Platforms are added for commerce, marketing, identity, payments and operations as each need arises, and they are not originally designed together. As digital channels grow, the business eventually requires a unified architecture so that customer data, consent, identity, commerce activity and engagement operate together. The omnichannel literature reaches the same conclusion from the demand side [2].

Figure 1 Sets Out The Two Clocks, The Flows Each Side Owns, And The Failure Modes That Occupy The Gap Between Them.

Figure 1. The commerce loop is bounded by a waiting customer; the marketing cycle is bounded by an interval and an audit requirement. The centre band lists the failure classes that arise when a flow is designed for the wrong clock. The failure classes are analytical, derived from the mechanism, and are not incidence counts.

5. REFERENCE ARCHITECTURE AND THE SOURCE OF TRUTH OWNERSHIP MATRIX

5.1 The governing principle

The recommendation in this section is deliberately narrow, and it is a recommendation rather than a report.

The goal is not to make every platform own everything. The goal is to define clearly what each platform owns, and to make information flow between systems in a secure, reliable and governed way.

That formulation rules out the two failure patterns that dominate practice. The first is the aggregating platform, where whichever system is newest or most visible gradually accretes authority over data it was not designed to hold, most often commerce becoming the source of truth for everything, or a marketing platform becoming the source of truth for customer identity. The second is negotiated ownership, where authority is decided per integration by whichever team is delivering, which produces an estate in which no domain has a single owner and every discrepancy is arguable. Section 3.4 established the literature's position that this assignment is the architecture rather than an input to it [9], [13], [14].

5.2 Nine domains, one owner each

Table 2 states the proposed assignment. It is a proposal about how a large consumer brand estate should assign ownership, using the platform categories present in the estate described in this paper. It is not a report of an audited configuration.

Table 2. Proposed source of truth ownership across nine customer data domains

#	Data domain	Owning platform	Principal consumers	Why ownership sits there
1	Customer identity and authentication	Customer identity platform, OKTA in the estate described	Commerce, CRM, registration, marketing	Authentication outcome and credential state must have one authority; a second authority is a security defect, not a data defect
2	Customer registration and onboarding	Registration platform, integrated with identity and business applications	Identity, CRM, consent, commerce	Onboarding is the event that creates the customer everywhere else; its record is the origin of the others
3	Customer account and service data	Salesforce CRM, sales and service	Commerce, marketing, service	Profile and service history accumulate through human and service interaction, which is what CRM is designed to hold
4	Consent and communication preferences	Preference centre and consent management layer, connected to Marketing Cloud	Marketing, CRM, service, commerce	Consent is legally meaningful and enforced at send time; the record must be authoritative and separable from the platform that acts on it
5	Marketing engagement and campaign execution	Salesforce Marketing Cloud	CRM, analytics	Engagement history and campaign state are produced by the executing platform and are meaningless outside it
6	Commerce activity and ordering experience	B2B or D2C commerce platform	OMS, ERP, CRM, marketing	Cart, browse and ordering interaction state belong to the interactive platform that produces them
7	Order lifecycle	OMS and ERP, by stage of the order	Commerce, CRM, service, marketing	Lifecycle authority legitimately moves during the order; the

				stage boundary must be stated rather than assumed
8	Pricing, product, invoice, credit and transaction data	ERP	Commerce, OMS, CRM	Financial and product truth carry audit and control requirements that only the system of financial record satisfies
9	Payment processing	Payment platform, CyberSource or Adyen in the estate described	Commerce, OMS, ERP	Authorisation and settlement state is owned by the processor; replicating it invites reconciliation error and scope creep

Three features of the table carry the argument.

Ownership is singular but not exclusive. Every domain has one owner and several legitimate consumers. A consumer holding a copy is expected and normal; a consumer treating its copy as authoritative is the failure.

Consent is separated from marketing execution. Domain 4 is not assigned to Marketing Cloud even though Marketing Cloud is the platform that enforces it. The separation is deliberate. Consent is legally meaningful, is captured before any marketing decision, constrains CRM and service as well as marketing, and outlives any particular engagement platform [17], [20]. Holding the authoritative record outside the executing platform means that replacing the executing platform does not put the consent record at risk, and that consent can be enforced by systems other than the one that sends campaigns.

Domain 7 admits a moving boundary honestly. Order lifecycle authority genuinely does transfer between order management and enterprise resource planning as an order progresses. Pretending otherwise produces the most common reconciliation dispute in commerce estates, where two systems both believe they hold current order state and both are partly right. The recommendation is not to eliminate the transfer but to state where it occurs, so that a query about order state has a determinate answer at every stage.

5.3 The flow that the matrix implies

The information flow should begin at customer onboarding. When a customer registers, identity is created or validated, profile data synchronises to CRM or the customer data layer, consent preferences are captured, and commerce access is enabled. From there, commerce activity connects to order management, enterprise resource planning, payments and service channels. Marketing automation then operates on approved customer identity, profile, consent and engagement data to support personalised and compliant communication.

That paragraph is prescriptive and is written as such. Its content is a recommended ordering, not a trace of a delivered pipeline. The reason the ordering matters is developed in Section 7: the sequence is not arbitrary, because each step establishes a precondition for the next, and reordering it produces specific, predictable defects.

FIGURE 2 RENDERS THE MATRIX.

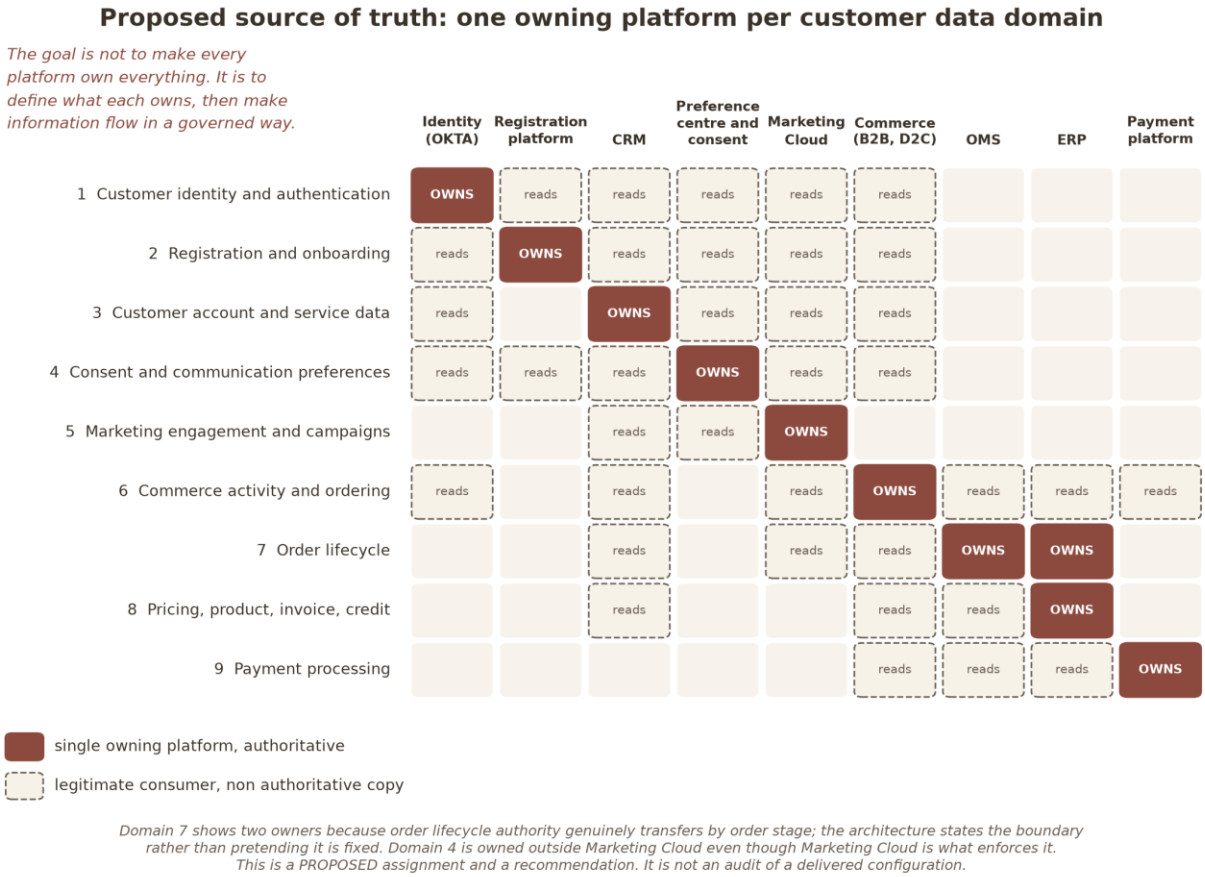

Figure 2. Each row is a data domain and each column a platform. A filled cell marks the single owning platform; an outlined cell marks a legitimate consumer holding a non authoritative copy. Domain 7 is shown with two owners because lifecycle authority transfers by order stage, which the architecture states rather than resolves.

6. THE COMMON CUSTOMER KEY

6.1 The central claim

The paper's central technical claim is stated in the author's own terms and is a position, not a measurement.

The most important step is building a common customer key, or an identity relationship across platforms. Without that, marketing and commerce will always struggle to stay aligned.

The claim is stronger than it looks. It says that the customer key is not one integration among the others but the precondition for all of them, because every cross platform statement about a customer presupposes agreement about which customer is meant. Segmentation is a statement about a set of customers. Personalisation is a statement about one. Consent enforcement is a statement about what may be done to one. Cross channel engagement is a statement that two interactions belong to the same one. If the systems disagree about identity, each of those statements is not merely inaccurate, it is undefined.

6.2 Four keys, four purposes

The platforms are not failing to agree through carelessness. Each holds the identifier its own function requires, as Section 4.2 set out and Figure 3 shows: a contact key carrying communication permissions, an account access

relationship carrying ordering rights, a credential subject, and a profile record carrying service history. The commerce key is the awkward one, because the relationship between an authenticated person and an account they may act on is many to many and time varying in a business to business setting. The identity key has the strictest integrity requirement of the four, because it gates access.

No pairwise mapping between these four resolves the problem, for the reason entity resolution research makes precise: with n systems, pairwise reconciliation requires a quadratically growing number of mappings, each with its own decision rule and error profile, and their composition has no coherent error semantics [11]. The resolved key is what collapses that.

6.3 What the key strategy has to decide

Fellegi and Sunter's formalisation frames the match decision as a likelihood ratio with two error rates, false matches and false non matches, traded against one another [10]. That trade off is what most enterprise key programmes leave unstated, and the asymmetry in a customer engagement estate is severe. A false match merges two customers, which can expose one person's account, order history or service record to another, and is close to irreversible once downstream systems have consumed the merged key. A false non match leaves one customer as two, producing duplicate records, split engagement history, contradictory consent state and the manual reconciliation effort Section 4 describes. The first is a security and privacy event; the second is an operational cost. They are not comparable, and the threshold should be set accordingly: a customer key strategy in an engagement estate should be conservative about merging and should treat unresolved duplicates as a managed backlog rather than as an error to be cleared aggressively.

The surrounding apparatus, blocking, comparison functions per attribute, a classification rule and evaluation, is well documented [11], [12]. Citing it here is not to recommend a technique but to establish that this is a solved problem class with a literature, and that enterprises frequently rebuild it informally inside integration code where its error behaviour is neither stated nor testable.

6.4 The relationship, not the identifier

One design consequence deserves emphasis. What the estate needs is not a new universal identifier imposed on every platform, which requires every platform to accept a foreign key it was not designed for and fails at the first platform that cannot.

What is needed is an explicit, governed identity relationship: a maintained record of which local identifiers, in which platforms, refer to the same customer. Each platform keeps the key its function requires. The relationship is owned as a domain in its own right, under identity per Table 2, and is queryable by every consumer. This is the structural move master data practice makes when it distinguishes the master record from the local records that reference it [13], and it is why the ownership matrix places identity first: identity owns not only authentication but the correspondence that makes the rest of the estate addressable.

In the estate described in this paper, the author's work on customer identity modernisation, registration redesign, and the customer key alignment workstream referred to internally as SubKey alignment was work of this kind. Its object was not moving a field between platforms. It was making the same customer recognisable consistently across identity, consent, CRM, commerce and marketing.

Figure 3 Shows The Resolution.

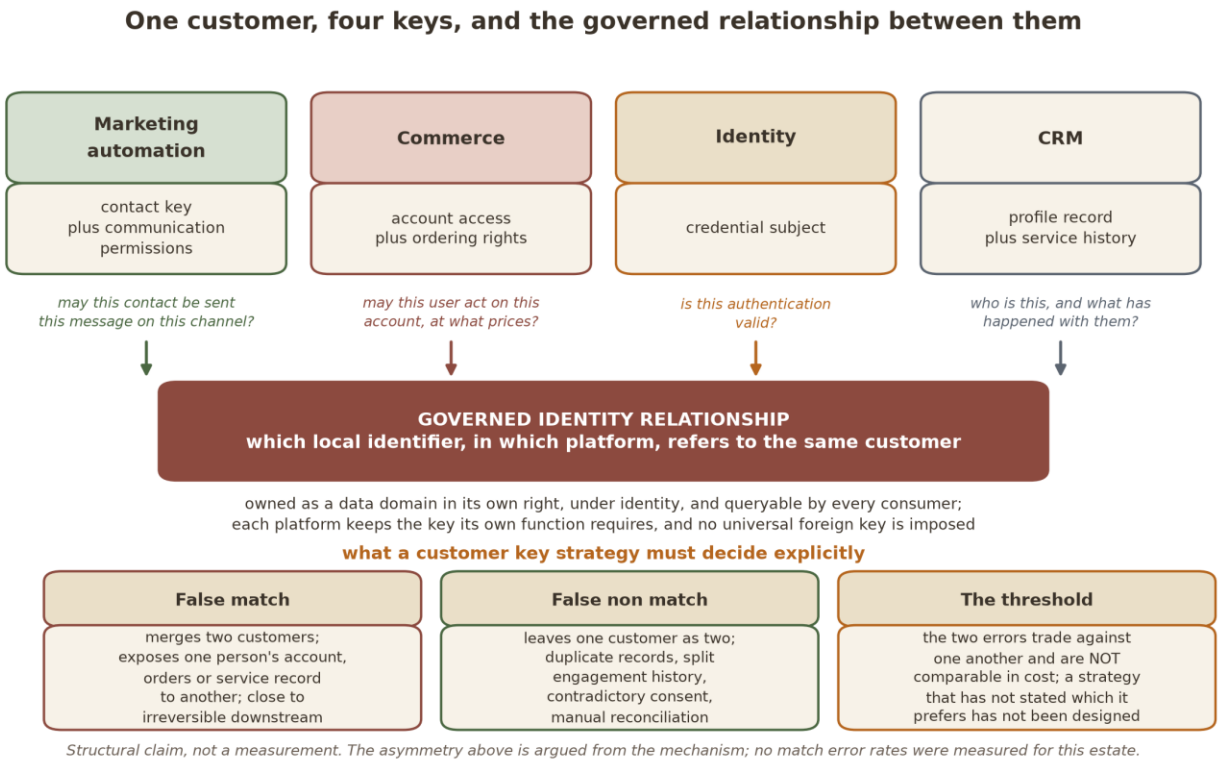

Figure 3. Each platform retains the identifier its function requires. The governed relationship records the correspondence between them and is owned as a data domain rather than embedded in integration code. The lower band names the decisions a key strategy must make explicit, including the asymmetric cost of a false match against a false non match.

7. IDENTITY AND CONSENT AS THE HARD INTEGRATION

7.1 The integration that was harder than expected

Of the integrations in the estate described here, customer identity and consent related integration proved harder than expected. This is a descriptive claim about that estate, and the author reports it as the recurring pattern in this class of work rather than as a one time surprise.

At first it presents as a straightforward data synchronisation problem. In practice it touches registration, login, profile management, CRM records, marketing automation, consent preferences, customer communication, commerce access and compliance expectations at once. The real challenge in the identity modernisation, registration redesign and customer key alignment work was not moving a field from one platform to another. It was making sure the same customer could be recognised consistently across identity, consent, CRM, commerce and marketing.

7.2 Why it is structurally hard

The explanation that follows is analytical. Four properties compound. Each system has its own identifier, per Section 6, which is why the problem is not a mapping. Each has its own data model: a contact, an account relationship, a credential subject and a profile are different entities with different cardinality, so onboarding one customer legitimately produces several records that are not copies of one another. Each has its own timing: identity state changes at authentication, consent when the customer acts on a preference, commerce access when entitlement changes, campaign audiences on a schedule, so a single onboarding lands in these systems at different moments and the

intermediate states are real states that something may observe. Each has its own business rules, since what constitutes a valid customer, a permitted communication or an entitled account differs per platform by design.

Compounded, these produce the failure modes the author names: duplicate records, incorrect communication preferences, and broken customer journeys. A business process spanning systems that cannot share a transaction has to be modelled as a sequence of local transactions with compensating actions rather than as one atomic operation [8], and the practical difficulty is not the happy path but the states between steps and the correctness of the compensations [9].

7.3 Onboarding as an ordered establishment of preconditions

Onboarding is the clearest instance, and its ordering is not arbitrary. Each step establishes a precondition for the next, which is why reordering produces specific defects rather than merely a different sequence.

Identity must be established or validated first, because every subsequent record needs something to attach to. Consent must be captured at the point the customer provides it and recorded authoritatively, because a marketing decision made before consent state is known is a decision made without a required input [17], [20]. The CRM record must exist before service or engagement history can accumulate against it. Commerce access must be provisioned against the established identity and the account relationship, not against a locally created user, or the estate acquires a second identity authority at the moment of onboarding. Marketing enrolment comes last among these, because it consumes all of the above.

Two properties of that sequence deserve stating. First, several steps can proceed concurrently and the ordering is a dependency ordering, not a serialisation requirement. Second, the sequence has compensations. If consent capture fails after identity creation, the correct behaviour is not to proceed with a default, it is to hold marketing enrolment and leave the customer in a defined, observable, incomplete state. That is the saga discipline applied to onboarding [8], and it is the difference between an estate where an incomplete onboarding is visible and one where it becomes a customer who receives communications nobody can account for.

Consent deserves a final note on integrity rather than mechanism. Empirical work shows that recorded consent is substantially shaped by the interface that captured it, and that many implementations record a preference the customer did not intend to express [18], [19]. Architecturally, this means the consent record is not a passive fact to be transported. It is a claim that downstream systems will act on, and whichever platform is designated authoritative for it, per Table 2, owns responsibility for how faithfully it represents what the customer chose.

Figure 4 shows the onboarding order.

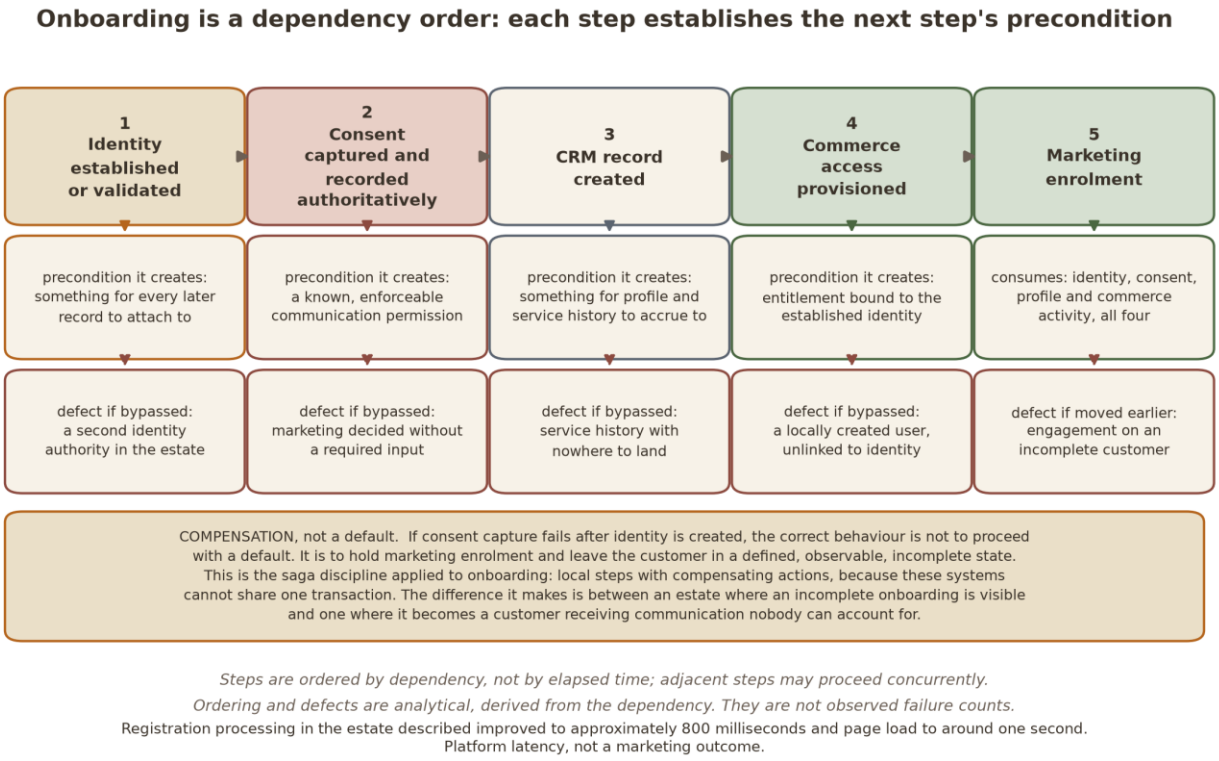

Figure 4. Steps are ordered by dependency rather than by elapsed time; adjacent steps may proceed concurrently. Each step's precondition is named, together with the defect produced if the step is bypassed or reordered. Marketing enrolment consumes every prior step, which is why it is last.

8. INTEGRATION TIMING: WHAT MUST BE REAL TIME AND WHAT MAY BE SCHEDULED

8.1 A per flow decision, scoped to this boundary

The classification in this section is prescriptive and is scoped narrowly to flows crossing the marketing, commerce, identity and consent boundary. It is not a general integration taxonomy for an enterprise estate. The governing position is that integration timing must match the business need, decided per flow and justified rather than inherited from whatever the previous integration did, per Section 4.1. Where a flow is scheduled, the divergence window it creates is a design parameter to be stated explicitly rather than discovered in production.

8.2 The classification

The flows and their required timing in Table 3 are as the author states them. The consequence column is analytical, derived from the mechanism, and is not a measured incidence.

Table 3. Timing classification for flows across the marketing and commerce boundary

Flow	Required timing	Why	Consequence if scheduled when it should be real time	Consequence if real time when it could be scheduled
Identity updates	Real time or near real time	Gates authentication and every downstream	Customer is denied or wrongly granted	Adds runtime coupling to a flow that already must be

		authorisation decision	access on stale credential state	available; limited additional risk, so real time is the correct default
Registration status	Real time or near real time	Determines whether onboarding may proceed and whether commerce access may be provisioned	Onboarding stalls or proceeds on an unconfirmed customer	Little downside; volumes are bounded by registration rate
Account access and entitlement	Real time or near real time	A waiting customer is attempting to act on an account	Wrong prices, wrong ordering rights, or visible access errors	Increases dependency on the entitlement source during checkout; must be designed for availability
Consent changes	Real time or near real time	Consent constrains what may be sent, and a stale opt out is a compliance failure, not a data lag	A customer receives communication after opting out	Negligible downside relative to the risk avoided; real time is strongly preferred
Key customer lifecycle events	Real time or near real time	Trigger service and engagement decisions whose value is time bound	Triggered engagement fires late enough to be irrelevant or wrong	Introduces load spikes on marketing platform ingestion that must be absorbed
Campaign audience refreshes	Scheduled	Audience assembly benefits from a stable, reproducible snapshot	None; scheduling is correct here	Audience mutates during send, which is unreproducible and unauditable
Selected engagement attributes	Scheduled	Aggregate behavioural attributes are population level and interval stable	None material; the attribute is not decision critical within the interval	Continuous recomputation of derived attributes for no decision benefit

8.3 Reading the table

Two asymmetries are worth naming. Consent is the flow where the two error directions are least comparable: sending after an opt out is a compliance and trust failure, while recomputing consent state more often than strictly necessary costs compute. There is no reason to trade the first against the second, which is why consent belongs on the real time side despite being conceptually a marketing attribute.

Campaign audience refresh is the one flow where scheduling is not a compromise but the correct design. A stable snapshot is what makes a send reproducible, auditable and explicable to the person accountable for it. An architecture that pushes audience membership to real time in pursuit of freshness makes campaign execution harder to reason about and gains nothing the business asked for. This is the clearest case where the marketing clock is right and commerce intuitions should not be imported.

The estate described comprised approximately 35 enterprise application programming interfaces, approximately 15 real time integrations and approximately 10 scheduled integration processes. That distribution is

reported as the shape of the estate; it is not offered as validation of the classification above, which is a recommendation and has not been evaluated against an alternative.

9. DETECTING AND RESOLVING DIVERGENCE

9.1 The governing rule

Divergence between platforms is not an eliminable condition in an estate with scheduled flows and independent business rules. It is a steady state to be bounded, observed and repaired. What makes repair possible is a single governance rule, stated by the author as the most important one in this section.

Every important data flow must have a defined source of truth. Without that, teams may fix the symptom in one system while the actual issue remains in another.

That is a decision rights statement rather than a technical one, which is what governance research identifies as the substance of governance [14], [15]. Its practical force is in triage. If consent data is inconsistent, the team must know whether the preference centre, the consent platform or Marketing Cloud is authoritative. If customer identity is mismatched, the team must know whether the issue originates in registration, in the identity platform, in CRM or in a downstream marketing system. Without that designation, both systems present evidence, the discrepancy is arguable, and the resolution taken is whichever repair is cheapest rather than whichever is correct.

The first step in practice is to locate the mismatch. In an estate of this shape the recurring boundaries are registration against identity, CRM against Marketing Cloud, commerce against enterprise resource planning, order management against commerce, and payment systems against order records. Each needs clear ownership and independent visibility, because a boundary monitored only from one side reports only the errors that side can see.

9.2 Practical methods

The practices below are recommended, and were used in the estate described here. No incident counts, error rates, mean time to detection or reconciliation volumes are available, so no claim is made about their effectiveness, only about their purpose.

- Integration logs, to identify failed application programming interface calls and data transfer errors.

- Error queues and exception reports, so that failed records are held in a defined place rather than lost.

- Retry mechanisms where the operation is idempotent and a retry is safe.

- Reconciliation reports between source and target systems, which are the only mechanism in the list that detects silent divergence, meaning records that no integration reported as failed but which do not agree. This is the data quality comparison method applied at the integration boundary [16].

- Data validation before and after major releases or migrations. Before and after rather than after only, because a discrepancy found after a change cannot be distinguished from one that pre existed it without a baseline.

- Monitoring for scheduled jobs and batch processes, since a batch that fails to run produces no error, only an increasingly stale target.

- Business user validation for customer, order, consent and account level issues, which detects the class of problem where every system is internally consistent and the aggregate is still wrong.

- Root cause review after recurring failures, which is what converts a repair into a change to the architecture.

Two of these carry more weight than the others. Reconciliation reports are the only detection mechanism for silent divergence, the failure class most likely to reach a customer precisely because nothing raised an alarm. Business user validation catches the residual class no automated comparison can, since it requires knowing what the data should mean.

9.3 The loop

These practices form a loop rather than a checklist: detect, locate, classify against the source of truth designation, repair at the authoritative system, propagate, verify, and where the failure recurs, change the flow rather than repeating the repair. The final step is the one most often skipped, and skipping it is what turns a reconciliation practice into permanent operating cost. Governance is what prevents small synchronisation issues from becoming larger customer experience problems.

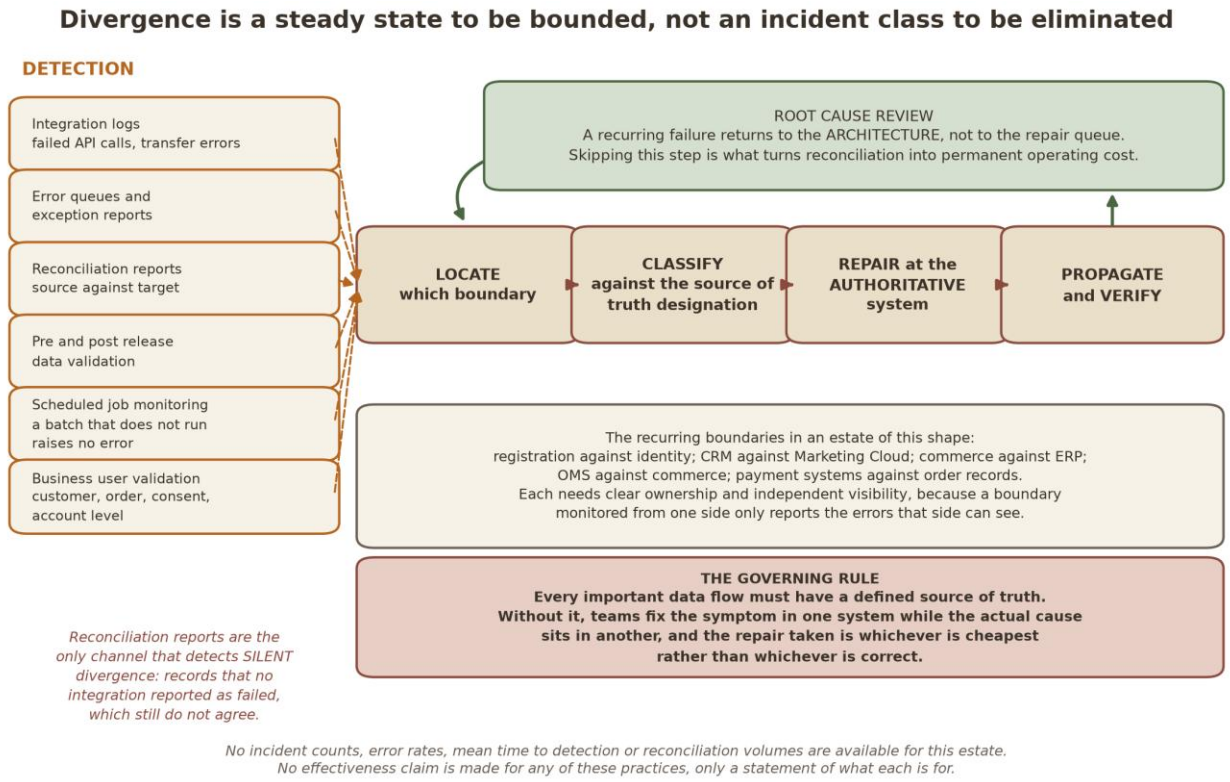

Figure 5. Detection channels feed a triage step that classifies the discrepancy against the source of truth designation. Repair is applied at the authoritative system and propagated, never at the symptom. The outer path returns recurring failures to the architecture rather than to the repair queue.

10. ORGANISATIONAL ALIGNMENT AROUND THE CUSTOMER JOURNEY

The organisational side is as consequential as the technical side. It is reported here descriptively and briefly, because what the author can attest to is an approach and a set of coordination relationships rather than a measured intervention.

Marketing and commerce teams routinely hold different goals, timelines and success measures. Commerce teams focus on ordering, customer access, platform stability, pricing, product data and transaction flow. Marketing teams focus on segmentation, personalisation, consent, engagement, campaigns and customer communication. Both sets are legitimate, and neither is a view of the journey.

The approach taken was to create alignment around the customer journey rather than around individual systems: instead of asking only what each system needed, asking what the customer and the end to end business process required. Customer onboarding is not only a registration issue, since it affects identity, commerce access, CRM records, consent and marketing communication. Customer engagement is not only a marketing automation issue, since it depends on accurate customer data, consent status, account activity and commerce behaviour. Framing the work at journey level is what makes the ownership matrix in Section 5 negotiable, because a matrix presented as a systems

question invites each team to defend its platform, whereas the same matrix presented as a journey question asks each team where the customer is best served.

This was handled by working closely with product owners, marketing stakeholders, commerce teams, integration teams, quality assurance, infrastructure, security and external vendors. The recurring task was translating business priorities into technical requirements and, in the other direction, explaining technical dependencies in terms business teams could act on. A large part of it was setting expectations explicitly: what should happen in real time, what can be scheduled, which system owns which data, what risks exist, and how changes should be tested before launch. Those are the four artefacts of this paper, restated as a conversation, and the observation worth generalising is that they function organisationally as much as technically. A stated ownership matrix and a stated timing classification remove a category of recurring disagreement, not because they are optimal, but because they are decided and written down.

11. A SEQUENCED ADOPTION PATH

This section is a recommendation for an organisation planning this work. It is prescriptive throughout, and it is not a report of a sequence that was executed and evaluated. A phased approach is recommended over attempting to integrate everything at once.

1. Map the customer journey. How a customer registers, logs in, accesses commerce, orders, receives communication, interacts with service and updates preferences [1], [2].

2. Define the source of truth for identity, consent, customer profile, orders, product data, pricing, payments and marketing engagement, per Section 5 [13], [14].

3. Clean up customer identity, per Section 6 [10], [11], [12].

4. Modernise consent and preference management before increasing marketing automation [17], [19], [20].

5. Integrate commerce with CRM and customer data, so service and engagement operate on what the customer actually did.

6. Connect marketing automation using governed data flows, consuming approved, accurate and consented data only.

7. Integrate order management, enterprise resource planning and payments with strong testing, monitoring and rollback planning.

8. Add automation and artificial intelligence or assistant capabilities last, once identity, consent, data ownership and integration are stable.

Two positions in that ordering do the work. The first is that step 1 precedes step 2: the ownership decisions are only answerable against a journey, because asked against a systems inventory they have no determinate answer. The second is the position of marketing automation at step 6 of 8, after identity, consent and data ownership, which is the opposite of the order in which these programmes are usually funded [3]. Increasing send volume across an unreliable consent record scales a compliance exposure rather than an engagement capability.

Step 8 carries the strongest claim, and it is structural rather than cautionary. An automation or assistant layer over an estate whose customer identity is unresolved, whose consent record is contested and whose data ownership is undefined does not have an automation problem. It has an identity and governance problem that will present as an automation problem the first time the layer is pointed at the estate, because it will confidently act on whichever of several conflicting customer representations it happened to read. The layer inherits the estate's ambiguity and adds speed to it. Nothing in the eventual consistency literature suggests that reasoning over convergent soft state is unsound [6], [7]; the point is narrower, that a system acting autonomously on customer data needs to know which copy is authoritative, and that is exactly what an undefined ownership scheme fails to supply.

The overall recommendation reduces to one line. Do not start with campaigns or tools. Start with identity, data ownership, consent and integration architecture.

12. OBSERVED SCALE AND OPERATIONAL OUTCOMES

This section reports what is measured. It is short, and its brevity is the honest result of the constraint stated in Section 2.2.

The architecture described supports a large North American enterprise environment serving more than 150,000 customers and more than 100,000 annual customer orders. Customer registration modernisation supported approximately 110,000 customer registrations. The integrated ecosystem comprised approximately 35 enterprise application programming interfaces, approximately 15 real time integrations and approximately 10 scheduled integration processes across commerce, CRM, enterprise resource planning, identity, payments, marketing and related platforms. The connected estate spans business to business commerce, direct to consumer commerce, Salesforce CRM, Salesforce Marketing Cloud, customer registration, customer identity on OKTA, consent management, enterprise resource planning, order management and order to cash workflows, payment platforms including CyberSource and Adyen, MuleSoft and enterprise integration services, and Azure cloud and data services.

Registration processing time improved to approximately 800 milliseconds, and page load performance improved to around one second. These are the only latency figures available. They are reported because onboarding is the entry point to both commerce and marketing engagement, so registration performance is a legitimate part of an engagement architecture's operational profile. They are platform scale and latency figures, not evidence of marketing effectiveness, and no marketing measurement exists for this work.

The improvements the author attributes to this work that cannot be quantified are stated as such: reduced manual effort, better data consistency, stronger consent alignment, improved platform reliability, better production readiness, and improved capacity to support subsequent customer experience work. For marketing and commerce specifically, the value claimed is foundational rather than performance related: cleaner customer identity, better consent handling, stronger integration between customer facing and backend systems, and improved readiness for personalised engagement. Readiness is not performance, and the paper does not treat it as such.

13. LIMITATIONS

Single enterprise, single practitioner. The account is drawn from one estate, described by one architect who worked in it. There is no second organisation, no control condition, no comparison estate and no counterfactual.

No controlled evaluation. None of the four artefacts, the ownership matrix, the customer key position, the timing classification or the reconciliation loop, has been evaluated against an alternative. Their warrant is argument and mechanism, not comparative result.

No campaign, conversion or marketing performance measurement is available, and none is claimed. This is the most consequential limitation and it constrains what the paper can be. A reader may not conclude from anything here that a unified customer engagement architecture improves campaign performance, conversion, email engagement or marketing return. The measurements that would support such a conclusion do not exist for this work, and the author declined to estimate them rather than state unconfirmed numbers. Anyone seeking that evidence should treat this paper as describing the precondition, not the effect.

The operational figures are platform scale and latency figures. More than 150,000 customers, more than 100,000 annual orders, approximately 110,000 registrations, approximately 800 milliseconds registration processing, around one second page load, and the integration counts, describe the environment the architecture supports and the speed of one interaction. They are not evidence of marketing effectiveness, and the registration latency figures in particular measure a technical interaction rather than a customer outcome. Nor is there any before and after baseline for anything marketing related: Section 4.3's account of the before state is a characterisation of a condition, not a measurement, and it cannot support an improvement claim in either direction.

The reference architecture is derived from one estate and its transferability is argued, not demonstrated. The nine domains and their assignment reflect the platform categories present in a large consumer brand estate running both business to business and direct to consumer commerce. An organisation with a materially different estate, for example one without a separate order management system, or one where consent is held inside the marketing platform for reasons of its own, will need to re derive the matrix rather than adopt it. The claim is that the assignment discipline transfers, not that this particular assignment does.

Prescriptive content is not validated by the descriptive content. Sections 5, 8, 9 and 11 are recommendations. That the estate in Section 12 operates at the stated scale is not evidence that those recommendations produced it, and the paper does not present it as such.

The estate is generalised, and consent is treated architecturally rather than legally. Products are named because an architecture paper without platform categories is not actionable, but the employer is not named, which limits independent verification of every descriptive claim here. The paper takes positions on where consent state should be owned and how promptly it should propagate; it offers no compliance assessment against any specific regulatory regime, and the regulatory material is cited for the constraint it establishes rather than for legal analysis [17], [18], [19], [20].

14. CONCLUSION

Marketing automation and commerce are usually integrated as though the problem were correspondence between fields. The argument here is that the problem is correspondence between clocks and between entities, and that field level work cannot reach either.

Commerce needs current state because a customer is waiting. Marketing automation needs stable, reproducible, consented population state because it is assembling an audience and will have to account for what it sent. Both are correct. The reconciliation is not a shared database. It is four decisions made explicitly: which platform owns each data domain, what the customer key relationship is, which flows the business genuinely constrains in time, and how divergence is detected and repaired at the authoritative system rather than at the symptom.

The nine domain ownership matrix is a concrete starting point rather than a finished answer, with two features worth carrying forward independently of the rest: consent held authoritatively outside the platform that executes on it, and order lifecycle authority stated as transferring by stage rather than pretended to be fixed. The customer key is the paper's central claim, and it is a claim about precondition rather than about difficulty. Without a governed identity relationship across platforms, statements about a customer are not inaccurate, they are undefined, and no amount of downstream integration work repairs that.

Future work follows directly from what is missing. The timing classification in Section 8 is testable: a flow moved between real time and scheduled, with the divergence window and the resulting business error measured on both sides, would convert a recommendation into a result. The customer key position is testable through match error rates and their downstream cost, with the asymmetry between false match and false non match measured rather than argued. And the effectiveness question this paper cannot answer, whether a governed engagement architecture improves marketing performance, requires exactly the instrumentation the author declined to substitute assumptions for: a campaign measurement baseline established before the architecture changes, which is a discipline to be adopted at the start of such a programme rather than reconstructed at its end.

The closing practical thesis is the author's own and is unchanged by anything in the analysis. Do not start with campaigns or tools. Start with identity, data ownership, consent and integration architecture.

REFERENCES

1. Lemon, K. N., & Verhoef, P. C. (2016). Understanding customer experience throughout the customer journey. Journal of Marketing, 80(6), 69-96. https://doi.org/10.1509/jm.15.0420

2. Cui, T. H., Ghose, A., Halaburda, H., Iyengar, R., Pauwels, K., Sriram, S., Tucker, C., & Venkataraman, S. (2021). Informational challenges in omnichannel marketing: Remedies and future research. Journal of Marketing, 85(1), 103-120. https://doi.org/10.1177/0022242920968810
3. Järvinen, J., & Taiminen, H. (2016). Harnessing marketing automation for B2B content marketing. Industrial Marketing Management, 54, 164-175. https://doi.org/10.1016/j.indmarman.2015.07.002
4. Gilbert, S., & Lynch, N. (2002). Brewer's conjecture and the feasibility of consistent, available, partition-tolerant web services. ACM SIGACT News, 33(2), 51-59. https://doi.org/10.1145/564585.564601
5. Brewer, E. (2012). CAP twelve years later: How the "rules" have changed. Computer, 45(2), 23-29. https://doi.org/10.1109/MC.2012.37
6. Vogels, W. (2009). Eventually consistent. Communications of the ACM, 52(1), 40-44. https://doi.org/10.1145/1435417.1435432
7. Pritchett, D. (2008). BASE: An ACID alternative. Queue, 6(3), 48-55. https://doi.org/10.1145/1394127.1394128
8. Garcia-Molina, H., & Salem, K. (1987). Sagas. In Proceedings of the 1987 ACM SIGMOD International Conference on Management of Data (pp. 249-259). ACM Press. https://doi.org/10.1145/38713.38742
9. Laigner, R., Zhou, Y., Salles, M. A. V., Liu, Y., & Kalinowski, M. (2021). Data management in microservices: State of the practice, challenges, and research directions. Proceedings of the VLDB Endowment, 14(13), 3348-3361. https://doi.org/10.14778/3484224.3484232
10. Fellegi, I. P., & Sunter, A. B. (1969). A theory for record linkage. Journal of the American Statistical Association, 64(328), 1183-1210. https://doi.org/10.1080/01621459.1969.10501049
11. Elmagarmid, A. K., Ipeirotis, P. G., & Verykios, V. S. (2007). Duplicate record detection: A survey. IEEE Transactions on Knowledge and Data Engineering, 19(1), 1-16. https://doi.org/10.1109/TKDE.2007.250581
12. Christen, P. (2012). Data matching: Concepts and techniques for record linkage, entity resolution, and duplicate detection. Springer Berlin Heidelberg. https://doi.org/10.1007/978-3-642-31164-2
13. Otto, B. (2012). How to design the master data architecture: Findings from a case study at Bosch. International Journal of Information Management, 32(4), 337-346. https://doi.org/10.1016/j.ijinfomgt.2011.11.018
14. Khatri, V., & Brown, C. V. (2010). Designing data governance. Communications of the ACM, 53(1), 148-152. https://doi.org/10.1145/1629175.1629210
15. Abraham, R., Schneider, J., & vom Brocke, J. (2019). Data governance: A conceptual framework, structured review, and research agenda. International Journal of Information Management, 49, 424-438. https://doi.org/10.1016/j.ijinfomgt.2019.07.008
16. Batini, C., Cappiello, C., Francalanci, C., & Maurino, A. (2009). Methodologies for data quality assessment and improvement. ACM Computing Surveys, 41(3), Article 16, 1-52. https://doi.org/10.1145/1541880.1541883
17. Martin, K. D., & Murphy, P. E. (2017). The role of data privacy in marketing. Journal of the Academy of Marketing Science, 45(2), 135-155. https://doi.org/10.1007/s11747-016-0495-4
18. Utz, C., Degeling, M., Fahl, S., Schaub, F., & Holz, T. (2019). (Un)informed consent: Studying GDPR consent notices in the field. In Proceedings of the 2019 ACM SIGSAC Conference on Computer and Communications Security (pp. 973-990). ACM. https://doi.org/10.1145/3319535.3354212
19. Nouwens, M., Liccardi, I., Veale, M., Karger, D., & Kagal, L. (2020). Dark patterns after the GDPR: Scraping consent pop-ups and demonstrating their influence. In Proceedings of the 2020 CHI Conference on Human Factors in Computing Systems (pp. 1-13). ACM. https://doi.org/10.1145/3313831.3376321
20. California Consumer Privacy Act, California Civil Code, Division 3, Part 4, Title 1.81.5, sections 1798.100 to 1798.199.100. California Legislative Information. Retrieved 2026-08-10, from https://leginfo.legislature.ca.gov/faces/codes_displayText.xhtml?division=3.&part=4.&lawCode=CIV&title=1.81.5