

SchemaShift

Flat files in. A real relational database out. Every run on the record.

V0.4.0 · WINDOWS & LINUX · GPG-SIGNED RELEASES

SchemaShift is a desktop application for Windows and Linux that reverse-engineers legacy flat files — CSV, Excel, pipe-delimited, and fixed-width extracts — into fully relational PostgreSQL schemas. It infers tables, keys, and foreign-key relationships from the data itself, shows you an editable plan before a single line of DDL is written, and records every run and every emitted script to a tamper-evident, hash-chained audit log. Deterministic and rule-based — runs entirely on your hardware, no cloud, no telemetry.

WHAT IT DOES

Schema inference from the data itself

Point SchemaShift at a flat file and it detects entities, dimensions, and relationships — a single wide extract becomes a set of properly normalized tables with primary keys assigned.

An editable plan before generation

Every inference lands in a reviewable schema plan first. Rename tables, correct entities, and approve the design before any DDL is generated — the operator catches bad guesses, not production.

Composable, automation-ready outputs

DDL, load scripts, and direct execution are independent steps you combine as needed — from a reviewed script handoff to fully unattended, zero-prompt migrations. GUI and CLI in one install.

Foreign keys, proven before enforced

Cross-file relationships are auto-detected by matching columns to existing keys, then verified with a zero-orphan check before any constraint is enforced. Explicit manual links fail loudly on bad data.

Quarantine and reconciliation

Malformed rows are never silently skipped — or worse, silently mutated. They're preserved verbatim in a quarantine sidecar with a count that must balance: rows in equals rows staged plus rows quarantined.

Tamper-evident audit chain

Every run and every emitted script is logged to a hash-chained audit database. A broken or missing chain writes a permanent discontinuity record, and archived logs are GPG-signed and re-verified.

WHY SCHEMASHIFT

Built for real production data

Streams million-row files in constant memory and has been proven on multi-file bank core extracts — hundreds of thousands of rows landed into an enforced relational schema.

The operator stays in charge

No content filtering, no redaction, no guessing. Empty means NULL but “NA” is data, account numbers are text rather than math, inferred constraints are suggestions you approve — and anything ambiguous stops instead of assuming.

Sibling to DataShift

Same signing key, same audit discipline, same house. DataShift makes legacy data examiner-ready; SchemaShift gives that data a real relational home to live in.

WHO IT'S FOR

COMMUNITY BANKS


Leaving a legacy core? Land your extracts in a modern, queryable SQL database without shipping data to anyone's cloud.

MODERNIZATION & CONVERSION TEAMS

Repeatable, scripted migrations with an audit trail — the same input and plan produce the same schema, every run.

ANALYSTS AND OPERATORS

Turn a mainframe dump into a browsable relational schema in minutes — then query it with any standard SQL tool.



The Analyzer: point at a flat file, review the inferred schema plan, then emit DDL and load scripts — or execute straight into PostgreSQL.

Give your flat files a real home.

Request a demo and watch your own legacy extract become a fully relational database — keys enforced, rows reconciled, on the record.

info@cobblestone-digital.com
cobblestone-digital.com/schemashift/