P A R C E N O M I C S

Better data. Better decisions.

CASE STUDY · APRIL 2026

Donor Intelligence at Scale

How Parcenomics resolved a 9,518-record donor universe into 7,194 campaign-ready households for a local services organization — and what it found along the way.

9,518 → 7,194	Source records resolved to mailable households
25	Deceased-donor resolutions (21 auto-applied, 4 operator-review)
1,776	Data quality findings categorized for CRM cleanup
548	Duplicate records merged at the household level

Note: This case study is anonymized. The client is described generically as a "local services organization." All outcome numbers reflect the actual engagement and are reconcilable in the pipeline's audit trail.

01 · EXECUTIVE SUMMARY

A Single Appeal, Done Right

A local services organization engaged Parcenomics for a county-wide donor intelligence build ahead of its annual fundraising appeal. Like most small and mid-sized nonprofits, its donor information lived across multiple disconnected sources: a CRM with nearly two decades of giving history, multiple wealth-screen prospecting exports purchased at different times for different campaigns, and assessor databases covering every property in its service area.

Before each annual appeal, staff and board volunteers reconciled these sources by hand — a process that consumed weeks, produced duplicate sends, failed to catch deceased donors, and left prospects hiding behind LLC ownership structures entirely invisible. The cost of a single misdirected letter, in both postage and relationship damage, was not trivial. The cost of the prospects they never identified was likely an order of magnitude greater.

Parcenomics built a single analytical pipeline that ingested all source files, resolved identity across sources at the person and household level, cross-referenced the result against regional obituary data, and produced a tier-segmented, campaign-ready mail universe — together with a categorized data quality sidecar and a full audit trail. This paper documents what we delivered, in detail.

> *For the cost of one wasted direct mail send, we got a resolved, segmented, appeal-ready donor universe — and the data quality intelligence to keep it that way. The infrastructure Parcenomics built compounds in value for every nonprofit that follows in the same service area.*

02 · WHAT WE STARTED WITH

Twenty Files, Five Formats, Fourteen Years

The engagement began with twenty files delivered to Parcenomics:

- One CRM export containing approximately 2,000 historical donor records — current active donors, lapsed donors stretching back fourteen years, and a mix of organizational and individual contacts.
- Five county-area assessor databases covering the organization’s service area — every property owner, every mailing address, every ownership structure.
- Fourteen wealth-screen exports purchased at various points for various campaigns: high-net-worth residents by town, top-100 giving capacity by geography, affinity lists tied to regional educational and land-conservation organizations, out-of-state donors with local second homes, and property-owners of specific interest to the organization.

Every file had a different schema. Column names varied. Name conventions differed. Some files had merged couples in a single record; others had couples split across two rows. Some recorded "Elizabeth" where another recorded "Betsy" for the same person. Assessor data spelled city names in local abbreviations; the CRM sometimes did too. Mailing addresses included PO boxes, care-of routing to family members, apartment numbers with and without the word "Apt."

The total ingested record count was 9,518 rows across all sources. The development office’s goal was a mail universe of fewer than 8,000 unique households, segmented by relationship strength and ready to mail in the next appeal drop. Parcenomics delivered 7,194.

03 · DELIVERABLE ONE

The Campaign-Ready Universe

The primary deliverable is a single resolved, deduplicated, tier-segmented mail file that the development office can hand to a fulfillment house with zero cleanup. In this engagement, this deliverable consisted of five spreadsheets:

Mail File (primary)

7,194 rows. One row per mailable household. Columns for Letter_Tier (Current / Lapsed / Prospect), envelope-line name, salutation line, is_org flag, First/Last/Spouse_First/Spouse_Last/Company/Entity_Name, fully USPS-normalized address, household key for cross-reference, data source, and a salutation review flag for operator-check records.

Tier-Segmented Mail Files (three)

1,074 Current donors. 716 Lapsed donors. 5,404 Prospects. The development office sends different letter copy to each tier; the pre-segmented files let the fulfillment house print and mail in discrete batches without further sorting.

Email File

909 rows. One row per person with a valid, deliverable email address, deduplicated against gmail-dot-and-plus-tag canonicalization (so "john.smith@gmail.com" and "johnsmith+newsletter@gmail.com" resolve to a single contact). Tier-tagged, but kept in a single file so the digital-only appeal can go out from one send.

What Made the Universe "Campaign-Ready"

Three categories of cross-source work produced the final universe. Each one solves a problem the development office had been living with and working around:

Entity Resolution Across Sources

A single donor appeared in as many as nine different input rows. The pipeline resolved these into single household records using three strata of match confidence: (1) exact match on normalized email or wealth-screen profile URL; (2) strong match on (accent-stripped surname, nickname-expanded first name, ZIP5) triple; (3) fuzzy match on token-set ratio within the same ZIP block. The nickname engine covers 60+ common substitutions: Bob/Robert, Betsy/Elizabeth, Kate/Katherine, Dick/Richard, so the same person does not appear twice under two different forms of the same name. 548 records were identified as secondary contacts in a household already represented by a higher-tier primary — not deleted, but correctly suppressed from the mail file while remaining fully visible in the audit trail.

Couple Resolution with Woman-First Salutation

Mail convention in modern development communications places the female partner first on the salutation line when known. The pipeline applies this rule automatically: given sex-tagged records from wealth-screen sources (and a fallback dictionary of 400+ common U.S. given names categorized by sex), the salutation reorders to woman-first even when the envelope retains the deed or registration order. 985 couples in this engagement were reordered in this fashion. 1,356 same-last-name couples were collapsed to the compact form ("Dorothy & David Collari") rather than the verbose form ("Dorothy Collari & David Collari"), matching the style of personal correspondence rather than legal notice.

Entity Classification and Reclassification

Not every donor in a CRM is a person. Trusts, LLCs, and corporate contacts need formal envelope-line treatment and a "Dear Friend" salutation by default. The pipeline identifies these by structural keywords ("LLC", "Trust", "Inc.", "Foundation") and by common-term patterns ("Realty", "Associates", "Holdings"). 778 of the 7,194 resolved households were organizational contacts. Of these, 26 had been incorrectly entered as individual contacts in the CRM — an LLC’s registered agent entered as the donor, or a corporate donor name split awkwardly across person fields. The pipeline caught these structural mismatches and reclassified them, producing correct organizational envelopes without requiring the development office to locate and fix the source records.

04 · DELIVERABLE TWO

The Data Quality Report

Every nonprofit CRM contains silent failures — records that do not error loudly but cannot actually deliver mail, families where one spouse has died and the development office never learned, duplicates spawned by imports that were never reconciled. The Data Quality Report is not a diagnostic dump. It is a categorized cleanup list, prioritized by impact, with every entry tied to a specific record the development office can correct.

Address Completeness Findings

1,776 records could not be mailed as received. The pipeline categorizes each one by which address component was missing or malformed:

- 1,280 records missing ZIP5 — the most common single failure. Typically a data-entry artifact where a 4-digit ZIP lost its leading zero in an Excel import. The pipeline reports the source row so the data steward can correct it with full context.
- 474 records missing multiple fields (no street, no city, no state, no ZIP) — effectively orphan records that cannot route. These are the records that cost postage on every prior appeal without ever reaching a mailbox.
- 22 records with partial failures: missing city while other fields are complete, missing state alone, missing street alone. Each is typically correctable from the other present fields plus a ZIP lookup.

Deceased-Donor Resolutions

The most costly silent failure a small nonprofit makes is mailing a solicitation letter to a household where the named donor has died. The emotional harm to the surviving family is real. The reputational cost to the organization is also real — and unlike other data errors, it is never discoverable from within the CRM itself. The pipeline addresses this by cross-referencing the resolved household universe against a regional obituary registry — funeral home notices, regional newspaper obituaries, community death records — scoped to the organization’s service area.

Twenty-five deceased-donor resolutions were identified in this engagement. The pipeline assigns each a confidence score and takes action accordingly:

21 HIGH-Confidence Matches — Auto-Applied

These are donors whose obituary was published in the same town-group as their donor-record mailing address, with a matching first name and surname. The pipeline auto-applies the correction:

- If the deceased is the spouse, the spouse's first and last names are removed from the household record. The envelope and salutation are rebuilt to address the surviving partner alone.
- If the deceased is the primary donor and a surviving spouse exists, the spouse is promoted to primary. Mail continues; the relationship with the organization transitions cleanly.
- If the deceased is a sole donor with no surviving spouse on file, the record is marked do-not-mail. The development director has the opportunity to make a condolence outreach rather than send a solicitation letter that arrives after a death.
- If both spouses have matching obituaries, the household is marked do-not-mail with both obituary references preserved in the audit trail.

4 LOW-Confidence Matches — Operator Review

These are matches where the obituary town and the donor mailing town are in different town-groups within the region. They could be the correct person — a donor who moved towns after their most recent mail update — or a namesake, a different person with the same name living in a different part of the region. The pipeline does not auto-apply these. It surfaces them to the development director with full obituary context (name, date, age, city, source URL) so a human can confirm or dismiss before the next mail run.

Household Duplicate Resolution

548 records were resolved as secondary contacts within a household already represented by a higher-tier primary record. A typical example: a Current donor appears in the CRM at one address; a Lapsed donor appears at the same address as a separate record. They are not duplicates of each other — they are two different people sharing a household. The pipeline preserves both in the audit trail, but sends only one letter to the household (to the Current donor, per the tier-rank ordering), with the mail contents acknowledging both. This avoids the common small-nonprofit error of mailing duplicate appeals to the same address without the development office realizing the recipients are related.

Structural Name-Parse Review

Five records in the input had name structures the pipeline could parse but not with full confidence. Examples include business names entered with one part in First and the other in Spouse fields, when in fact they were a single organization. Hyphenated or compound family names split awkwardly across person fields. Orphan ampersands that suggest a lost first name. These are surfaced in a Structural Review file with the original source data intact, so the CRM can be corrected at the source rather than papered over in every downstream export.

05 · DELIVERABLE THREE

The Appeal Intelligence

The third deliverable is the one that converts the engagement from "list-cleaning" into "campaign strategy." It surfaces the opportunities hiding in the data that the development office could not have found by working from the CRM alone.

Tribute-Gift Opportunities

> *The obituary of a longtime supporter explicitly requested memorial donations to the organization in lieu of flowers. Our pipeline surfaced this through regional obituary cross-reference — catching a tribute-gift opportunity the CRM could not produce on its own, and a relationship-building moment with the surviving family that would otherwise have required a solicitation letter addressed to the deceased.*

Tribute gifts surfaced this way tend to have unusually high conversion rates. The donor’s family — often themselves supporters or candidates for future giving — has already been primed by the obituary to expect outreach. The appropriate response is not a solicitation letter but a condolence note with a tribute-giving option, and the pipeline provides the data to write exactly that note with full context.

Planned-Giving Candidates

Legacy gifts (bequests, charitable remainder trusts, life-insurance gifts) are the highest-value revenue most small nonprofits ever see and the lowest-cost to cultivate. But most small organizations cannot systematically identify candidates because the required data — age, giving depth, life circumstances, wealth indicators — lives across three different systems.

The pipeline composes candidate scores from four signals: donor age (from wealth-screen sources), total giving over tenure (from CRM), giving consistency (from CRM timestamps), and life-event signals such as recent widowhood (from obituary cross-reference). Candidates are surfaced ranked, not filtered — the planned-giving officer applies human judgment to the ranking. The organization received a ranked list of planned-giving candidates for the engagement, each with the supporting evidence that justifies outreach.

Service-Area Prospect Expansion

Small-nonprofit prospecting usually stops at the donors the CRM already knows about. But in a defined geographic service area — as this organization has — public assessor records reveal every property-owning household whether or not

they are in the CRM. Those households are strong prospects: they live in the service area, they own real estate (a proxy for both wealth and commitment to place), and their mailing addresses are already public record.

The complication is that many valuable properties are held in LLC or trust structures that mask the human owner. The pipeline surfaces the principals behind these structures by cross-referencing registered agent filings, parcel owner history, and care-of-name mail routing. In this engagement, this produced a universe of 5,404 qualified prospects — geographically proximate, property-owning households in the service area, with known or inferred human principals and mailable addresses — none of which had been in the existing CRM.

Lapsed-Donor Re-Engagement Signals

The CRM classifies 716 donors as "lapsed" — last gift more than two years ago. Lapsed lists are notoriously hard to work because many have moved, died, or simply aged out of direct mail. The pipeline cross-checks every lapsed donor against assessor records to confirm they still own property in the service area, and against obituary records to catch the deaths that would otherwise produce the classic small-nonprofit error of mailing condolence-free re-engagement letters to a widow. What remains is a lapsed list the development office can work with confidence: these are donors who are still here, still reachable, and whose lapse is a relationship problem, not a location problem.

06 · THE NUMBERS, IN FULL

Reconciled Outcomes

The engagement produced the following verified outcomes, reconcilable in the audit trail:

Source records ingested across 20 files	**9,518**
Resolved households delivered as campaign-ready	**7,194**
Current donor households	**1,074**
Lapsed donor households	**716**
Prospect households (service-area expansion)	**5,404**
Email-deliverable records	**909**
Couple households (resolved across sources)	**1,663**
Couples reordered woman-first on salutation	**985**
Same-last-name couple collapses	**1,356**
Organizational households (entities)	**778**
Individuals reclassified as organizations	**26**
Data quality findings requiring CRM cleanup	**1,776**
Household duplicates merged (same address)	**548**
Deceased-donor resolutions (HIGH auto-applied)	**21**
Deceased-donor matches (LOW operator review)	**4**
Structural name-parse review items	**5**

Every row in every deliverable is traceable back to its source record through a 9,518-row audit trail that the development office can query at any time. No record is lost; no decision is invisible.

07 · WHAT THE DEVELOPMENT OFFICE GOT BACK

Time, Accuracy, and Strategic Capacity

Time

The annual appeal-prep workload, previously measured in weeks of hand-reconciliation by the development director and a rotating cast of board volunteers, became a two-hour review of pre-reconciled deliverables. The development director spent her returned time on donor cultivation conversations — the work that only she can do and that only grows revenue when she has the capacity to do it.

Accuracy

Twenty-one families that would have received a solicitation letter addressed to a deceased spouse, did not. Four additional possible deceased matches were caught before mailing rather than after. 548 duplicate sends that would have cost postage and created household confusion were suppressed while the underlying people remained trackable. A longtime supporter’s tribute-gift opportunity was surfaced to the development director automatically, permitting a condolence note rather than a solicitation.

Strategic Capacity

The 5,404 service-area prospects — none previously in the CRM, each resolved to a human principal behind whatever LLC or trust structure held their property — represent a universe the organization can now work with. They are not anonymous strangers; they are mapped to address, tier, and relationship context. The development director now has a prospect pipeline that a board development committee can actually vote on.

> *Parcenomics did not replace the development director. It replaced the three weeks of data-wrangling that stood between her and her real job. The appeal went out on schedule, to a cleaner list than the organization has ever had. The prospect universe it produced will be worked for years.*

08 · HOW IT WORKS, IN BRIEF

The Pipeline Behind the Deliverables

The Parcenomics Donor Intelligence pipeline runs in eight sequential stages, each of which produces artifacts the next stage consumes:

1. Ingest

Source-specific adapters read each input file in its native format. Vendor-specific adapters for common CRM systems; wealth-screen adapters for the major prospecting services; an assessor adapter that handles both Excel and delimited-text files with the quirks specific to municipal CAMA systems.

2. Normalize

Field-level cleanup. Whitespace collapse, "nan" artifact removal, accent handling, title-casing with more than 50 exception rules (McDonald, O’Brien, Proctor’s, Mr. and Mrs., LLC and Inc., Roman numeral suffixes, professional designations, surname particles, acronym allowlists). USPS address normalization. Currency and date parsing.

3. Person Entity Detection

Classifies every normalized record as a person, an organization, or a crammed-couple pattern requiring split. Handles the case where a CRM put a first name into one field and a last name into another, when the record was actually a couple with different surnames.

4. Cross-Source Match

Three strata of identity resolution: exact (email or profile URL), strong (surname + nickname-expanded first + ZIP), fuzzy (token-set ratio within block). 60+ nickname variants. Thresholds tuned to avoid false merges.

5. Household Dedup

Records sharing a normalized street address and ZIP5 collapse to a household. Within a household, one record wins the primary slot based on a tier-rank ordering (Current > Lapsed > Prospect) with secondary tiebreakers on gift count and last gift date. Other records are preserved in audit but suppressed from mail.

6. Tier Assignment

CRM records with a recent gift classify as Current; older CRM records as Lapsed; non-CRM records as Prospect. Edge cases (CRM records with no gift date) default to Lapsed out of caution.

7. Obituary Cross-Reference

Regional obituary sources are scanned continuously (funeral homes, regional papers, community death notices). Each household is checked against the obituary registry using the same nickname-aware, town-group-aware match logic. Matches are scored HIGH, MEDIUM, or LOW and actions are taken accordingly.

8. Salutation and Output

Envelope lines and Dear-lines built with couple-ordering, same-last collapse, woman-first reordering, and organization-appropriate fallbacks. Nine Excel deliverables produced with internally consistent row counts and a full audit trail documenting every decision.

ABOUT PARCENOMICS

Parcenomics builds analytical infrastructure for small and mid-sized organizations that need enterprise-grade data work at grassroots budgets. Our nonprofit vertical serves organizations in the $500K to $15M annual revenue range that have outgrown doing data work by hand but cannot yet justify an enterprise CRM implementation. Our municipal vertical serves communities of all sizes on tax classification, assessment equity, and exemption policy analysis.

Every engagement is scoped to a specific deliverable with a specific timeline and a specific cost agreed before the work begins. If we do not see a path to meaningful return on the engagement fee, we say so.

Patrick White, Founder
patrickwhitestockbridge@gmail.com · 413-441-5231 · Stockbridge, MA