LOCUS/ANALYTICA
DATA SCIENCE & ENGINEERING

WHITE PAPER

WHITE PAPER / REVENUE ANALYTICS · DATA ENGINEERING

Credit the Account, Not the Lead

A warehouse-native pattern for account-based campaign influence on Salesforce, BigQuery and Power BI

AUDIENCE CMOs, CROs, CIOs, RevOps and data leaders

AUTHOR Adrián del Valle Rodríguez · Founder & Data Lead, Locus Analytica

PUBLISHED November 2026

6–10	4	1	7	0	1
buyers per decision[1]	CRM objects joined	winning campaign	governed rules	PII in the model	number both teams trust

[1] Gartner, as reported in reference [2].

Contents

Executive Summary 03

1. Introduction and Context 05

2. Problem Statement 05

3. Technology Overview 07

4. Reference Architecture and Methodology 08

5. Strategic Analysis 14

6. Case Study: One Number for Marketing and Sales 16

7. Implementation Roadmap 18

8. Business Impact and KPIs 19

9. Recommendations and Conclusions 20

10. References 21

ABOUT THIS PAPER

The organization in the case study is anonymized. Schemas, field names and code are reference implementations prepared by Locus Analytica. Bracketed numbers refer to the sources listed in Section 10.

Executive Summary

B2B purchases are made by groups, but most campaign attribution is still built around individuals. Gartner reports that the buying group for a complex B2B solution typically involves six to ten decision makers [2], and Forrester found that more than 94% of B2B buying decisions are made by a group of three or more [3]. When the person who attends an event is not the person a sales rep attaches to the deal, contact-level attribution loses the connection, and marketing's influence on pipeline goes unrecorded.

Central thesis. Account-based campaign influence should be decided **once, upstream, in the data warehouse**, as a governed, time-bounded and deterministic bridge between opportunities and campaigns. Its rules should be owned jointly by marketing and sales leadership, its output kept at one row per opportunity, and every reported number should drill back to the opportunities behind it. CRM-native influence features and BI-layer calculations each solve part of the problem; neither, on its own, produces a number both the CMO and the CRO will defend.

Key findings

- **The gap is structural, not a data-entry problem.** In Salesforce, a campaign member is either a contact or a lead [7], and native campaign influence connects campaigns to opportunities through opportunity contact roles [8, 9]. Buyers who engage but are never added as contact roles disappear from the analysis.
- **Account qualification closes the gap.** Reaching the account through the contact's own account link, and then testing a campaign-level influence window against the opportunity's creation date, attributes influence to the account as a whole while keeping the logic auditable.
- **Rule-based models still have a role in B2B.** Google retired four rule-based models from Google Ads and GA4 in 2023 in favor of data-driven attribution [5]. That suits high-volume digital conversions; low-volume, high-value B2B deals with offline touchpoints still benefit from explicit, agreed rules.
- **The modern stack is sufficient.** Salesforce data can be landed in BigQuery with Google's own connector [10], the rules expressed in standard SQL with window functions [11], and the result served through a star-schema Power BI model [13]. No proprietary attribution engine is required.

Recommendations

#	Recommendation	Owner
1	Write the attribution rules down and have marketing and sales leadership sign them off before any code is written.	CMO · CRO
2	Build one attribution bridge in the warehouse at opportunity grain; let every report read from it.	Data engineering
3	Keep the attribution date (opportunity created) separate from the reporting date (pipeline recognized).	RevOps · Data
4	Govern credit splits as versioned configuration, applied in measures, never hard-coded in visuals.	RevOps
5	Validate against the CRM with the business owners before adoption, and keep automated quality checks afterwards.	RevOps · Data
6	Make every aggregate drill to the opportunity list it summarizes.	BI team

1. Introduction and Context

Gartner describes today's B2B buying journey as nonlinear, closer to a set of distinct buying jobs that a group completes in parallel than to a funnel [1]. Forrester has urged B2B organizations to move from lead-centric demand management to insight about accounts, opportunities and buying groups, and connects account-based marketing measurement directly to opportunity measurement [4].

Revenue leaders have adapted their go-to-market motions to this reality: field events, executive dinners, account-based advertising and partner programs are planned around target accounts rather than individual leads. Measurement has lagged. In many organizations, the systems that record who attended an event and the systems that record which deals were created were designed around individual people, and the account sits between them as an implicit relationship rather than an explicit join.

At the same time, the attribution tooling landscape has shifted. In April 2023 Google announced it would remove first-click, linear, time-decay and position-based models from Google Ads and Google Analytics 4, keeping data-driven and last-click models, and noted that fewer than 3% of Google Ads web conversions used the retired models [5]. For digital advertising at scale, algorithmic attribution is now the default. For B2B pipeline, where deal counts are small, cycles are long and many touches happen offline, organizations still need a transparent rule set that sales and marketing can inspect, debate and agree on.

Scope of this paper

This paper addresses **account-based campaign influence**: determining, for each opportunity, which campaign was influencing the account when the opportunity was created, and reporting the resulting pipeline by campaign and period. It presents a reference architecture on Salesforce, Google BigQuery and Microsoft Power BI, the analytical trade-offs behind it, a case study of a B2B software company that adopted the pattern, and a roadmap for organizations considering the same path. Multi-touch weighting and incrementality testing are discussed as extensions, not as the core.

2. Problem Statement

2.1 The contact-role gap

Salesforce models campaign participation in the `CampaignMember` object, and each record contains either a `ContactId` or a `LeadId`, never both [7]. The native Customizable Campaign Influence feature creates influence records when a campaign member is also a contact role on an opportunity; the `CampaignInfluence` object itself carries the opportunity contact role it was derived from [8], and practitioner guidance is consistent that without contact roles no influence is tracked [9]. The consequence is shown below.

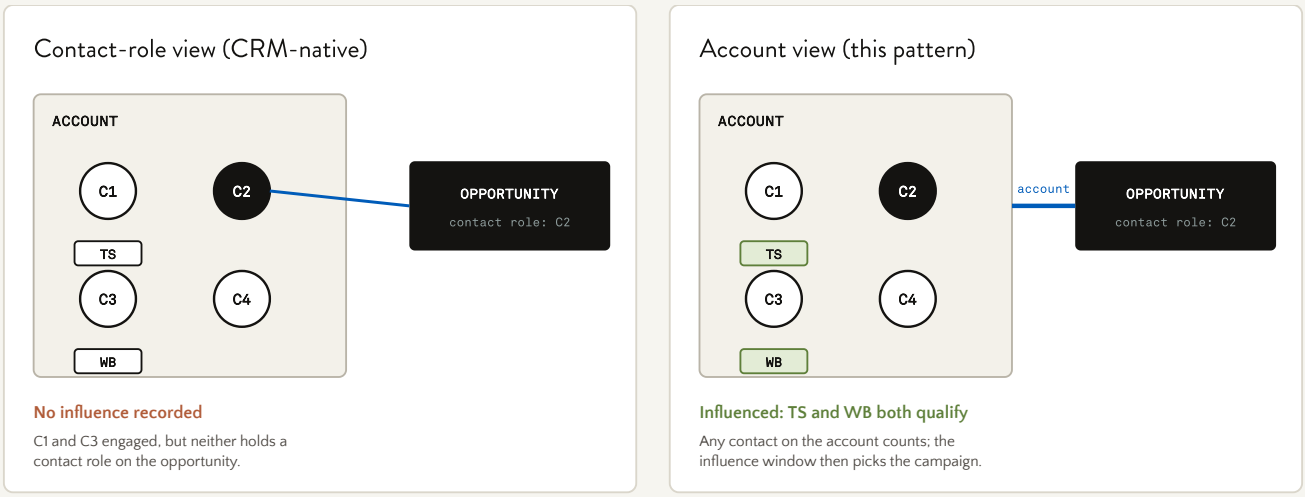

Figure 1 · Two engaged contacts, one opportunity, two different answers.

A related nuance is easy to misread. Salesforce offers an optional *Accounts as Campaign Members* setting that adds an `AccountId` to `CampaignMember` [6]. That field identifies an account enrolled in a campaign as a member in its own right; it does not record the employer of a contact who attended. For contact members, the account must still be reached through `Contact.AccountId`.

2.2 Why the usual fixes fall short

Approach	What it does well	Where it breaks for account-based influence
Enforce contact roles in the sales process	Improves CRM hygiene and native reports	Depends on rep behavior; buyers who never meet sales still drop out
CRM-native influence models (first/last touch, even distribution)	Native, no data pipeline	Contact-role dependency; limited control over account logic and time windows
Spreadsheet reconciliation	Fast to start, flexible	Not repeatable or auditable; every refresh is manual
Calculation in the BI layer	Close to the report	Logic duplicated per report; hard to test; the attribution decision is remade on every query
Dedicated attribution platform	Rich multi-touch features	Cost, another system of record, and rules that may not match the agreed business definition

There is also a quieter failure mode: **SQL that looks right and silently drops records**. A common first implementation chains opportunity → contact → campaign member → campaign with outer joins and then filters on campaign dates in the WHERE clause. Filtering the right-hand side of a LEFT JOIN in WHERE removes the unmatched rows and turns the join into an inner join, so opportunities vanish without an error. The pattern in Section 4 avoids this by construction.

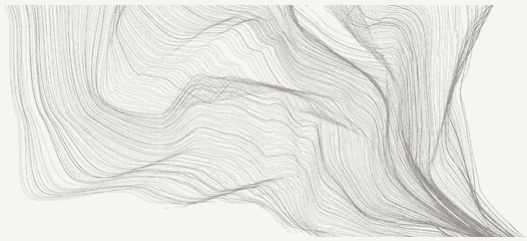

3. Technology Overview

3.1 Attribution models in brief

Family	How credit is assigned	Typical fit
Single-touch (first or last)	100% to one touch selected by a rule	Clear ownership questions; sourced vs influenced
Rule-based multi-touch (linear, U-shaped, time decay)	Fixed weights across touches	Long journeys where every stage matters
Data-driven / algorithmic	Weights learned from conversion paths	High-volume digital conversions [5]
Qualified first-touch, time-window (this paper)	100% to the earliest campaign whose influence window contains the opportunity's creation, qualified at account level	B2B influence on pipeline with offline programs

3.2 Key terms

Term	Definition used in this paper
Influence window	A start and end date stored on each campaign, defining when the campaign is considered to influence its accounts.
Qualifying account	An account with at least one contact who is a member of a campaign.
Attribution bridge	A warehouse table with one row per influenced opportunity and its credited campaign.
Attribution date	The date tested against influence windows: the opportunity's creation date.
Reporting date	The date an opportunity counts as pipeline (for example, entering a discovery stage).
Credit share	The fraction of an influenced opportunity's value credited to marketing, agreed by leadership.

3.3 The platform components

- **Salesforce Sales Cloud** is the system of record for accounts, contacts, campaigns, campaign members and opportunities.
- **BigQuery Data Transfer Service for Salesforce** schedules recurring loads from Sales Cloud using the Salesforce Bulk API, with full or incremental transfers [10]. Managed ELT vendors are a valid alternative.
- **BigQuery** hosts staging views, the attribution bridge and quality checks. GoogleSQL supports window functions and the QUALIFY clause for filtering on their results [11].
- **Power BI** serves a semantic model in star-schema form, which Microsoft recommends for performance and usability [13], with drillthrough pages for opportunity-level detail [15].

4. Reference Architecture and Methodology

The architecture separates four concerns: **landing** CRM data unchanged, **standardizing** it, **deciding** attribution once, and **serving** the decision to people. Each concern lives in its own layer so that a change in one (a new report, a new rule) does not ripple through the others.

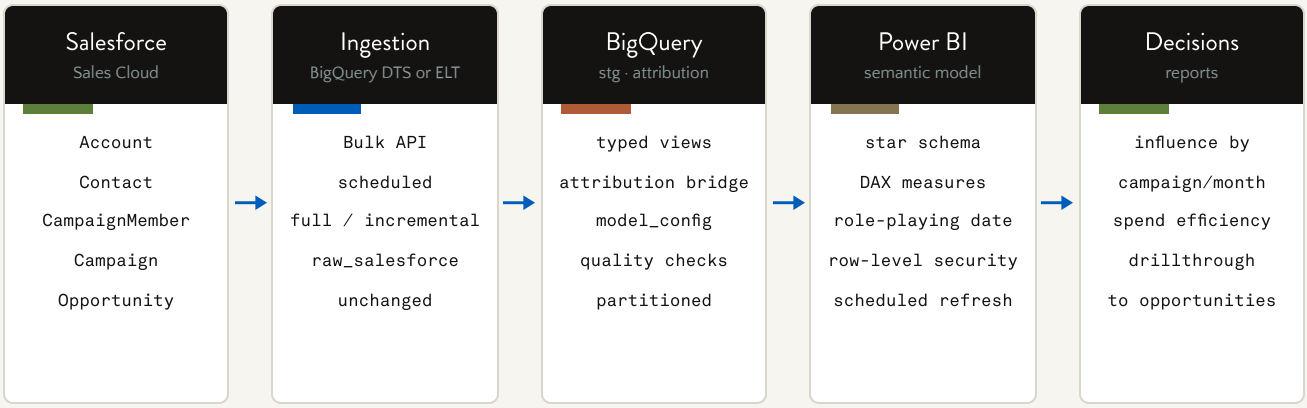

Figure 2 · Reference architecture.

4.1 Source data model

Five Salesforce objects are involved. The influence window and the pipeline date are organization-specific custom fields; their names below are illustrative.

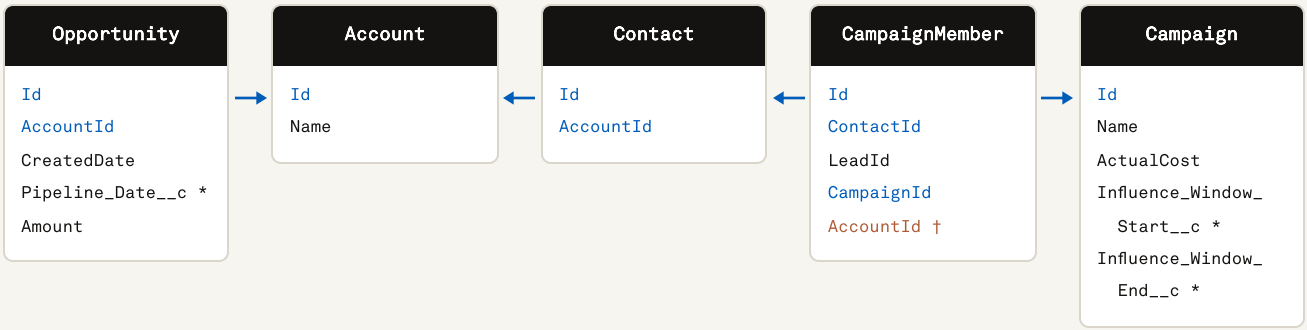

* illustrative custom fields † present only when 'Accounts as Campaign Members' is enabled; the member is then the account itself

Figure 3 · Objects and keys used by the pattern.

4.2 The seven rules

The model is defined in business language first. Each rule maps to one construct in the code that follows.

#	Rule	Implemented by
1	**Account-based.** Any contact on the account can qualify it.	Join through `Contact.AccountId`
2	**Contacts only.** Lead memberships never count.	Inner join on `contact_id`; leads have none
3	**Time window.** The campaign's window must contain the opportunity's created date.	`BETWEEN` inside the join condition
4	**First touch.** Earliest window start wins; ties are broken deterministically.	`QUALIFY ROW_NUMBER() … = 1`
5	**One value field.** Opportunities without the agreed amount are excluded.	`WHERE o.amount IS NOT NULL`
6	**Reporting date.** Pipeline is reported when it is recognized.	Active date relationship in Power BI
7	**Credit share.** Marketing receives an agreed share of influenced value.	`model_config` + DAX measure

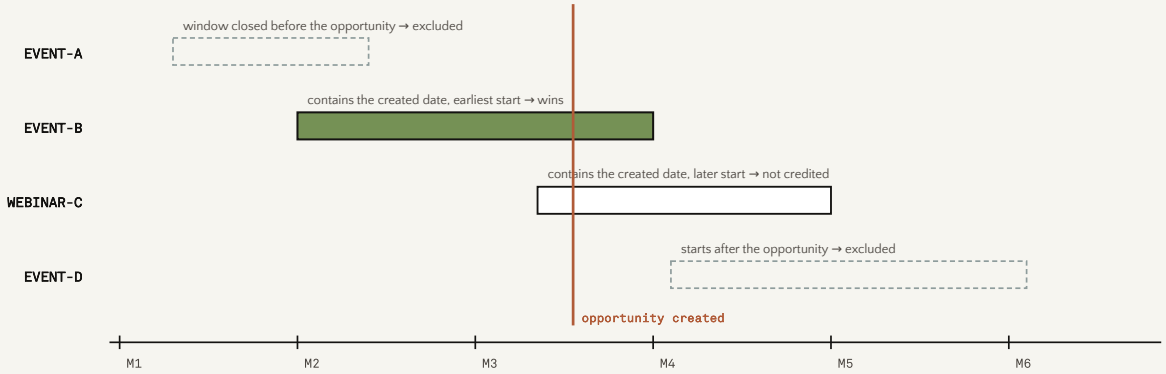

Figure 4 · Rules 3 and 4 in practice.

4.3 Staging layer

Staging views rename and type the raw columns and drop deleted records. Only identifiers, dates and amounts move downstream, so the attribution model never handles contact names or e-mail addresses.

01_staging.sql · BigQuery lines 1–38

```sql
-- Staging layer: typed, renamed views over the raw Salesforce landing tables.
-- Only identifiers and dates travel downstream; no contact PII is needed.
CREATE OR REPLACE VIEW stg.campaign_member AS
SELECT
  Id          AS campaign_member_id,
  CampaignId  AS campaign_id,
  ContactId   AS contact_id,   -- NULL when the member is a lead
  LeadId      AS lead_id
FROM raw_salesforce.CampaignMember
WHERE IsDeleted = FALSE;

CREATE OR REPLACE VIEW stg.campaign AS
SELECT
  Id                              AS campaign_id,
  Name                            AS campaign_name,
  Type                            AS campaign_type,
  ActualCost                      AS program_spend,
  DATE(Influence_Window_Start__c) AS influence_start_date,
  DATE(Influence_Window_End__c)   AS influence_end_date
FROM raw_salesforce.Campaign
```

```sql
WHERE IsDeleted = FALSE;

CREATE OR REPLACE VIEW stg.contact AS
SELECT Id AS contact_id, AccountId AS account_id
FROM raw_salesforce.Contact
WHERE IsDeleted = FALSE;

CREATE OR REPLACE VIEW stg.opportunity AS
SELECT
  Id                     AS opportunity_id,
  AccountId              AS account_id,
  DATE(CreatedDate)      AS created_date,
  DATE(Pipeline_Date__c) AS pipeline_date,    -- date the deal counts as pipeline
  StageName              AS stage_name,
  Type                   AS opportunity_type,
  Amount                 AS amount
FROM raw_salesforce.Opportunity
WHERE IsDeleted = FALSE;
```

4.4 The attribution bridge

The bridge is a table rebuilt on a schedule. Two design choices deserve explanation. First, qualification runs **from campaign membership to accounts**, producing (account, campaign, window) combinations, and only then meets opportunities; the window is a join condition, so no outer-join filter can drop rows. Second, `ROW_NUMBER()` with a secondary sort on `campaign_id` guarantees exactly one winner even when two windows start on the same day; `RANK()` would keep both and double-count the opportunity.

02_bridge.sql · BigQuery lines 1–49

```sql
-- Attribution bridge: one row per influenced opportunity, rebuilt on a schedule.
CREATE OR REPLACE TABLE attribution.opportunity_campaign_influence
PARTITION BY DATE_TRUNC(opportunity_created_date, MONTH)
CLUSTER BY account_id, campaign_id
AS
WITH account_campaign_windows AS (
  -- Rule 1-2: an account qualifies through ANY of its contacts (never leads).
  SELECT DISTINCT
    ct.account_id,
    cm.campaign_id,
    cp.influence_start_date,
    cp.influence_end_date
  FROM stg.campaign_member AS cm
  JOIN stg.contact AS ct ON ct.contact_id = cm.contact_id
  JOIN stg.campaign AS cp ON cp.campaign_id = cm.campaign_id
  WHERE cp.influence_start_date IS NOT NULL
    AND cp.influence_end_date   IS NOT NULL
),

first_touch AS (
  -- Rule 3: the window is part of the join, not a filter after it.
  -- Rule 4: earliest window start wins; campaign_id breaks exact ties.
  SELECT
    o.opportunity_id,
    o.account_id,
    o.created_date AS opportunity_created_date,
    o.pipeline_date,
    o.amount,
    w.campaign_id,
    w.influence_start_date,
    w.influence_end_date
```

```sql
  FROM stg.opportunity AS o
  JOIN account_campaign_windows AS w
    ON w.account_id = o.account_id
    AND o.created_date BETWEEN w.influence_start_date AND w.influence_end_date
  WHERE o.amount IS NOT NULL
  QUALIFY ROW_NUMBER() OVER (
    PARTITION BY o.opportunity_id
    ORDER BY w.influence_start_date ASC, w.campaign_id ASC
  ) = 1
)

SELECT
  ft.*,
  cfg.model_version,
  cfg.marketing_credit_share        -- Rule 7: governed, not hard-coded in reports
FROM first_touch AS ft
CROSS JOIN attribution.model_config AS cfg
WHERE cfg.is_active;
```

WHY A SCHEDULED TABLE RATHER THAN A MATERIALIZED VIEW

BigQuery's incremental materialized views support a limited SQL subset, while non-incremental materialized views support most queries but rerun the full query on every refresh [12]. Because the bridge depends on a window function and multiple joins, a scheduled rebuild (orchestrated with scheduled queries, Dataform or dbt) is simpler to reason about, cheap at opportunity grain, and produces a stable snapshot for each refresh.

4.5 Governed configuration

The credit share is a business decision, not a constant. Storing it as versioned configuration records who approved it and when it took effect, and lets the share change without editing reports.

05_config.sql · BigQuery lines 1–8

```sql
-- Governed model parameters: changed by agreement, versioned, auditable.
CREATE TABLE IF NOT EXISTS attribution.model_config (
  model_version          STRING  NOT NULL,
  marketing_credit_share NUMERIC NOT NULL,   -- e.g. 0.50 for an equal split
  approved_by            STRING,
  effective_from         DATE    NOT NULL,
  is_active              BOOL    NOT NULL
);
```

4.6 Automated quality checks

Four assertions run after every rebuild and must each return zero. They catch the problems that most often erode trust: duplicate attribution, campaigns that can never qualify, malformed windows and ambiguous configuration.

03_quality.sql · BigQuery lines 1–21

```sql
-- Data-quality assertions run after every rebuild. Each query must return 0.

-- A. Grain: exactly one row per opportunity.
SELECT COUNT(*) - COUNT(DISTINCT opportunity_id) AS duplicate_rows
FROM attribution.opportunity_campaign_influence;

-- B. Campaigns with contact members but no influence window (cannot qualify).
SELECT COUNT(DISTINCT cp.campaign_id) AS campaigns_without_window
```

```sql
FROM stg.campaign AS cp
JOIN stg.campaign_member AS cm ON cm.campaign_id = cp.campaign_id
WHERE cm.contact_id IS NOT NULL
  AND (cp.influence_start_date IS NULL OR cp.influence_end_date IS NULL);

-- C. Inverted windows (end before start).
SELECT COUNT(*) AS inverted_windows
FROM stg.campaign
WHERE influence_end_date < influence_start_date;

-- D. Exactly one active model configuration.
SELECT ABS(COUNTIF(is_active) - 1) AS config_errors
FROM attribution.model_config;
```

4.7 Semantic model in Power BI

The bridge becomes a single fact table at opportunity grain, surrounded by conformed dimensions. The date dimension plays two roles. Microsoft's guidance is to duplicate role-playing dimension tables when visuals must filter by both roles at once, and to use inactive relationships in specific circumstances, such as when reports do not need to filter by different roles simultaneously [14]. Here the pipeline date drives nearly every report, so it is the active relationship and the created date is activated only inside one measure.

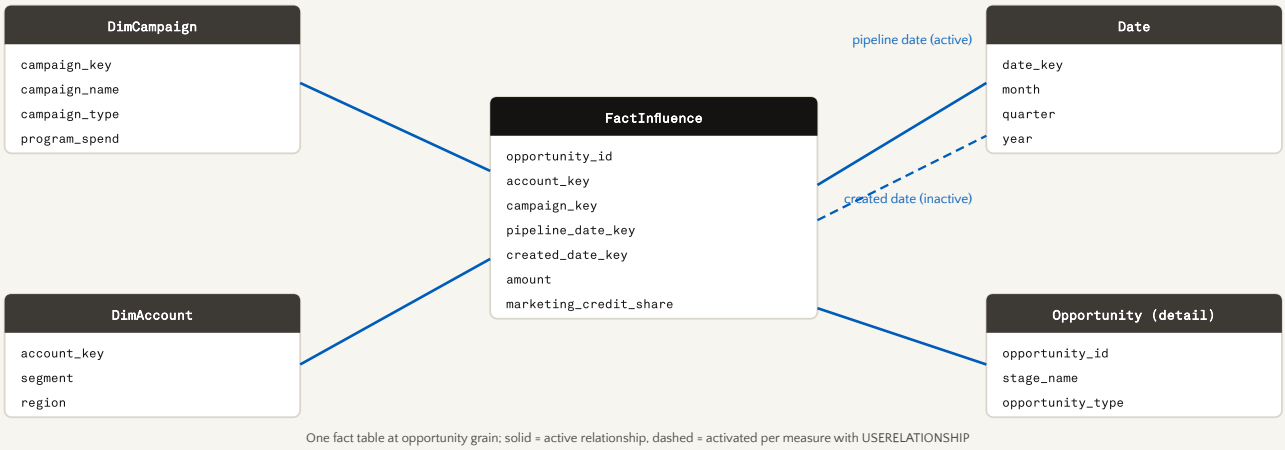

Figure 5 · Power BI semantic model.

04_measures.dax · Power BI lines 1–27

```dax
-- Core measures on the FactInfluence table (Power BI semantic model)

Influenced Opportunities =
DISTINCTCOUNT ( FactInfluence[opportunity_id] )

Influenced Pipeline (Gross) =
SUM ( FactInfluence[amount] )

Influenced Pipeline (Marketing Credit) =
SUMX (
    FactInfluence,
    FactInfluence[amount] * FactInfluence[marketing_credit_share]
)

-- Same measure, sliced by the date the opportunity was created
-- (inactive relationship activated only for this calculation)
Influenced Pipeline by Created Date =
```

```
CALCULATE (
    [Influenced Pipeline (Marketing Credit)],
    USERELATIONSHIP ( FactInfluence[created_date_key], 'Date'[date_key] )
)

Program Spend =
SUM ( DimCampaign[program_spend] )

Pipeline per Program Dollar =
DIVIDE ( [Influenced Pipeline (Marketing Credit)], [Program Spend] )
```

A drillthrough page filtered by campaign and pipeline month lists the opportunities behind any number [15]. The drill filters must use the **same date role** as the report that launches them; a summary built on the pipeline date that drills into a list filtered by created date shows a different set of deals and undermines trust.

4.8 Security, governance, scalability and reliability

Dimension	Design response
Security	No contact PII leaves staging. Least-privilege IAM per dataset; reports consume only the attribution and mart layers.
Access control	Power BI row-level security by region or segment where pipeline visibility is restricted.
Governance	Rules documented and signed off; credit share versioned in `model_config`; quality assertions gate publication.
Scalability	Opportunity grain keeps the bridge small relative to activity data; partitioning by month and clustering by account and campaign bound scan costs.
Reliability	Deterministic tie-breaking makes every rebuild reproducible; a failed assertion stops the refresh instead of publishing a wrong number.
Portability	Standard SQL with a window function; the same logic ports to other cloud warehouses with minor syntax changes.

5. Strategic Analysis

5.1 Options compared

Dimension	CRM-native campaign influence	Dedicated attribution platform	Warehouse-native bridge
Business value	Good for contact-level questions	Broad multi-touch views	Account-level influence on agreed rules
Feasibility	High; configuration only	Medium; implementation project	Medium; needs data engineering
Scalability	Bound to CRM limits	Vendor-managed	Warehouse-scale
Performance	Native reports	Vendor-dependent	Precomputed, fast to query
Cost	Included in license	Subscription	Warehouse compute and team time
Security	Inside the CRM	Data shared with a vendor	Inside existing cloud controls
Maintainability	Admin-managed	Vendor-managed	Owned code; needs tests and documentation
Interoperability	CRM reports	Vendor connectors	Any BI tool or downstream model
Key risk	Contact-role dependency	Rules diverge from internal definition	Key-person dependency without documentation
Best when	Contact roles are reliable	Digital multi-touch at volume	Account-based programs need one defensible number

5.2 Benefits

- **One definition.** Marketing and sales report the same influenced pipeline because both read the same table.
- **Auditability.** Every aggregate resolves to a list of opportunities and the campaign that won each one.
- **Separation of concerns.** Rules change in one place; reports, dashboards and downstream models inherit them.
- **Reuse.** The same bridge can feed program-level KPIs, such as event performance, without new attribution logic.

5.3 Limitations and trade-offs

- **Single-touch by design.** First touch is transparent but ignores later campaigns. Organizations that need multi-touch weights can extend the bridge to keep all qualifying campaigns with a weight column.
- **Account-level generosity.** Qualifying through any contact can credit campaigns for deals the engaged person never touched. Influence windows, campaign-type scoping and the credit share are the levers that keep this honest.
- **Dependence on campaign metadata.** Campaigns without influence dates cannot qualify; the quality checks surface them, but the fix belongs to campaign operations.
- **Correlation, not causation.** Influence measures association in time. Causal lift requires holdout or incrementality designs.

5.4 Build vs. buy, lock-in and organization

Building is justified when the attribution rules are specific to the business, the warehouse already holds CRM data, and leadership wants the definition owned internally. Buying is justified when the main need is digital multi-touch at volume and the organization accepts the vendor's model. Lock-in risk for the pattern described here is low: the core is a few hundred lines of standard SQL and a handful of measures. The organizational implications matter more than the technical ones. The model needs a named business owner for the rules, a data owner for the pipeline, and a regular forum where marketing and sales review exceptions.

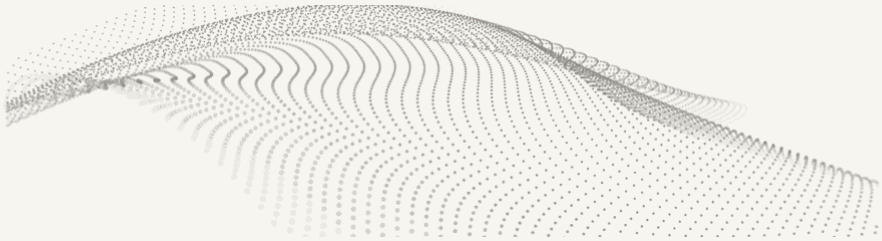

6. Case Study: One Number for Marketing and Sales

The Company is a global B2B software business selling to mid-market and enterprise accounts. Its marketing and sales leaders had already done the hardest part of attribution, which is agreeing on what it should mean. This is the story of how their revenue operations team turned that agreement into a number both sides trust, with Locus Analytica as their engineering partner.

The challenge

The CMO and CRO had designed a new **Marketing Influenced** model for trade shows and other field programs. Influence would be judged at the account level, campaigns would count only within an explicit influence window, and credit for opportunities not already sourced by marketing would be shared equally between marketing and sales. The definition was sound. The difficulty was that the CRM's native influence logic was contact-role based, and a first implementation that chained contacts to campaigns with outer joins was quietly missing opportunities.

The turning point: the business validated the data

The decisive moment came from the Company's own marketing operations lead. Working through campaign records by hand, the team found opportunities they could see in the CRM that the report did not show, and asked for the logic to be reviewed. That instinct to validate before adopting is what made the eventual model trustworthy.

What the team built

- **Account-first qualification.** The bridge was rebuilt to start from campaign membership, reach the account through each contact, and only then meet opportunities inside the influence window, exactly as the leadership definition read.
- **Hypotheses tested in isolation.** A plausible alternative explanation, that converted leads were the missing piece, was prototyped as separate delta queries without touching production. It did not hold, and the business confirmed the rule: contacts only, never leads.
- **The right date for each question.** At the business owner's request, pipeline was reported in the month it was recognized, while the creation date remained the test for influence.
- **Credit split last.** The equal split was applied as the final step, after attributed volumes had been validated against the CRM in full.
- **Every number drills.** Each aggregate links to the list of opportunities and campaigns behind it.

The outcome

The Company now operates a single definition of Marketing Influenced pipeline, owned by its leadership and computed once in the warehouse. Marketing and sales review the same numbers, and disagreements turn into conversations about specific opportunities rather than about whose spreadsheet is right. Because the decision lives in one bridge, the team reused it for event-level performance reporting without writing new attribution logic.

WHAT MADE IT WORK

Leadership agreed on the rules before the build. The business owner validated the output against the CRM. Engineering kept the decision in one place and made every number traceable. Locus Analytica supplied the architecture and the code; the Company supplied the definition, the validation and the discipline to adopt it.

6.1 Other applications of the pattern

Program	How the pattern applies	Expected impact
Field events and trade shows	Event campaigns carry influence windows around the event dates	Event budgets judged on influenced pipeline, not badge scans
Account-based advertising	Target-account programs qualify accounts during flight dates	Spend tied to pipeline in target accounts
Partner co-marketing	Partner campaigns qualify shared accounts	Clear view of partner-influenced pipeline
Executive programs	Dinners and councils qualify senior contacts' accounts	Visibility for low-volume, high-value touches

7. Implementation Roadmap

A typical sequence is shown below. Durations are indicative and depend mostly on data readiness and on how quickly leadership converges on the rules.

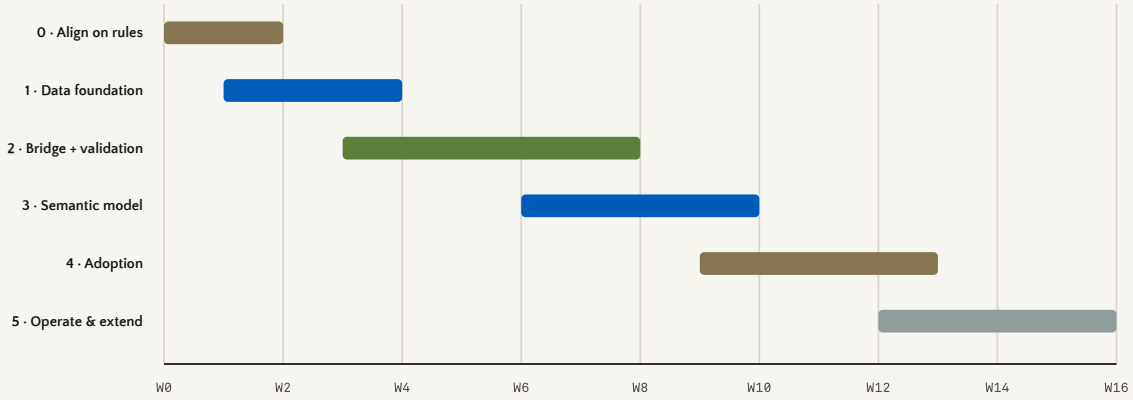

Figure 6 · Indicative phasing in weeks.

Phase	Key activities	Exit criteria
0 · Align on rules	Workshop with marketing and sales leadership; write the seven rules; define dates and credit share	Signed rule document
1 · Data foundation	Land Salesforce objects; build staging views; populate influence windows on campaigns	Staging reconciles to CRM counts
2 · Bridge + validation	Build bridge and quality checks; business owners validate sample campaigns against the CRM	Owners accept validation; all assertions pass
3 · Semantic model	Star schema, measures, drillthrough, row-level security	Reports match bridge totals
4 · Adoption	Train users; publish rule document; set the exception-review forum	Leadership uses the report in planning
5 · Operate & extend	Monitor quality; extend to multi-touch weights or program KPIs	Stable refresh; change log maintained

7.1 Capabilities required

A RevOps or marketing operations owner who can arbitrate rules; a data engineer comfortable with SQL window functions and orchestration; a BI developer for the semantic model; and a CRM administrator to maintain campaign metadata. In smaller teams one person may hold two of these roles; the business owner role should not be merged into engineering.

7.2 Risks and mitigations

Risk	Mitigation
Rules change mid-build	Version the rule document and the configuration table; rebuild, don't patch.
Campaigns lack influence windows	Quality check B; campaign-creation checklist in the CRM.
Silent record loss in joins	Window inside the join; reconciliation counts against the CRM.
Duplicate attribution on ties	Deterministic `ROW_NUMBER` ordering; grain assertion A.
Report and drill use different dates	One date role per report; drill filters reuse it.
Key-person dependency	Documented rules, code in version control, tests in the pipeline.

8. Business Impact and KPIs

The primary value is decision quality: program budgets can be allocated on influenced pipeline that marketing and sales both accept. Secondary value comes from operational efficiency (no manual reconciliation each cycle) and risk reduction (one audited definition instead of several competing ones). The following KPIs make the impact measurable without inventing targets the organization has not set.

KPI	Definition	Purpose
Influenced pipeline (marketing credit)	Σ amount × credit share for influenced opportunities, by reporting period	Headline outcome
Influenced opportunities	Distinct opportunities with a credited campaign	Volume behind the value
Pipeline per program dollar	Influenced pipeline ÷ program spend	Budget efficiency by campaign
Influence coverage	Influenced opportunities ÷ all new opportunities	Reach of marketing programs
Window completeness	Campaigns with valid windows ÷ campaigns with contact members	Metadata health
CRM reconciliation variance	Difference between bridge and CRM-validated samples	Trust and data quality
Refresh SLA	Share of refreshes completed with all assertions passing	Operational reliability

9. Recommendations and Conclusions

Account-based go-to-market needs account-based measurement. The pattern in this paper is deliberately simple: qualify accounts through their contacts, test campaigns against an explicit window, pick one winner deterministically, and serve the result through a governed semantic model. Its strength is not sophistication but ownership. The rules belong to leadership, the decision is made once, and every number can be explained.

Conditions for successful adoption

- Executive sponsorship from both marketing and sales, with a written, signed definition.
- Campaign operations that maintain influence windows as part of campaign setup.
- Validation by business owners before the first executive review.
- Engineering discipline: version control, automated assertions and documented configuration.

Future outlook

Two developments are likely to shape this space. Algorithmic attribution will keep expanding in digital channels, as Google's 2023 consolidation shows [5], which raises the value of a transparent account-level model as a reference point. And as buying-group thinking matures [4], organizations will want to move from "which campaign influenced the account" toward "which members of the buying group engaged, and when". A clean, governed bridge at opportunity grain is the foundation for both.

10. References

[1] Gartner, "B2B Buying: How Top CSOs and CMOs Optimize the Journey," Gartner Insights. Accessed Nov. 19, 2026. https://www.gartner.com/en/sales/insights/b2b-buying-journey

[2] LeanData, "B2B Buying Groups & Today's Go-to-Market Motions" (citing Gartner on buying-group size). Accessed Nov. 19, 2026. https://www.leandata.com/blog/buying-groups-todays-b2b-go-to-market-motions

[3] Forrester, "Demand And Account-Based Marketing: Planning Assumptions 2022," Sept. 7, 2021 (citing the 2021 Forrester Revenue Operations and Buying Groups Survey). https://reg.techtarget.com/rs/095-PCN-759/images/Demand%20And%20AccountBased%20Marketing%20Planning%20Assumptions%202022.pdf

[4] M. Threadgill, "Aligning Account-Based Marketing Measurement To The B2B Revenue Waterfall," Forrester, Aug. 19, 2021. https://www.forrester.com/report/aligning-account-based-marketing-measurement-to-the-b2b-revenue-waterfall/RES176112

[5] Google, "First click, linear, time decay, and position-based attribution models are going away," Google Ads Help, Apr. 6, 2023. https://support.google.com/google-ads/answer/13427716

[6] Salesforce, "CampaignMember," SOAP API Developer Guide (current version). Accessed Nov. 19, 2026. https://developer.salesforce.com/docs/atlas.en-us.api.meta/api/sforce_api_objects_campaignmember.htm

[7] Salesforce, "CampaignMember," Object Reference for the Salesforce Platform (archived release 228), usage notes. https://developer.salesforce.com/docs/atlas.en-us.228.0.object_reference.meta/object_reference/sforce_api_objects_campaignmember.htm

[8] Salesforce, "CampaignInfluence," Object Reference for the Salesforce Platform. Accessed Nov. 19, 2026. https://developer.salesforce.com/docs/atlas.en-us.object_reference.meta/api/sforce_api_objects_campaigninfluence.htm

[9] L. Mazalon, "Complete Guide to Salesforce Campaign Influence," Salesforce Ben, updated May 11, 2026. https://www.salesforceben.com/salesforce-campaign-influence-marketers-guide/

[10] Google Cloud, "Load Salesforce data into BigQuery" and "Introduction to Salesforce data transfers," BigQuery documentation. Accessed Nov. 19, 2026. https://cloud.google.com/bigquery/docs/salesforce-transfer

[11] Google Cloud, "Window function calls," GoogleSQL for BigQuery reference. Accessed Nov. 19, 2026. https://docs.cloud.google.com/bigquery/docs/reference/standard-sql/analytic-function-concepts

[12] Google Cloud, "Introduction to materialized views," BigQuery documentation. Accessed Nov. 19, 2026. https://docs.cloud.google.com/bigquery/docs/materialized-views-intro

[13] Microsoft, "Understand star schema and the importance for Power BI," Microsoft Learn. Accessed Nov. 19, 2026. https://learn.microsoft.com/power-bi/guidance/star-schema

[14] Microsoft, "Active vs inactive relationship guidance," Microsoft Learn. Accessed Nov. 19, 2026. https://learn.microsoft.com/power-bi/guidance/relationships-active-inactive

[15] Microsoft, "Set up drillthrough in Power BI reports," Microsoft Learn. Accessed Nov. 19, 2026. https://learn.microsoft.com/power-bi/create-reports/desktop-drillthrough

LOCUS/ANALYTICA

DATA SCIENCE & ENGINEERING

ABOUT LOCUS ANALYTICA

Locus Analytica is a data science and engineering practice that helps revenue teams turn fragmented go-to-market data into decisions they can defend. The organization described in this paper is anonymized, and all schemas, field names and code are reference implementations prepared by Locus Analytica for illustration.

© 2026 Locus Analytica · Data Science & Engineering