iCustomer

GUIDE · PERSISTENT DATA GOVERNANCE

Switzerland, Not Walmart

Persistent data governance for the agentic era: how RevOps, marketing ops, data teams and the CIO office manage contacts, accounts and signals across every internal system and external data vendor.

Abhi Yadav, Co-founder and CEO, iCustomer · November 2026 · icustomer.ai

Contents

For RevOps, sales ops, marketing ops, data teams, data product managers and the CIO office at enterprise B2B companies.

01	Agents find your data problems in days	3
02	The Walmart of data won't fix it	4
03	What teams need: a Switzerland of data, run as an identity router	5
04	Accounts aren't flat: hierarchies and buying centres	6
05	One data model for sales, marketing and CX	7
06	Bring your own data, under one set of rules	8
07	Start with the contacts you already have	8
08	Eight rules every team signs	9
09	Five measures that set priority	9
10	What each team and each agent gets	10
11	Use cases that run on the same IDs	11
12	What data product teams get	13
13	How to evaluate an identity layer	13
14	How Agentic ID works for your teams	14
15	Appendix: running the program	18

Click any section to jump to it. Your PDF reader's bookmarks panel lists the same sections.

Agents read your CRM, your warehouse and your vendor feeds at machine speed, so every gap in your identity and data quality shows up within days. Buying more coverage won't close those gaps. What closes them is a neutral, persistent identity layer that every team and every agent works from.

Signs your contact data foundation is unreliable or thin. If three or more of these are true, the rest of this guide is for you:

- Different teams bought different data vendors, and nobody owns the full list.
- The same buyer exists as separate records in CRM, the marketing cloud and the ABM platform.
- Reps don't trust titles or emails in Salesforce, so they look people up again by hand.
- Account hierarchies live in spreadsheets, or nowhere.
- You can't say which vendor supplied a given field, or what its license lets you do with it.
- Champion job changes surface after the renewal is already at risk.
- An agent pilot stalled because the data underneath it wasn't good enough.
- **Data quality:** required fields are missing or unverified on more than half your contacts, and nobody can say how much of what's there is correct.
- **Decay rate:** you don't know how fast your contact data goes stale, so reverification happens once a year, if at all.
- **Data model:** sales, marketing and CX each model accounts and contacts their own way, so the same buyer can't be followed across systems.
- **Accuracy:** vendors disagree on the same contact's email or title, and nothing decides which answer is right.
- **Decision-maker coverage:** most target accounts have one or two known contacts, and whole buying centres have nobody mapped.

Agents find your data problems in days

A human SDR who hits a dead email shrugs and moves on. An outbound agent hits 2,000 of them before lunch. A research agent asked to brief a rep on an account pulls three records for the same VP of Finance, with two emails and three titles, and picks one. A scoring agent ranks an account as cold because half its activity sits on a duplicate.

None of these errors is new. Agents make them faster, at higher volume, and in front of customers. Problems that used to surface in a quarterly CRM cleanup now show up in your bounce rate, your sender reputation and your reps' trust by the end of the week.

The source data is part of it. Vendors that look identical in a demo return very different numbers of wrong emails once you test them on contacts you know are right. At human speed that's an annoyance. At agent speed it can burn a sending domain.

The deeper problem is identity. Your agents can't tell which record is the real person, which account a contact belongs to, or which activity belongs to whom, because nothing in your stack decides that once and keeps deciding it.

The Walmart of data won't fix it

The default response is to buy more. Large contact and firmographic vendors sell the Walmart of data: every aisle stocked, hundreds of millions of records, coverage as the headline. Most enterprises run several of them, plus a waterfall on top.

Coverage helps you find people. It doesn't tell you which answer is right when two vendors disagree, how fast each source goes stale, or which contact in your own CRM is the same person as the one in your marketing cloud. And it comes with three structural costs:

- **Rented snapshots.** Titles, emails and reporting lines change every month. A record enriched in January is a guess by June, and the vendor's next refresh overwrites your history.
- **Someone else's IDs.** Each vendor keys records to its own identifiers. Switch vendors and every downstream join, list and model has to be rebuilt.
- **Volume without priority.** More records means more for agents to act on, with nothing telling them which records are trustworthy or which accounts matter now.

So most enterprises stack vendors: one for contacts, one for technographics, one for intent, a few regional specialists. Often it's one vendor per tool: an enrichment app in the CRM, another in the marketing automation platform, an intent feed inside the ABM platform. Each tool ends up with its own copy of the buyer. That's ten or more contracts, each with its own data model, field names and refresh cycle. The data team spends months mapping fields and unifying records, and the business ends up using a small slice of what it paid for. By the time the unified file is ready, part of it has already decayed.

Add agents to that stack and you get more wrong answers, delivered faster.

Example: one global account, four buying centres. Take a 40,000-person manufacturer headquartered in Chicago, with subsidiaries in Germany and Singapore, rolling out a new software platform. In most CRMs it's one account record. In reality it looks like this:

Location and unit	Buyers	Users	Influencers
Global HQ, Chicago (parent)	CFO signs the enterprise agreement; VP Procurement negotiates	Almost none	CIO and CISO run the security review
North America sales BU	VP Sales Ops owns the budget	400 reps and managers	RevOps director, the champion
EMEA subsidiary, Frankfurt	Regional finance director signs the local entity contract	150 legal and HR staff	Data protection officer, who can block on consent terms
APAC region, Singapore	Country GM holds the expansion budget	60 ops staff on a free trial, signed up with personal emails	Regional IT lead

The CFO signs but never logs in. The 60 trial users in Singapore are the strongest expansion signal in the account, and they're invisible if their sign-ups don't resolve to the APAC unit. If the RevOps champion in North America changes jobs, the renewal is at risk and a new-logo lead opens at their next company. Each of these people needs one persistent contact ID, tied to the right place in the hierarchy, with a role per buying centre.

What teams need: a Switzerland of data, run as an identity router

A Switzerland of data is a neutral, persistent identity layer that you own. It sits in your data cloud (Snowflake, Databricks or BigQuery), takes no side between vendors or internal systems, and gives every contact, account and activity one ID that lasts.

Every source connects to it: Salesforce, the marketing cloud, product data, web activity, events, and as many external vendors as you want. It doesn't care where a record came from. It judges each source on coverage, quality and decay rate, keeps the best current answer for each field, and attaches every new signal to the right ID. Intelligence and scores build on top, and they persist when you add or drop a vendor.

	Walmart of data	Switzerland of data
What you buy	Volume and coverage	Persistent identity and judgment across sources
Whose IDs	The vendor's	Yours, anchored on your CRM
When a vendor changes	Rebuild joins, lists and models	Plug the new source in; IDs stay put

	Walmart of data	Switzerland of data
Decay	Hidden until the next refresh	Measured per source and per field
History	Overwritten	Kept, so intelligence compounds
What agents get	Raw records to sort out themselves	One resolved ID, its history and its priority

In practice, that Switzerland is an aggregator that sits between you and every vendor: an identity router. It connects to 100+ data vendors and sources, normalizes their data models into one schema, and resolves everything to your persistent IDs. You stop buying data in bulk. You ask for what a record needs, when it needs it.

- **Waterfall on demand.** A record is missing a verified email or a key buying role. The router asks the vendor most likely to have a correct answer for that field, checks it, and moves to the next vendor only if it has to. You pay for the answer you used, not for ten copies of a database.
- **Ongoing upkeep.** The same router reverifies records on a schedule set by each source's decay rate, so the contacts you care about stay current without a yearly re-enrichment project.
- **Identity that survives job changes.** When a champion moves to a new company, the old and new records link to the same persistent contact ID. Their history comes with them, and both accounts get the signal.
- **One ID across platforms.** That ID maps to Salesforce, the marketing cloud, ad platforms, product data and agent tools, so every system and every agent refers to the same contact and account.
- **Your context graph, under your control.** Resolved identities, relationships, signals and history live in your data cloud. Vendors supply answers. They don't own the graph, and dropping one doesn't break it.

Accounts aren't flat: hierarchies and buying centres

Enterprise accounts are trees. A global customer has a parent, subsidiaries, regions and business units, and each of them can buy on its own. The one-size-fits-all model, one account with a list of contacts scored as leads, misses how these companies actually buy.

Treat hierarchy and buying centres as data you maintain every week:

- **Hierarchy.** Every account ID sits in a tree (parent, subsidiary, region, business unit) that stays current through reorgs, acquisitions and spin-offs. Territories, ownership, rollups and renewal reporting all hang off it.
- **Buying centres.** Each business unit or product line can have its own buying centre: economic buyer, champion, technical evaluator, IT and security, procurement, finance. Each one changes on its own schedule.

- **Buyers and users.** Users show up in product data. Buyers show up in deals and meetings. Link both to the same contact IDs and you can see who uses the product, who signs and who influences.
- **Roles per buying centre.** The same VP can be the economic buyer for one product line and an influencer for another, so roles live on the contact and buying centre pair, not on the contact alone.

That's what turns a contact list into a treasure map. Coverage gets measured per buying centre, so a rep can see that a customer's finance team is fully mapped while its new APAC unit has nobody at all.

One data model for sales, marketing and CX

Sales, marketing and customer success usually model the customer three different ways: accounts and opportunities in CRM, leads and lists in the marketing cloud, users and tickets in CX tools. A shared identity layer needs one model all three read from. Think of it as a world data model. It covers every company and person in your market, not only the ones already in your CRM, so a new account or a job change lands in a structure that already exists.

Entity	What it holds	Who uses it
Account	Persistent ID, place in the hierarchy, firmographics, TAM, SAM, ICP or customer tag	Everyone
Buying centre	Business unit or product line, roles it needs, coverage	Sales, CS, ABM
Contact	Persistent ID that survives job changes, verified channels, consent	Everyone
Role	A contact's role in a buying centre: buyer, champion, user, influencer	Sales, CS, marketing
Activity and signal	Visits, product use, events, replies, intent, sales and support touches, each tied to IDs	Data science, every downstream team
Source	Which vendor or system supplied each value, when, and under what license	CIO office, data team

Each tool keeps its own view. Salesforce shows accounts, contacts and roles. The marketing cloud shows contacts, consent and segments. CX tools show users and health. Underneath, every view points to the same IDs.

Bring your own data, under one set of rules

Nobody has to give up their tools. Sales keeps working in Salesforce. Marketing keeps its marketing cloud and its list uploads. ABM keeps its ad platforms. Data science keeps its notebooks. Each team brings data in through its own systems and workflows, in a common format, under rules everyone agreed to.

The CIO office, or whoever runs enterprise data governance, owns those rules: how records enter, how identity resolves, which source wins for each field, who can see what, and how consent is enforced. It also tracks what each vendor's license allows: which fields can go to which tool, which can be uploaded to ad platforms, and which have to stay in the warehouse. Because the identity layer runs in your own cloud, contact data stays under your security and residency controls. It doesn't own every list or every score.

Each team owns its data products inside the rules. RevOps owns account hierarchy and ownership. Marketing ops owns contactable audiences. ABM owns target lists and ad audiences. Data science owns the scores. The identity layer is what lets all of that line up: once unique contacts, accounts and activities are resolved across systems, sales feels empowered in Salesforce, marketing has contacts it can trust for campaigns, and ABM and data science work from the same accounts.

Start with the contacts you already have

New records are the easy part. Native CRM enrichment fills firmographics and shell accounts at creation, and most teams won't rip that out. The mess is the installed base: years of contacts in CRM, the marketing cloud and old list uploads, with stale titles, dead emails, the same person under two accounts, and no history of what they've done.

That installed base is also where agents and signals land. A web visit, a webinar sign-up or a badge scan is only useful if it attaches to the right person and account. Attach it to a duplicate and the intelligence splits in two.

- Mirror every existing contact and account from CRM and the marketing cloud into the warehouse, with its source and last-verified date.
- Resolve them to persistent contact and account IDs. Merge duplicates in the warehouse, then push the merges back.
- Reverify stale records on a rolling schedule, starting with contacts at target accounts.
- Attach every new signal to those IDs as it arrives, so each record keeps building history across uploads, campaigns and quarters.

Eight rules every team signs

This is the data contract. Write it down, get RevOps, marketing ops, data science and the CIO office to sign it, and review it every quarter.

1. **One door for decisions.** Native CRM enrichment can keep filling new records in Salesforce. Every record, whatever wrote it, is mirrored into the warehouse with its source, and identity, survivorship and scores are settled there, then written back to CRM and the marketing cloud.
2. **One persistent ID per contact, account and activity.** Anchor the account ID where records are born. For most B2B teams that's the Salesforce account ID. The identity layer maps every internal and vendor ID back to it.
3. **Provenance on every field.** Each value carries its source, capture date and confidence. If you can't say where a title came from, neither a rep nor an agent should trust it.
4. **Survivorship by field, agreed in writing.** First-party usually wins for email, since a reply or a login is proof. A recently verified vendor value might win for title. Write it once and apply it everywhere.
5. **Decay is measured, not assumed.** Track how fast each source goes stale, per field. Set reverification windows from those numbers, and flag expired values instead of deleting them.
6. **Vendors earn their place against ground truth.** Test each one on about 500 contacts you know are right (replies, logins, closed-won, current customers), looking each one up by LinkedIn URL and by name plus company. Score coverage, wrong answers and decay. A vendor with less coverage and fewer wrong answers can be worth more.
7. **Consent travels with the ID.** Consent and suppression status live on the contact ID and flow to every tool and agent.
8. **Agents read and write through the identity layer.** No agent acts on a raw vendor feed or an unresolved record. Agents get resolved IDs, history and priority through one API, and what they learn writes back to the same IDs.

Five measures that set priority

The identity layer scores two things on the same five measures: every source, and every record. Sources that score well win survivorship. Records that score well get acted on first, by people and by agents.

Measure	For a source	For a record
Coverage	Share of your target accounts and buying roles it can fill	Are all the key buying roles at this account filled?
Quality	Correct-answer rate against ground truth	Is the email verified and the title confirmed?

Measure	For a source	For a record
Decay rate	How fast its values go stale, per field	How long since this record was last verified?
Intelligence	Does it add fit, hierarchy or role data you use?	Fit, engagement and role scores
Signals	Does it bring intent, job changes or other external signals you can attach to your IDs?	Website visits, product use, intent, job changes, events, replies, sales touches

Signals come from inside and outside the company, and they only add up when they land on the same IDs:

Signal type	Examples	Source
First-party (1P)	Website visits and pages viewed, product sign-ups and usage, form fills, email and ad engagement, sales activity, support tickets, event check-ins	Internal: your own systems
Second-party (2P)	Partner webinar and co-sponsored event lists, marketplace and integration partner data, agency and reseller data	External, shared by partners
Third-party (3P)	Topic intent, job and title changes, hiring, funding, tech installs, executive moves, news	External: data vendors and the open web

Every signal gets resolved to a contact ID and an account ID, with its source, timestamp and consent status. A job change from a vendor, a visit to your pricing page and a spike in product use then read as one story on one account.

Once activity attaches to persistent IDs, data science can score across the funnel on observed behavior: account fit, account engagement, and the likelihood that a contact is in the buying group. Scores refresh as signals arrive, push to Salesforce, the marketing cloud and ad audiences, and retrain on meetings, pipeline and closed-won. That loop is what makes the intelligence persistent.

What each team and each agent gets

Who	Owns	Gets from the identity layer	Acts in
RevOps / Sales Ops	Account hierarchy, territories, ownership	Deduplicated accounts, full buying groups with missing decision makers filled, priority scores	Salesforce
Marketing ops	Contactable audiences, segments	Verified, consented contacts with current titles and engagement history	Marketing cloud, MAP

Who	Owns	Gets from the identity layer	Acts in
Growth, ABM and paid media	Target lists, ad audiences	Account and contact lists with clean match keys, refreshed on a schedule	LinkedIn, Google, Meta, programmatic
Data science	Scoring and models	Unified activity history per account and contact across the funnel	Writes scores back to the warehouse
Customer success / CX	Health, adoption, renewals	Users and buyers per buying centre, champion changes, usage tied to contact IDs	CS platform, Salesforce
CIO office / governance	The contract, access, consent, audit	Lineage, source quality and decay metrics	Data cloud
Sales and marketing agents	Nothing; they consume	Resolved IDs, history, consent and priority through one API	Wherever they run

Sales usually feels it first. Accounts stop duplicating, reps can see who else to call in the buying group, and the priority list matches the accounts marketing is running ads against.

Use cases that run on the same IDs

Once accounts, contacts and activities share persistent IDs, every use case reads and writes the same records. Champion tracking, visitor identification, renewal risk and lead scoring stop being separate tools with separate copies of the truth.

Start with the account universe. Tag every account ID with where it sits, and update the tags as firmographics change:

- **TAM:** every company that could buy from you.
- **SAM:** the part you can serve today, by region, segment and product.
- **ICP accounts:** SAM accounts that pass Fit.
- **Customers:** with renewal date, products owned and the buying centres already sold to.

Every signal below rolls up to these tags, so you can see movement: a SAM account that starts to fit, an ICP account that goes active, a customer that goes quiet.

Score Fit at the door. Every new account or contact, whatever the source (form fill, list upload, vendor, rep entry, resolved website visitor), gets resolved to an ID and scored for Fit before anyone routes it. Account Fit checks firmographics, tech stack and hierarchy against the ICP. Contact Fit checks role against your buying centres. Records under the floor go to nurture. Records above it

enter the ranking. Fit gets rescored when the company changes, so it never freezes at the value it had on entry.

Rank with FIRE, persistently. FIRE stands for Fit, Intent, Recency and Engagement. Fit decides who's eligible. Intent covers research on your category, visits to high-intent pages, hiring and executive changes. Recency measures how fresh the last meaningful signal is, so stale accounts drop on their own. Engagement measures depth across the buying group: how many contacts are active, meetings, replies, product use. Scores live on the account and contact IDs, roll up to the buying group, and keep their history. Reps and agents see the rank and which component moved it.

Use case	Signals it uses	What gets scored or flagged	Where it lands
Champion tracking	Job and title changes, engagement of known champions at customers and past deals	Champion left (renewal risk), champion joined an ICP account (new-logo lead), new champion emerging	Salesforce task for the AM or AE
Website visitors	Visits resolved to accounts, and to contacts where consented; pages viewed	Raises Intent and Recency; flags unworked ICP accounts that show up	SDR queue, Salesforce, retargeting audiences
TAM, SAM and ICP signals	Third-party intent, hiring, funding, tech installs, executive hires	Accounts moving up a tier; SAM accounts that newly fit	ABM tiers, ad audiences, SDR priority
Renewals	Product usage by user, champion and economic buyer in seat, support history, engagement	Renewal risk per customer, 6 to 12 months out	AM and CS in Salesforce
Expansion	Users in adjacent teams, buying centres not yet sold, usage growth, new executives	Expansion readiness per buying centre	AE plays, targeted campaigns
New accounts and contacts	Any source, at creation	Fit score and routing decision	Owner assignment or nurture
Events and offline lists	Badge scans, sponsorship and dinner lists, booth conversations	Matched to existing contact IDs before upload; new contacts scored for Fit	Salesforce, marketing cloud, follow-up sequences
Product-led signups	Free and self-serve users, work and personal emails, product usage	Users resolved to accounts and buying centres; accounts with many active users and no enterprise contract	AE and SDR queues, Salesforce

Use case	Signals it uses	What gets scored or flagged	Where it lands
Prioritization	All of the above	FIRE rank across ICP accounts and customers, in separate queues	Rep daily lists, agent work queues

The point of running all of this on one identity layer is that the signals reinforce each other. A website visit from a contact at a customer, where the champion just left and a new VP just joined, is one story on one account. Spread across four tools, it's four alerts nobody connects.

What data product teams get

If your company runs customer and account data as an internal product, the identity layer becomes a dependency for every product downstream: scoring models, territory tools, ABM audiences, agent tools. Hold it to the same bar as any other upstream data product:

- **A contract.** A versioned schema for accounts, buying centres, contacts, roles, signals and sources, with notice before any change that could break a consumer.
- **Service levels.** Freshness targets per field, response times for real-time resolve and verify calls, and API uptime.
- **Observability.** Resolution rate, decay and vendor accuracy published as tables in your warehouse, so product teams monitor them like any other pipeline.
- **One interface.** The same API, SDK and MCP endpoint for internal apps and agents, so no team builds its own matching logic.
- **Clear ownership.** The data team owns the identity product. RevOps, marketing ops and CX are its customers, and the CIO office sets policy.

How to evaluate an identity layer

Whether you build it or buy it, hold it to the same questions. The CIO office, RevOps and the data team should all sign off on the answers.

Question	What a good answer looks like
Where does it run?	In your data cloud (Snowflake, Databricks or BigQuery), under your security and residency controls
Whose IDs are they?	Yours, anchored on your CRM, and they stay if you change providers

Question	What a good answer looks like
Is it vendor-neutral?	It connects to many data vendors and your own systems, and you can add or drop any of them without breaking IDs
How is data bought?	By waterfall, per field, on demand or on a decay schedule, paid by use
Can it explain every value?	Source, capture date, confidence and license on every field
Does identity survive job changes?	The same contact ID follows a person to a new company, with their history
Does it model complex accounts?	Hierarchies, buying centres and roles per buying centre, beyond a flat account and contact list
Can agents use it safely?	One API or MCP endpoint that returns resolved IDs, consent and priority, and writes back what agents learn
Can you measure it?	Coverage, depth, completeness, duplicates, staleness and per-vendor decay, reported by tier
Can your data team build on it?	A versioned schema, documented API and SDK, and change notices, so internal data products can depend on it
Who is controller and who is processor?	You stay the controller of your data; the provider acts as your processor
How is personal data handled for agents?	Agents get masked values by default, and raw values only through scoped keys
What happens if you leave?	Your IDs and mappings stay in your warehouse; you lose only updates and new matches

How Agentic ID works for your teams

Agentic ID is the identity router this guide describes. It runs inside your data cloud and acts as the glue between your internal systems, your external data sources, and every person and agent that acts on them. Everything it produces stays first-party: your IDs, your graph, your rules.

Open source, in your data cloud, managed by you. Agentic ID is open source. The ID spec, ledger protocol, match-key format, SDK and CLI are public, and you can run your own node. Your engineers can inspect exactly how IDs are minted and matched. The IDs live in your data cloud and your team manages them under your own rules. Each person or company owns its own ID and consent, and your brand owns its relationship records. iCustomer runs the enterprise layer on top: the crosswalk across 100+ sources and verified contact and company data.

One ID to pull any external signal. Every contact and account carries the same Agentic ID, so you use it to pull signals and attributes from any connected vendor without remapping each vendor's schema. Sources come pre-configured and unified at the entity level (person, company) and the identity level (email, phone, device, domain). A job change from one vendor, intent from another and a tech-stack update from a third all land on the same record.

Where it fits in iCustomer. Agentic ID is the identity layer of one platform with three ways in. iReveal adds first-party website and agent-traffic signals. Growth Brain runs FIRE scoring, decisions and activation on the same IDs. Teams can start with any one of the three, and all of them share one credit system.

How it works

1. **Connect.** Point it at Salesforce, your marketing cloud, product data, website and event sources. Agentic ID already connects to 100+ data vendors and 200+ apps through one API, so you don't sign and integrate each vendor yourself.
2. **Resolve.** Every contact, account and activity resolves to a persistent ID anchored on your CRM. Duplicates merge, hierarchies and buying centres attach, and when a contact changes jobs the same ID follows them.
3. **Verify.** Before a rep or an agent acts on a record, the verify service checks it: is the email live, is the person still at the company, is the title current, what's the consent status.
4. **Enrich by waterfall.** Missing fields get filled from the vendor most likely to be right for that field, on demand or on a decay schedule. You pay for the answers you use.
5. **Attach signals.** First-party website visits and product use, second-party partner lists, and third-party intent and job changes all land on the same IDs, with source, time and consent.
6. **Activate.** Resolved audiences sync to Salesforce, the marketing cloud and ad platforms. Agents read and write through the same API or MCP endpoint.

Data vendor vs enrichment workflow tool vs Agentic ID, verb by verb

An identity layer for people and agents has ten jobs: identify, discover, find, resolve, validate, enrich, tag, score, match and monitor. A big data vendor, an enrichment workflow tool and an identity router built on your data cloud handle each one differently. Agentic ID is also an alternative way to get the data itself: use its 100+ sources, or bring your own keys and vendor contracts and run them through the same router.

Verb	Big data vendor	Enrichment workflow tool	Agentic ID: identity infrastructure on your data cloud
Identify	Visitor identification inside the vendor's platform	Visitor data comes from other tools you connect	First-party, server-side visitor resolution on your own domain, to contact and account IDs where consent allows. AI agents visiting your site are tied to the person they act for.

Verb	Big data vendor	Enrichment workflow tool	Agentic ID: identity infrastructure on your data cloud
Discover	Search the vendor's own company and contact database	Build lists from providers in the tool's marketplace	Discover new and look-alike accounts, plus the buyer contacts in each account's decision-making unit (US contacts, companies globally), across 100+ independent sources, written to your graph
Find	Contacts from one database	Waterfall across providers, set up workflow by workflow	Fill missing buying-centre roles per account, by waterfall, only where there are gaps
Resolve	Your records matched to the vendor's records and IDs	Matching logic you build per table	One persistent ID per contact, account and activity, anchored on your CRM and kept through job changes
Validate	The vendor verifies its own data	A verification step you add to each waterfall	Checks the entity is real and the email is live before any person or agent acts
Enrich	Fields from one dataset	Fields from many providers, paid in credits	Each field from the best-scoring source, on demand or on a decay schedule
Tag	The vendor's taxonomy	Columns you define per table	One taxonomy (industry, persona, role, buying centre) across internal and external records
Score	The vendor's fit and intent scores	Formulas or AI columns you build	Fit at entry and FIRE ranking over time, with history on the ID
Match	Audience exports to ad platforms	Ad audience sync from tables	Multi-key matching of verified IDs to ad platforms, with match rates up to 90% in documented results
Monitor	Job-change and intent alerts on the vendor's data	Signals and job-change tracking on records inside the tool	Watches your IDs for job changes, decay and new signals, and updates every downstream system
Built for	Sellers and marketers looking up records	GTM engineers building enrichment workflows	Enterprise identity infrastructure shared by RevOps, marketing ops, data teams, the CIO office and agents

Verb	Big data vendor	Enrichment workflow tool	Agentic ID: identity infrastructure on your data cloud
Where data lives	The vendor's platform, licensed to you	The tool's platform	Your data cloud: Snowflake, Databricks or BigQuery
Whose IDs	The vendor's	Rows in the tool's tables	Yours, and you keep the mappings if you leave
Data sources	One vendor's database	Providers in the tool's marketplace, plus your own keys	100+ independent sources, plus your own keys and existing vendor contracts
Neutrality	The vendor is the data source	Neutral across providers inside the tool	Neutral across vendors, apps and your own systems
Agent access	API or MCP grounded in the vendor's data	API or MCP to query and trigger workflows	One API or MCP endpoint for resolve and verify that returns consent and priority; works with Claude, ChatGPT, Gemini and Grok agents
Consent	Managed in your own tools	Opt-outs and lawful basis managed by your team	Stored on the contact ID and checked on every sync
Pricing	Annual license by package	Credits plus platform actions	Credits used per call: lookups, verification and persistent monitoring are all billed by usage, with no charge for misses

Consent built in. Consent and suppression status come in from your consent tools and forms, live on the contact ID, and get checked before anything syncs to email, ads or an agent. Website visitors are identified only where consent allows. When someone opts out, every downstream system gets the change.

Persistence you can measure. Website visits, product activity, ad audiences and CRM records all point to the same IDs, so history accumulates instead of resetting. In documented customer results, Agentic ID identified about 70% of visits from consented visitors and reached paid-media match rates of up to 90%.

One service for people and agents. Reps see verified buying centres in Salesforce. Marketing ops gets consented segments. Data science gets one activity history. Agents call the same resolve and verify services through one API or MCP endpoint, so a person and an agent working the same account see the same contacts, the same consent and the same priority.

Security and compliance. Your brand stays the data controller and iCustomer acts as processor. Agents see masked contact values by default and raw values only through scoped keys. If you ever

leave, your IDs and mappings stay in your warehouse. SOC 2 Type II is complete, with the report available under NDA, and ISO 27001 and ISO/IEC 42001 certification are in progress.

What rollout looks like. iCustomer Growth Engineer FDEs run enterprise deployments inside your data cloud, in three phases. Foundation: connect your systems, resolve IDs and agree the data contract. Activation: scores, buying centres and audiences flow to Salesforce, the marketing cloud and ad platforms. Optimization: monitoring, decay tuning and agent access. Your data stays your property, and your data team owns the schema and policy. The FDE does the build.

Your teams keep their tools. The identity, the signals and the learning stay in your first-party data.

If you want a starting point, we can run the five-number scorecard on a sample of your in-scope accounts and contacts and show you where coverage, depth and decay stand today.

Get your free data scorecard

Appendix: running the program

A 90-day plan for the data you already have

Enterprise teams don't start from zero. They start with years of accounts and contacts in Salesforce and the marketing cloud, and the job is to audit, validate, enrich and then keep it current. Don't try to fix the whole database at once. Scope it to the accounts that drive revenue, fix those properly, then widen.

Days 1 to 15: audit and scope.

- Pull every account and contact from Salesforce, the marketing cloud and active list sources into the warehouse, with source and last-modified date.
- Baseline the damage: duplicate rate, contacts with no activity in 12 months, bounce rate, accounts with no parent or hierarchy, missing industry and size.
- Inventory vendors: what each one writes, where, and what it costs.
- Pick the first scope: customers renewing in the next 12 months, named expansion targets and your top-tier prospect accounts.
- Draft the eight rules and get sign-off from RevOps, marketing ops, data science and the CIO office.

Days 16 to 45: validate and resolve.

- Resolve accounts and contacts to persistent IDs anchored on Salesforce account IDs. Merge duplicates in the warehouse and push merges back after RevOps reviews them.
- Build account hierarchies (parent, subsidiaries, regions, business units) so buying centres roll up correctly.
- Validate contacts in scope: verified email, still at the company, current title. Flag leavers. A champion who left a customer is a renewal risk and a new-logo lead at their next company.
- Test vendors against a ground-truth set of 500 known-good contacts on coverage, wrong answers and decay.
- Add provenance and last-verified dates to every field.

Days 46 to 75: enrich gaps and map buying centres.

- Define buying centres for each account type. A large customer can have several, one per business unit or product line, each with its own economic buyer, champion, users, IT or security, procurement and finance.
- Map every account in scope: which roles are filled, which are missing, and who are users (from product data) versus buyers.
- Enrich only the gaps. Fill missing roles from the vendor that scores best for that field, instead of re-enriching the whole database.
- Build a treasure map per account: buying centres, roles, coverage, engagement and whitespace.

Motion	What the treasure map must show	Owner
Renewal	Economic buyer and champion still in seat, user adoption, risk where a key contact left	Account manager, CS
Expansion	Buying centres not yet sold to, users in adjacent teams, new executives	Account executive
Prioritization	ICP fit, key roles covered, recent intent and engagement	SDR, ABM

Days 76 to 90: activate and set the upkeep.

- Push treasure maps, coverage scores and missing-role tasks into Salesforce as account fields, contact roles and rep tasks.
- Send verified, consented segments by buying role to the marketing cloud, and target audiences for accounts with missing roles to ad platforms.
- Point your first agent, research or outbound, at the identity layer.
- Turn on upkeep: reverification on a schedule set by each source's decay rate, job-change alerts on key roles, new signals attached to IDs as they arrive, and a monthly coverage review by account tier. Then widen the scope to the next tier of accounts.

Track the scorecard in the next section weekly against your day-one baseline.

Measure your data like an asset

Your contact and account data is an asset that loses value every month nobody maintains it. Measure it the way finance measures any asset: a short scorecard, a baseline, a target, an owner and a review cadence.

The health scorecard. Five numbers tell you whether the data is usable. Measure them on accounts in scope (customers, ICP accounts, named targets) and break them out by tier. A whole-database average hides the problem.

Metric	Definition	Better	Example baseline	Example target
Coverage	Share of in-scope accounts with at least one verified, reachable contact in the buying centre	Higher	60.2%	80%
Depth	Share of in-scope accounts with every required role filled in each buying centre	Higher	8.0%	30%
Profile completeness	Share of required fields populated and verified: email, title, role, seniority, department, hierarchy, industry, size	Higher	41.7%	75%
Duplicate rate	Share of records that resolve to an ID another record already holds	Lower	2.9%	1.5%
Stale contact rate	Share of contacts past their reverification window, or with a job change not yet reflected	Lower	63.9%	35%

Look at what the example says. Coverage of 60% sounds respectable. But only 8% of accounts are fully mapped, and almost two in three contacts are stale. A team that tracks coverage alone thinks it's in decent shape while its reps and agents work from a map that's mostly out of date.

Persistence and governance metrics. The scorecard shows the state of the data. These show whether the system keeps it healthy:

Metric	What it tells you	Owner
Resolution rate	Share of incoming records and signals matched to an existing persistent ID	Data team
Signal attach rate	Share of activity events tied to both a contact ID and an account ID	Data team
Job-change continuity	Share of detected job changes where the contact ID and its history carried over	RevOps

Metric	What it tells you	Owner
Decay rate, per source and field	How fast each vendor's values go stale	CIO office, data team
Correct-answer rate, per vendor	Accuracy against your ground-truth set	CIO office
Time to reverify	Days from a record going stale to being verified again	RevOps, marketing ops
Consent coverage	Share of contacts with a known consent and suppression status	Marketing ops
Cost per correct, usable record	Vendor spend divided by records that are verified and actually used	CIO office, finance
Activation health	Ad platform match rate, CRM sync success, bounce and spam rates on sends	Marketing ops, ABM

Cadence. Ops teams review the scorecard weekly. The governance group (RevOps, sales ops, marketing ops, data and the CIO office) reviews trends and vendor performance monthly. Leadership sees the five scorecard numbers against target every quarter, next to pipeline and renewal results, so data health sits beside the revenue it drives.

Tie it to revenue. Data health only matters if it moves the numbers revenue and growth leaders already watch. Report these next to the scorecard: pipeline created from in-scope accounts, win rate on multi-threaded versus single-threaded deals, renewal and expansion rates where the champion and economic buyer are still in seat, CAC and ROAS on audiences built from verified IDs, and ad spend lost to unmatched or out-of-ICP records.