

Grid Data Simulator

A whole utility from a polygon — delivered in your MDMS's own data model, at millions-of-meters scale.

Grid Data Simulator (GDS) builds the utility your systems need to be tested against. Draw a territory on the map — or search a place, or describe it in plain English — and it generates a complete synthetic utility: electric, gas, and water networks from substation to meter, premises and service points on real addresses, partners, contracts and service agreements, meters and their channels, and interval, register, and SCADA telemetry. Then it delivers that utility the way the real estate would — master data in the Grid Data Management Platform's own entity model, reads on real cadences, CIS lifecycle transactions, outages, and a full command-and-control round trip — to each system on its own Kafka cluster. Every scenario is reproducible from its spec and seed.

Millions of meters

designed and proven at scale

Broker-speed

parallel publishing to Kafka

GDM-native

entities · relationships · history

Multi-commodity

electric · gas · water

KEY FEATURES

A utility from a polygon

Draw territories, search a place or ZIP code, or describe the scenario in plain English and review the spec Claude drafts. Premises land on real buildings, assets on real roads — urban, suburban, or rural density.

GDM's data model, natively

Premise, location, device, channel, partner, and contract entities with the full set of relationship families — device-channel, partner-contract, contract-location, supplier and replacement links among them. Every attribute is dated, and every change carries its history.

Integration-ready master data

Deliver per entity or as one combined message per service point; target GDM or SAP, with profile sync and stable measurement tasks for SAP; stamp every envelope with the source system a CI/CD pipeline expects.

Several brokers, one routing table

Meter master data and reads to GDM, asset master data and SCADA to Enhanced Analytics, commands to the MDMS — each topic routed to its own Kafka cluster, with a producer per broker opened only when needed.

Telemetry feeds & historical loads

AMI cohorts on staggered cadences with their own channels, lag, and message span; SCADA on its own clock; meters silent under outages. Historical backfills of any length run as monthly job chains.

Realism on demand

Storms restored by finite crews with ETRs and nested outages; DER and EV adoption curves; VEE fault patterns with an expected-findings manifest; and dirty data — duplicates, drift, late and out-of-order — injected at the broker edge.

Actions & lifecycle events

Change one record and publish exactly the transaction a CIS would send: meter exchange, reversal, removal and install, move-out and move-in, attribute changes — each with the reads that follow, in a replayable ledger.

Command & Control

The full CIS → MDMS → HES round trip — disconnect, reconnect, on-demand read, power status, ping — now, at a time, or recurring, with the simulator on both ends and the MDMS under test in the middle.

Operable at scale

Partition manifests, bucket-aligned band files, monthly job chains, and per-band pause and cancel; read leases let publishes run side by side; memory limits spill to disk instead of failing.

Deploys like the family

Docker Compose, or Kubernetes with Helm as a managed application on the Platform Manager's shared Postgres, Keycloak, Kafka, and Redis — built, imported, and upgraded through ControlPlane from GitHub Actions.

PERFORMANCE & SCALE

Engineered for utility scale	Reviewed against the largest utilities and a full year of fine-grained, multi-channel data, re-engineered for it, then proven end to end — from an empty territory to a whole utility on Kafka in minutes.
Publishing	Reads are assembled inside the database and published by parallel processes — many times faster than before, and scaling with every worker you add.
Generation	A rewritten generator produces interval data far faster with markedly less memory; banded, spill-to-disk processing keeps even the largest days within a worker's budget.
Real time & C&C	A real-time feed for every meter keeps pace with its cadence with ample headroom, and large command batches close end to end through a real MDMS.

Measured through the product's own jobs, scheduler, and consumers on a single workstation — no special hardware, no shortcuts.

USE CASES

- **MDMS acceptance & VEE certification** — every rule fires on exactly the meters it should — the manifest says which.
- **GDM and SAP integration testing** — master data, profile sync, and reads in the exact entity model.
- **Load and scale proving** — millions of meters, published at broker speed, before go-live.
- **CIS lifecycle & outage operations** — exchanges, move-ins, storms with crews and ETRs.
- **Command-and-control integration** — a real MDMS's flows — no head-end, no CIS.
- **CI/CD test environments** — deployed by ControlPlane, stamped with each pipeline's source system.

TECHNOLOGY

Angular + MapLibre GL · Python / FastAPI · DuckDB (spatial) · Postgres control plane + day-partitioned Parquet · Arq + Redis workers · Kafka, multi-broker · Claude (NL → spec) · Keycloak · OpenAPI · Docker Compose / Helm

BETTER TOGETHER

GDS is the family's proving ground: it speaks the Grid Data Management Platform's own entity model, routes SCADA to Enhanced Analytics on its own broker, gives the Testing Platform environments stamped with the source system each pipeline expects, deploys through ControlPlane, and registers with the Grid Data Platform Manager.