The Data 360 Credit Burn Kit

Identity Resolution costs ~50,000× more than a query. This is the full credit rate card, the 11 operations silently draining your credits, the SQL to cut identity-resolution cost, and a burn worksheet with real dollar math.

Rate card | Identity Resolution | Batch vs Streaming | SQL to cut cost | March 2026 model

THE 11 SILENT CREDIT-BURNERS

01 Editing an Identity Resolution ruleset

02 Resolving identity on junk profiles

03 Streaming when batch would do

04 Hourly Calculated Insights

05 Ingesting external what you could get native

06 Ingesting history no use case touches

07 Not pre-transforming upstream

08 Delta volume compounding

09 Over-frequent segment publishes

10 Streaming activation

11 Flying blind on consumption

Compiled by the Genetrix data engineering team · genetrix.tech

HOW THE CREDIT MODEL WORKS

Data 360 (formerly Data Cloud) bills most processing on **consumption credits**. Each operation has a multiplier applied per million rows processed — **every time it runs**, not just the first time.

```
Credits = (rows ÷ 1,000,000) × multiplier   |   ~$500 per 100,000 credits = $0.005 / credit
```

Two things make the bill bite: (1) there's **no friendly overage** — run out and operations get throttled or fail, so your "real-time" segment silently stops updating; (2) processes like Identity Resolution **re-trigger automatically** when you edit them. Monitor everything in the free **Salesforce Digital Wallet**.

THE MULTIPLIER RATE CARD — MOST EXPENSIVE FIRST

OPERATION	TYPE	CREDITS / 1M	$ AT LIST PRICE
Identity Resolution (Profile Unification)	Batch	100,000	$500 / 1M rows
Sub-second Real-Time events	—	70,000	$350 / 1M
External data ingestion	Streaming	5,000	$25 / 1M
Data Transforms	Streaming	5,000	$25 / 1M
Inferences (Einstein / BYOM)	—	3,500	$17.50 / 1M
External data ingestion	Batch	2,000	$10 / 1M
Activate DMO	Streaming	1,600	$8 / 1M
Real-Time Profile API	—	900	$4.50 / 1M
Calculated Insights	Streaming	800	$4 / 1M
Data Actions	Streaming	800	$4 / 1M
Data Transforms	Batch	400	$2 / 1M
Zero-Copy federation (rows accessed)	—	70	$0.35 / 1M
Segment rows processed	—	20	$0.10 / 1M
Calculated Insights	Batch	15	$0.075 / 1M
Activation	Batch	10	$0.05 / 1M
Data Queries (SQL)	—	2	$0.01 / 1M
Native Salesforce connector ingestion	Batch / Stream	0	Free

Multipliers per Salesforce's published Data Services and Segmentation & Activation rate cards (2025 rate sheets, carried into the March 2026 consumption model). Dollar figures use the $500 / 100,000-credit list price — **your contract's rate card and price-per-credit will differ**. Always confirm in your Digital Wallet and with your AE.

THE 11 SILENT CREDIT-BURNERS

01 Editing an Identity Resolution ruleset

BURNS CREDITS Every edit to an IR ruleset triggers a FULL re-run at 100,000 credits / 1M — not just the deltas.

CUT IT Finalize rules in a sandbox, batch your changes, and run IR on a schedule. Don't tweak-and-save on a live 5M-profile ruleset.

02 Resolving identity on junk profiles

BURNS CREDITS You pay $500 per 1M rows to unify duplicate, inactive, and un-contactable records that will never generate revenue.

CUT IT Filter to contactable, active profiles BEFORE Identity Resolution. Resolve clean data, not your whole raw lake.

03 Streaming when batch would do

BURNS CREDITS Streaming multiplies cost: ingestion 5,000 vs 2,000, transforms 5,000 vs 400, Calculated Insights 800 vs 15, activation 1,600 vs 10.

CUT IT Reserve streaming for genuine sub-minute use cases. Everything else runs batch — often 50×+ cheaper.

04 Hourly Calculated Insights

BURNS CREDITS Refreshing a CI hourly instead of daily runs it 24× as often for data nobody reads until tomorrow.

CUT IT Move non-time-sensitive CIs from hourly to daily — Salesforce's own guidance says this saves ~96% of CI credits.

05 Ingesting external what you could get native

BURNS CREDITS External pipeline costs 2,000 / 1M; the native Sales, Service, MCE, and Commerce connectors cost ZERO.

CUT IT Pull Salesforce data through native connectors, not an external pipeline. Same data, no ingestion credits.

06 Ingesting history no use case touches

BURNS CREDITS You pay to ingest, store, resolve, and re-process rows that no segment, CI, or activation ever uses.

CUT IT Scope ingestion to the use case and a sensible time window. If no use case needs it, don't bring it in.

07 Not pre-transforming upstream

BURNS CREDITS Running dedup and shaping inside Data Cloud (Transforms at 400–5,000 / 1M) that your warehouse could do for free.

CUT IT Clean, dedupe, and shape data in your existing warehouse before ingestion. Don't pay Data Cloud to fix data twice.

08 Delta volume compounding

BURNS CREDITS A 10% daily delta doesn't just hit ingestion — it re-runs IR, every related CI, and every segment that touches that DLO.

CUT IT Forecast delta across ALL downstream processes, not just ingestion. Small daily growth becomes a big monthly bill.

09 Over-frequent segment publishes

BURNS CREDITS Each segment publish costs 20 credits / 1M rows evaluated — and teams often schedule far more republishes than the use case needs.

CUT IT Size the segment first (a 2-credit query), then set the lowest publish cadence the campaign actually requires.

10 Streaming activation

BURNS CREDITS Activating a DMO via streaming costs 1,600 / 1M vs 10 / 1M for batch — a 160× premium for latency you may not need.

CUT IT Batch-activate wherever the downstream channel can tolerate a schedule. Stream only when minutes matter.

11 Flying blind on consumption

BURNS CREDITS With no monitoring, overage surfaces as throttled or failed jobs — your 'real-time' segment silently stops updating.

CUT IT Watch the free Salesforce Digital Wallet consumption cards, review multipliers by feature, and set renewal alerts.

SQL TO CUT YOUR IDENTITY-RESOLUTION BILL

> Identity Resolution is the single most expensive operation. The cheapest way to cut it is to resolve fewer, cleaner rows. Run these in the Data Cloud Query Editor / as Data Transforms before IR.

Find the junk profiles you're paying to resolve

```
-- Every row here costs 100,000 credits / 1M (= $500 / 1M) to unify.
-- Filter these out BEFORE Identity Resolution runs.
SELECT COUNT(*) AS wasted_profiles
FROM   Individual__dlm
WHERE  (Email__c IS NULL AND Phone__c IS NULL)     -- not contactable
   OR  LOWER(Email__c) LIKE '%@test.%'             -- test data
   OR  LOWER(Email__c) LIKE '%@example.com'
   OR  Is_Inactive__c = TRUE;                      -- dead records
```

Feed a CLEAN DLO into Identity Resolution

```
-- Data Transform (batch, 400/1M) that shapes a resolvable, clean DLO.
-- Resolve 2M clean rows instead of 5M raw = 300,000 credits saved per run.
SELECT Id__c, FirstName__c, LastName__c, Email__c, Phone__c
FROM   Individual_Raw__dll
WHERE  (Email__c IS NOT NULL OR Phone__c IS NOT NULL)
  AND  Is_Inactive__c = FALSE
  AND  LastActivityDate__c >= CURRENT_DATE - INTERVAL '24' MONTH
```

Size a segment before you schedule republishes

```
-- Each publish = 20 credits / 1M rows evaluated. Know the number first.
SELECT COUNT(*) AS segment_rows
FROM   UnifiedIndividual__dlm u
JOIN   Purchase__dlm p ON p.IndividualId__c = u.ssot__Id__c
WHERE  p.PurchaseDate__c >= CURRENT_DATE - INTERVAL '90' DAY
```

Catch duplicates that inflate every process

```
-- Duplicates inflate ingestion, IR, CIs AND segments - every run, forever.
SELECT Email__c, COUNT(*) AS dupes
FROM   Individual__dlm
WHERE  Email__c IS NOT NULL
GROUP  BY Email__c
```

```
HAVING COUNT(*) > 1
ORDER  BY dupes DESC
```

THE BURN WORKSHEET — REAL NUMBERS

Identity Resolution on 5M raw profiles

First full run: 5 × 100,000 = 500,000 credits = **$2,500**. Then a 10% daily delta (500k rows): 50,000 credits/day = $250/day ≈ **$91,000/year** — just for IR.

...the same job, cleaned to 2M contactable first

First run: 200,000 credits = $1,000. Delta 200k/day = 20,000 credits = $100/day ≈ **$36,500/year**. Saved: ~$54,000/year from one WHERE clause.

A 20M-row Calculated Insight: hourly streaming vs daily batch

Hourly streaming: 20 × 800 × 24 = 384,000 credits/day = **$1,920/day**. Daily batch: 20 × 15 = 300 credits/day = **$1.50/day**. Same insight, ~99.9% cheaper.

Ingesting 50M Sales Cloud rows

Native connector: **0 credits**. The same data via an external pipeline (batch): 50 × 2,000 = 100,000 credits = **$500/run**. Use the native connector.

Data 360 bill climbing faster than the value?

That's usually architecture, not usage. Genetrix audits your data model, Identity Resolution, and refresh cadences to cut credit burn — then keeps the pipeline lean across Data 360, Marketing Cloud, and Agentforce.

Book a 30-minute Data 360 cost review →