DemandFlow Workbooks

DEMANDFLOW®

SQL, grids, charts and plain-English questions over your live network records, in the browser, with governed edits back

DATA SHEET

The questions a network team asks every week rarely fit a fixed report. Which circuits will be hot by Friday? Which diverse pairs share a duct? Which closures are nearly full in the north region? The answer usually means an export to Excel, a pivot, a copy to someone's inbox and a number that is wrong by the time it is read.

DemandFlow® Workbooks put a SQL engine in the browser against every record the platform holds. Every entity is a table. A workbook is a set of cells, each a query, a grid, a chart, a KPI tile, an input or a question in plain English, that re-run when the cells they depend on change. Results stay in the tenant, edits go back through a reviewed changeset, and the workbook can be scheduled, shared, presented or exported.

What it is

A notebook-style analysis tool for operators: DuckDB SQL in the browser and on the server over the tenant's records and uploaded datasets; reactive cells; typed grids with statistics and regression; charts; KPI tiles with thresholds; parameters as inputs; an Ask cell that writes the SQL for you; governed edits back to records through changesets; sharing with roles; versions and comments; schedules with run history; live feeds; presentations and redaction; and twenty templates covering capacity, shared risk, faults, fibre, RAN, lifecycle, compliance, cost and emissions.

It sits inside DemandFlow, so the tables are the sites, links, cables, cells, platforms, tickets and budgets the rest of the platform maintains, and an edit made in a grid becomes a record change with a reason and an approver.

What a team gets

- **Real SQL, no proprietary language.** Joins, aggregates, window functions, dates, string aggregation and spatial functions on geometry, with every entity as a table and option labels as columns.
- **Cells that re-run themselves.** A dependency graph from the SQL parser marks dependants stale and re-runs only those, four at a time, with cycles refused.
- **Ask the data.** A question in plain English becomes SQL with an explanation and a written answer; insert it as a cell when it is right.
- **Grids and charts that mean something.** Bandwidth, duration, percentage, currency, address and status columns typed; eight chart types that follow a grid's filters.
- **Datasets alongside records.** CSV, Excel, JSON, Parquet, GeoJSON, KML and KMZ uploaded, versioned and queried beside the tenant's tables.
- **Edits that are governed.** Change a grid, see the diff, submit a changeset with a reason, and a second approver above 500 records.
- **Scheduled, shared, presented.** Daily to cron schedules with run history and notifications; viewer, runner, editor and owner roles; presentations and redaction for a wall or a customer.
- **Server scale when you need it.** Run on the server for two million rows per entity and a million rows per result, user-initiated.

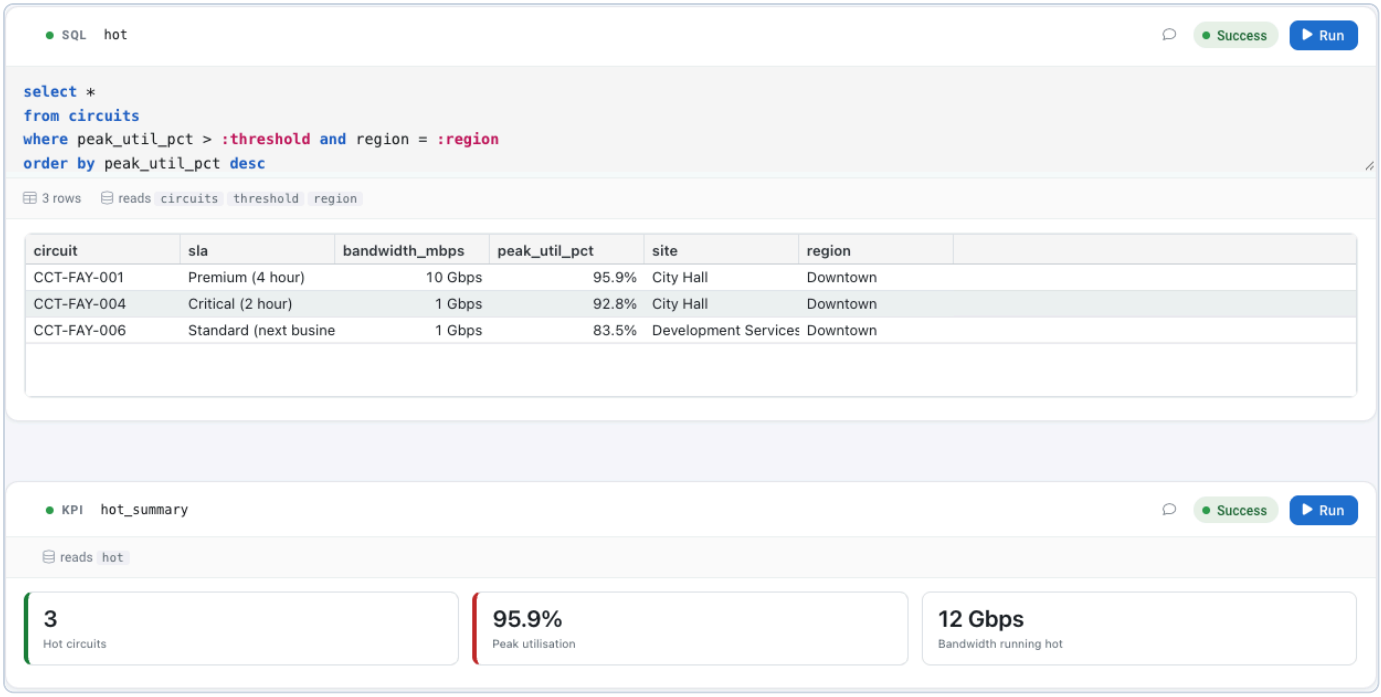

FIGURE 1. CIRCUITS ABOVE THE THRESHOLD IN THE CHOSEN REGION, WITH KPI TILES THAT UPDATE WHEN EITHER INPUT CHANGES

Cells and the engine

Cell types are text, input, SQL, grid, chart, KPI, Ask, router, JavaScript and Python. Inputs (select, number, text, date) bind into SQL as named parameters and cells read each other by name. The engine is DuckDB: DuckDB-WASM in the browser and the same engine on the server, with every tenant entity exposed as a table by its code, option fields as label columns and lookups as name columns. Spatial functions run on geometry loaded from GeoJSON, KML or KMZ. A no-SQL query builder covers single-table filters, ordering and limits for people who do not write SQL.

Cells move through not run, out of date, queued, running, success, error and waiting; a failed cell blocks its dependants. KPI tiles compute count, distinct count, sum, average, minimum, maximum, median or an expression, with warning and critical thresholds, a target, a baseline and the delta. Grids are Lattice grids with typed columns (bandwidth in Mbps or Gbps, duration, percentage, GBP, IPv4, dates, status dots, geometry), statistics and regression panels and paging pushed down to the engine. Charts are bar, stacked bar, horizontal bar, line, area, scatter, pie and heat map, and follow the filters of the grid they read.

On the server. Run on server and Open full table send the same SQL to a sandboxed DuckDB behind a Lambda: a single SELECT, locked extensions (Parquet, JSON, spatial, Excel), up to 6 GB and 900 seconds, two million rows per entity, one million per result in pages of 10,000, with results that expire after 24 hours. The tenant is taken from the token, never the request body, and results are visible only to the person who ran them.

Code cells. JavaScript and Python (Pyodide with NumPy and pandas) run in the browser only, in a sandboxed frame with a 30-second timeout.

Live feeds. A router cell subscribes to a WebSocket, SQS or Kinesis feed through a relay and keeps a windowed live table refreshing every second; the live fault triage template drives a NOC wall from it.

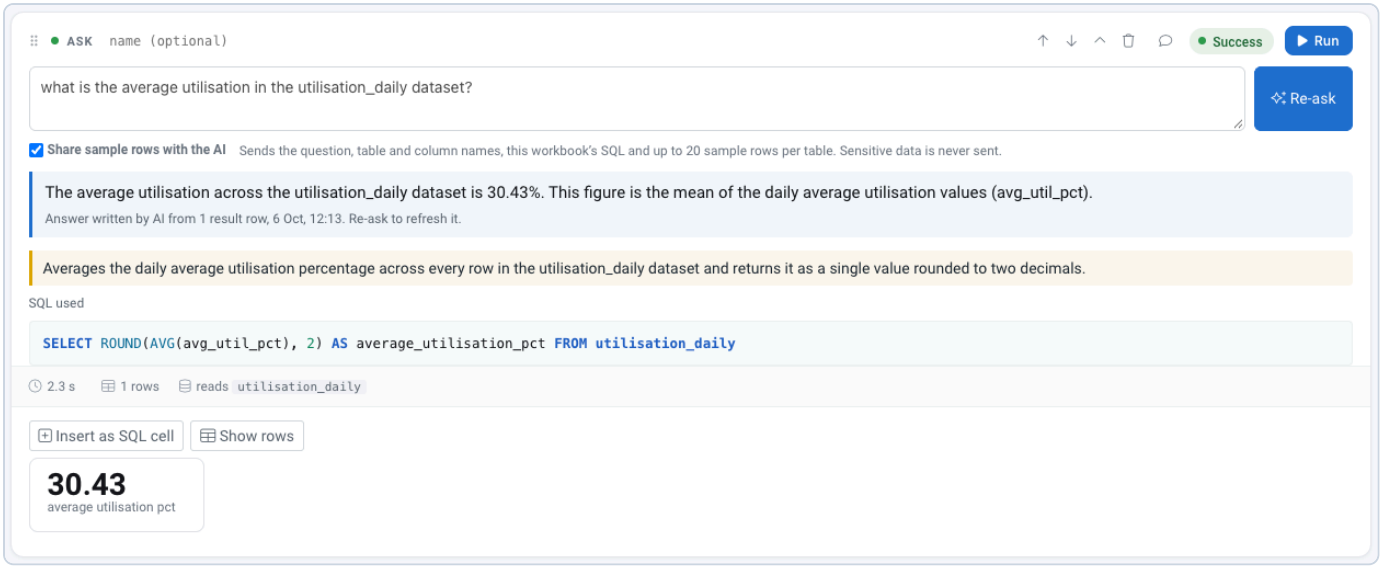

FIGURE 2. A PLAIN-ENGLISH QUESTION, THE WRITTEN ANSWER, WHAT THE QUERY DOES AND THE SQL BEHIND IT

Ask, datasets and governed edits

Ask. A question goes to Claude with the schema, option labels, the workbook's existing SQL and optionally twenty sample rows from up to five tables. It returns DuckDB SQL with an explanation, the SQL runs, and a plain-English answer is written from up to 200 result rows, with up to two automatic repairs if the SQL fails. The SQL is always shown and can be inserted as a cell. A grid's "ask your data" bar proposes a filter or sort and waits for you to apply it. Sensitive datasets and columns, and anything built on them, are removed from what the model sees. The feature is off by default per tenant and can be switched off per workbook; you can bring your own Anthropic key, held under KMS with only the last four characters shown; and every call is logged with model, tokens, latency and tables.

Datasets. Upload CSV, TSV, Excel, JSON, NDJSON, Parquet, GeoJSON, KML or KMZ (2 GB per file; Excel and spatial 500 MB). Files are converted to Parquet in the browser with a thousand-row preview to rename columns, set types, pick sheets and headers and mark sensitive columns. Datasets are versioned; a workbook pins a version or follows the latest. Promote to records inserts, updates by key or upserts up to 50,000 rows through a changeset.

Governed edits. A grid is editable when each row is one record. Edits are validated, staged and highlighted, with undo, redo and discard and a preview of downstream effects. Submitting creates a changeset with a reason and an approver; a record whose version has moved is refused individually; more than 500 records needs a second approver; the editor role is the minimum.

Sharing and history. Roles are viewer, runner, editor and owner. Workbooks autosave; Save bumps the version, and version history offers compare and restore with a confirmation. Comments attach to cells, rows and columns, and an activity stream records who did what.

Schedules, presentations, exports. Schedules run daily, on weekdays, weekly, monthly or on a cron expression in a chosen time zone, with inputs per schedule, snapshots in the run history and email or in-app notification. Present steps through cells with a dwell time and refresh (at least 10 seconds). Redact masks flagged columns and sensitive datasets everywhere, including the PDF. Exports are CSV, Excel (visible rows), PDF (60 pages, five per grid) and the workbook file; imports accept workbook files and Jupyter notebooks (5 MB, 500 cells). Export is restricted by role.



FIGURE 3. THE CAPACITY HOT SPOTS TEMPLATE CHARTING THE HOTTEST CIRCUITS BY PEAK UTILISATION

Templates

Twenty templates ship across the modules. Network: capacity hot spots, shared risk, fault triage, live fault triage, shared risk by strand. Fibre: fibre utilisation, closure and chamber capacity, ODF port utilisation, splice and test loss outliers. RAN: worst cells, congested sectors, availability by region and vendor, technology and band mix, spectrum licence expiries, EMF compliance. Lifecycle and compliance: end-of-support exposure, control evidence status. Finance and ESG: cost model variance, budget versus actuals, emissions by site. Each opens with its inputs set and its SQL visible, so it can be changed rather than rebuilt.

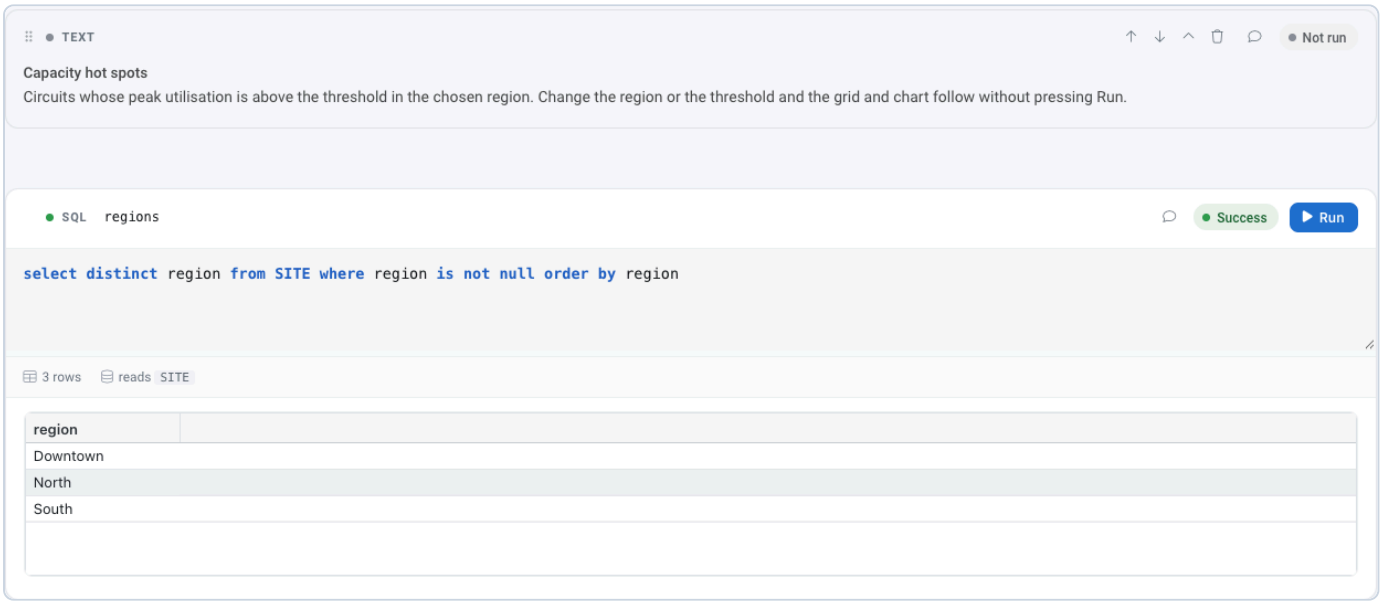

FIGURE 4. A SQL CELL READING THE LIVE SITE RECORDS, WITH THE RESULT GRID BENEATH IT

Summary of specifications

Item	Detail
Engine	DuckDB 1.4 in the browser (WASM) and on the server; every entity a table; option labels and lookup names as columns; spatial functions on geometry; no-SQL query builder
Cells	Text, input, SQL, grid, chart, KPI, Ask, router, JavaScript, Python; inputs select, number, text, date bound as parameters; reactive re-run of dependants, 4 in parallel; cycles refused

Item	Detail
Grids and charts	Typed columns (bandwidth, duration, percent, GBP, IPv4, date, status, geometry); statistics and regression; pushdown paging; bar, stacked bar, horizontal bar, line, area, scatter, pie, heat map following grid filters
KPI tiles	count, distinct count, sum, average, min, max, median, expression; warning and critical thresholds; target, baseline, delta
Server runs	User-initiated; single SELECT; extensions Parquet, JSON, spatial, Excel; 6,144 MB; 900 s; 2M rows per entity; 1M rows per result; 10,000-row pages; results expire after 24 h
Datasets	CSV, TSV, Excel, JSON, NDJSON, Parquet, GeoJSON, KML, KMZ; 2 GB per file (Excel and spatial 500 MB); 1,000-row preview; versions; promote to records up to 50,000 rows via changeset
Governed edits	One row per record; validated and staged; undo, redo, discard; downstream preview; changeset with reason and approver; per-record version conflicts refused; second approver above 500 records
Ask	Claude writes DuckDB SQL with explanation; up to 2 repairs; 200 rows for the answer; schema, labels, cell SQL and optional 20 sample rows from 5 tables as context; sensitive data excluded; off by default; own key under KMS; every call logged
Sharing	Viewer, runner, editor, owner; autosave; versions with compare and restore; comments on cells, rows, columns; activity stream
Schedules and feeds	Daily, weekdays, weekly, monthly, cron with time zone and per-schedule inputs; run history with snapshots; email and in-app notify; live feeds over WebSocket, SQS, Kinesis at 1 s refresh
Present and export	Presentations with dwell and refresh (10 s minimum); redaction of flagged columns and sensitive datasets; CSV, Excel, PDF (60 pages, 5 per grid), workbook file; import workbook and Jupyter (5 MB, 500 cells)
Code cells	JavaScript and Python (Pyodide, NumPy, pandas) in a sandboxed browser frame; 30 s timeout

www.demandflow.com

Book a demo: www.demandflow.com/workbooks/

DemandFlow® is a registered trademark of TOCLOCO LTD. Registered in England No. 09950859. 71-75 Shelton St, London, WC2H 9JQ. © TOCLOCO LTD. All rights reserved.