Data Guide

What is in the download files, which columns matter, which basis to use, how to add pixels up to your own areas, and what to watch for in this release.

Describes release 2026.10.4 Updated October 29, 2026

This guide describes the files of release 2026.10.4 and how to use them. For how the numbers are made, read the Methods page. Files are on the Downloads page or directly at

```
https://data.demandgrid.net/releases/2026.10.4/exports/<cc>/<file>
```

where `<cc>` is the country code: `gb` Great Britain (with Northern Ireland for 2021 to 2023), `fr` metropolitan France, `gp` Guadeloupe, `mq` Martinique, `gf` French Guiana, `re` Réunion, `yt` Mayotte, `ca` Canada, `kr` South Korea, `au` Australia. Years are 2021 to 2024 everywhere except Australia, which covers financial years 2017 to 2025.

The files

Each country has the same set. Sizes are per file: Great Britain's pixel table is about 33 MB a year, France's about 70 MB, Canada's and Australia's larger still.

File	What it is for	Format
`<cc>_electricity_pixels_<year>.parquet`	The main table, one row per 1 km cell: consumption by sector on both bases, provenance columns and the cell polygon. Use it for joins, filters and sums. One per year.	GeoParquet, 33 MB and up
`<cc>_electricity_gross_<year>.tif` and `<cc>_electricity_net_<year>.tif`	The same values as a four-band raster, one file per basis and year. Use these in GIS and raster workflows.	GeoTIFF, about 2 MB
`<cc>_electricity_gross_ratios_<year>.tif` and `<cc>_electricity_net_ratios_<year>.tif`	The per-person and per-built-area values as a raster, one band per column, one file per basis and year.	GeoTIFF, about 3 MB
`<cc>_facility_points.parquet`	One row per large facility and year: the site estimate, its range, status, sources, and what the raster carries for it. All years in one file.	GeoParquet
`<cc>_region_metadata.parquet` and `.csv`	One row per official region and year: what the region's totals follow, with citations, and the region polygon (parquet only). All years in one file.	GeoParquet, CSV
`<cc>_export_metadata.json`	Definitions, units, the facility-layer summary, license, citation and credit lines.	JSON

Hourly profiles exist for Great Britain on the map only; they are not published as downloads yet.

Units, sectors and bases

Every consumption value is **GWh per year per 1 km cell**, which is numerically the same as GWh per km². Values are consumption at the customer's meter: network losses are not included. `year` is the calendar year, except in Australia, where it is the financial year ending 30 June: 2025 means July 2024 to June 2025.

Three sectors, two bases, and one identity that holds in every cell:

```
elec_tot_gwh = elec_res_gwh + elec_non_res_gwh + elec_facility_gwh   (gross)
net_tot_gwh  = net_res_gwh  + net_non_res_gwh  + net_facility_gwh    (net)
```

`elec_` columns are **gross**; `net_` columns are **net**. Net is what the grid delivers to the premises. Gross is net plus the self-supply we can estimate: rooftop solar used on site where a rooftop-solar model exists, and a large site's own generation where known. Residential is homes; non-residential is everything else that is not a large facility; large facilities are the individually identified sites of 10 GWh a year or more.

Which basis to use

Your question	Use	Why
Compare with metered or network data; an emissions baseline from grid electricity; what the network carries	Net	Net is what the grid delivers.
Electrification planning; energy-use statistics; compare with a national energy balance	Gross	Gross is all the use we can account for, however it was supplied. Where there is no rooftop-solar model it equals net for homes and buildings.
Feed a grid or capacity-expansion model	Net if the model has no behind-the-meter generation; gross if it already represents rooftop solar and on-site plants as generators	Using net alongside those generators counts their output twice.
Estimate rooftop solar used on site	Gross minus net, residential and non-residential only	Only where a rooftop-solar model exists (Great Britain, Australia, Ontario). It excludes what the panels export, has no error band, and is not validated.

Where a country has no rooftop-solar model (`no_btm_model = true`: France, its overseas territories and South Korea), gross equals net for homes and buildings. Canada's model covers Ontario only: `no_btm_model` is `false` for the whole country, so check `btm_covered` in the region metadata file before using gross minus net there. Large facilities can differ on the two bases anywhere, by the electricity they generate for themselves.

The pixel table

One row per cell and year. GeoParquet 1.1 in longitude and latitude (EPSG:4326) with a covering bounding box, so a reader can fetch just the rows in an area.

Column	Meaning
`pixel_id`	The cell's index on the global 1 km equal-area grid, identical in every country, year and release: the key for joining across years. See "Cell ids" below.
`year`	Calendar year, or the financial year ending 30 June in Australia.
`region_uid`	The official region the cell belongs to. Joins to the region metadata file (Northern Ireland's blank 2024 rows carry a region with no 2024 row there).
`elec_res_gwh`, `elec_non_res_gwh`, `elec_facility_gwh`, `elec_tot_gwh`	**Gross** consumption by sector and the total.
`net_res_gwh`, `net_non_res_gwh`, `net_facility_gwh`, `net_tot_gwh`	**Net** consumption by sector and the total.
`anchor_case`	The case code: what training data and which official statistics stand behind the cell (see the case-code appendix of the Methods page).
`res_anchor_basis`, `non_res_anchor_basis`	The basis of the official figure the gross layer was scaled to: `gross` (the metered figure plus the rooftop-solar estimate), `approx_gross` (the official figure already counts some self-supply the grid never sees, so with the rooftop-solar estimate added the gross layer is approximately gross: Ontario, Australia), or `net_no_btm` (no rooftop-solar model, so gross equals net).
`facility_basis`	`grid_import` where every facility on the cell has a known own-generation split; `self_gen_unknown` where at least one does not and its net equals gross; `gross` where the cell holds no facility, or only facilities known to use no own generation.
`level_source`	The official statistics behind the cell's region, as source ids. Citations are in the region metadata file.
`no_btm_model`	`true` where the country has no rooftop-solar model (BTM: behind the meter, rooftop solar used on site). Per country, not per region.
`training_tier`	What local data trained the model where the cell lies: `admin_label` (the model learned from its official region's own figures: every cell in Great Britain, France and the overseas territories, Ontario's postal areas, the Australian council areas with billing), `unit_label` (a smaller published training unit only: municipalities in British Columbia and Quebec, postal areas in Queensland and Victoria, Ausgrid suburbs), or `none` (no local data, so the model trained elsewhere draws the pattern: South Korea, the rest of Canada and Australia). Both label values count as trained. Blank outside the published area.
`elec_tot_kwh_per_person`, `elec_res_kwh_per_person`, `elec_tot_kwh_per_built_m2`, `elec_res_kwh_per_res_built_m2`, `elec_non_res_kwh_per_nres_built_m2`, and the five `net_` twins	The map's per-person and per-built-area layers, value for value: consumption × 10^6 divided by the cell's residents, or by its built-up surface in m² (total, residential or non-residential). Blank where the cell has fewer than 5 residents or less than 100 m² of the relevant built surface (never infinity); 0 where consumption is 0. These are per-cell ratios: never sum them, and do not average them across cells. For an area's per-person or per-m² figure, sum the GWh columns and divide by your own population or floor area. The total ratios include large facilities, so a cell holding one reads very high. The population and built-surface counts themselves (GHSL) are not published; the country metadata JSON's `derived_metrics` names them.
`geometry`	The cell polygon in longitude and latitude. Cells are cut from an equal-area grid, so each is exactly 1 km² and looks slightly skewed in longitude and latitude.

Blank is not zero. The table holds only the cells inside a covered region (there are no rows for the sea). A value is blank (null) where the region has

no official figure: Northern Ireland in 2024, and a few French communes outside any intercommunality in every year. A value of exactly zero means no modeled consumption; the map draws it as transparent.

Cell ids. `pixel_id = row × 36,082 + column` on the global 1 km Mollweide grid (ESRI:54009), with `column = (x − (−18,041,000) − 500) / 1,000` and `row = (9,021,000 − y − 500) / 1,000` for the cell center (x, y) in meters. You never need this to use the files, since the geometry column carries each cell, but it lets you compute ids from your own grid.

The rasters

The GeoTIFFs stay on the equal-area grid (ESRI:54009, Mollweide) at 1,000 m, float32, no-data as NaN. Each file holds four bands in this order:

1. residential
2. non-residential
3. large facilities
4. total (the sum of the first three)

The gross and net files have the same layout. Band values are GWh per cell, the same numbers as the pixel table to float32 precision. The ratio files carry the five per-person and per-built-area columns of one basis as float32 bands, in this order: total per person, residential per person, total per built m², residential per residential built m², non-residential per non-residential built m². Each band's `name` metadata tag is the column name. No file mixes GWh with kWh per person or per m². Reproject your own layers onto the raster rather than the raster onto them, so each cell keeps its value and its identity.

The large-facility points

One row per facility and year, as points in longitude and latitude. The columns to read first:

Column	Meaning
`facility_id`	Permanent id of the site across years and releases. Join on it, not on the name.
`name`, `operator`, `facility_class`, `operating_status`, `status_from`, `status_to`	Who and what the site is, and the years its status applies. Names are unique within a country (a site that shares its name with another carries its town in brackets). A closed site carries zero in the years after `status_to`.
`estimate_status`	`estimated` for sites with an electricity estimate; `no_estimate` for large facilities we can locate and classify but cannot size yet (estimate columns blank, mapped columns 0, nothing in the raster).
`elec_facility_gwh` with `_low` and `_high`	The site's gross estimate and its range, the spread of the figures behind it; the range is blank where the source gives a single figure.
`net_facility_gwh` with `_low` and `_high`	The site's grid import. Equals gross where `facility_basis = self_gen_unknown`.
`onsite_generation_gwh`, `grid_export_gwh`	Own generation used on site, and what the site's plant exports to the grid. Exports are not subtracted from net. Blank means unknown, not zero.
`mapped_elec_facility_gwh`, `mapped_net_facility_gwh`	What the raster carries for the site. The sites inside a region's official statistics together never take more than half of that region's official non-residential figure in the raster, so where they would, each is scaled down and this is lower than the estimate.
`scaled_to_fit`	`true` where that scaling leaves the raster with less than half of the site's estimate. Sites scaled down by less are not flagged; compare the mapped columns with the estimate to see every case. A flag beside `confidence`, which describes the estimate itself.
`confidence`	`high`, `medium`, `low` or `unverified`. Unverified means a modeled estimate with no graded published figure behind it.
`frame`	How the official statistics count the site: `in_frame` (inside the regional figures), `label_gap` (inside the state or national total but outside the regional figures), `out_of_frame` (outside every official total), or `partial` (part inside, part outside).
`method`, `data_lead_source_ids`, `evidence_count`	How the estimate was made and which source datasets it rests on. Source names and links are in the country metadata JSON under `facility_layer.data_lead_sources`.

Two rules before you add things up. Do not add the points to the raster: the raster already contains each site's mapped value, spread over about 2 km around it within its region. And when you compare a site with a regional total, use the mapped column on the same basis as the total (`mapped_elec_facility_gwh` for gross, `mapped_net_facility_gwh` for net), which is what the total contains.

For a network model that wants sites as point loads: take `elec_res_gwh + elec_non_res_gwh` (or the net twins) from the pixels as the surface, and attach each site's `mapped_elec_facility_gwh` (or `mapped_net_facility_gwh`) at its point. Per region and per country this adds up to the total columns.

The region metadata file

One row per official region and year: what a region's total follows and how to cite it. A region is a local authority district in Great Britain and Northern Ireland (361 for 2021 to 2023, 350 in 2024), an intercommunality in metropolitan France (1,236, including four communes outside any intercommunality), a commune in the overseas territories, a Forward Sortation Area in Canada (1,634, the first three characters of the postal code), a district (sigungu) in South Korea (274 polygons sharing 229 published figures), and a local government area in Australia (544).

- `region_id` and `region_desc` name the official unit; `region_uid` is the join key to the pixel table.
- `level_source` and `level_source_citation` give the official statistics the region's totals follow, with publisher, license and URL. `shape_source` and its citation name any finer official source used to shape the split within the region.
- `res_anchor_basis`, `non_res_anchor_basis` and `basis_note` say on what basis the official figure was published and any caveat about it.
- `btm_covered` is `true` where a rooftop-solar estimate was applied in that region and year.
- `res_level_correction_factor` is the factor by which finer official residential figures were scaled to match the coarser official total, where both exist (Canada, Australia, South Korea). Blank elsewhere.
- `official_res_gwh` and `official_non_res_gwh` are the published figures for the unit, with `official_source` and `official_basis` saying where they come from and on what basis; `served_res_gwh` and `served_non_res_gwh` are the map's comparable sums (net residential; net non-residential plus the large facilities the figure counts, `in_frame_facility_gwh`); `facility_outside_official_gwh` is what the map adds on top, on the gross basis; `official_relation` is `equal` where the served sums reproduce the figure and `shape` where the figure only shapes the split of a coarser total. Where several regions share one published figure (Korean districts), `official_unit_id` names it: deduplicate on it before summing. Where a region has no published figure of its own (Canada outside Ontario, Australian council areas without billing), all of these columns are blank.
- The parquet version carries the region polygon; the CSV has the same columns without geometry.

Reconciling with the official figures you already use

Where your area is an official region, you do not need a spatial join: open the region metadata file (the CSV opens in a spreadsheet), filter to your region and year, and read the published figures, the map's comparable sums and `official_relation` side by side; to get the map's totals by sector and basis, sum the pixel rows that carry the region's `region_uid`. What you will find, by country:

- **Great Britain, France and the overseas territories** (`official_relation = equal`). `net_res_gwh` summed over a local authority, intercommunality or commune equals the published residential figure. `net_non_res_gwh` plus the grid draw of the large facilities the statistics count (`in_frame_facility_gwh`) equals the published non-residential figure. What the map adds on top, on the gross basis, is in `facility_outside_official_gwh`: in Great Britain, transmission-connected steelworks and smelters and every site's own generation; in France, own generation only. On the net basis, France's totals equal the published figures and Great Britain's exceed them only by the grid draw of the transmission-connected sites. On the gross basis, rooftop solar used on site is added to both sectors where it is modeled.
- **Canada, Australia, South Korea** (`official_relation = shape`). The published figure for a region shapes the split of a coarser official total (province, state, nation), so a region's served sum differs from its figure by design: in the latest year the median gap is about 6% for Ontario postal areas and 5% to 9% for Australian council areas (residential to non-residential), while Korean districts reproduce KEPCO's figures. The province, state and national sums follow the official figure on the basis it was published, to within a fraction of a percent. In Australia, as in Great Britain, net residential is the metered level and gross adds the rooftop-solar estimate; non-residential follows the Australian Energy Statistics as published.

The country metadata JSON

A small file worth opening once per country: `bases`, `sectors` and `additivity` in machine-readable form; `facility_layer` (the size threshold, classes present, how many sites have estimates and how many are scaled to fit, the yearly amount outside the official statistics in `out_of_frame_gwh_by_year` and `label_gap_gwh_by_year`, and the sources behind the estimates); `official_statistics` (one plain statement of which published figure the region

file's official columns carry and how the map relates to it); `derived_metrics` (the per-person and per-built-area columns, their denominators, thresholds and sources); `level_source_citation`; and `license`, `citation`, `attribution` and `upstream_attributions`, which lists the credit lines upstream sources require.

Adding pixels up to your own areas

One rule, used by the site and every recipe here: a cell counts once, in the polygon that contains its center. Cells are exactly 1 km² on the equal-area grid, so there is no area weighting to do. If your area is an official region, skip the join and sum by `region_uid` (previous section). Open the recipe for your tool:

▾ QGIS, no code

1. Add the GeoTIFF for the year and basis you want (Layer, Add Layer, Add Raster Layer). Do not reproject the raster; leave it in its own projection and let QGIS reproject your boundary on the fly.
2. Add your boundary layer, in any projection.
3. Processing Toolbox, Raster analysis, Zonal statistics: input layer = your boundary, raster layer = the GeoTIFF, raster band = 4 for the total (1 residential, 2 non-residential, 3 large facilities), statistics = Sum, and set the output column prefix (for example `gross_tot_`) so the runs do not overwrite each other. Run once per band you need, and again with the net file. Use Sum on the GWh rasters only, never on the ratio rasters.

Zonal statistics assigns each cell to the polygon that contains its center, the same rule as below. A QGIS build whose GDAL has the Parquet driver (check Help, About; many installers do not include it) also opens the GeoParquet pixel table as a vector layer, where you can filter by `region_uid` and read sums from the Statistics panel. Otherwise the GeoTIFF route above is the one to use.

▾ Python

Geopandas 1.0 or later. Install with `pip install geopandas pyarrow fsspec aiohttp`.

```python
import fsspec
import geopandas as gpd

RELEASE = "2026.10.4"
url =
f"https://data.demandgrid.net/releases/{RELEASE}/exports/gb/gb_electricity_pixels_2024.pa

areas = gpd.read_file("my_areas.gpkg").to_crs("EPSG:4326")
pixels = gpd.read_parquet(
    url,
    bbox=tuple(areas.total_bounds),  # fetches only the row groups that overlap
    columns=["pixel_id", "region_uid", "elec_tot_gwh", "net_tot_gwh", "geometry"],
    filesystem=fsspec.filesystem("https"),
)
# each cell counts once, in the polygon that holds its center
centers =
pixels.set_geometry(pixels.geometry.to_crs("ESRI:54009").centroid.to_crs("EPSG:4326"))
joined = gpd.sjoin(centers, areas, predicate="within")
totals = joined.groupby("area_name")[["elec_tot_gwh", "net_tot_gwh"]].sum()
```

▾ R

Download the parquet file first. `sfarrow::st_read_parquet()` reads it; with GDAL 3.5 or later, `sf::st_read("gb_electricity_pixels_2024.parquet")` does too.

```r
library(sf)
library(sfarrow)
library(dplyr)

pixels <- st_read_parquet("gb_electricity_pixels_2024.parquet")
areas <- st_read("my_areas.gpkg") |> st_transform(st_crs(pixels))

centers <- pixels |>
  st_transform("ESRI:54009") |>
  st_centroid() |>
  st_transform(st_crs(pixels))
joined <- st_join(centers, areas, join = st_within, left = FALSE)
totals <- joined |>
  st_drop_geometry() |>
  group_by(area_name) |>
  summarise(gross_gwh = sum(elec_tot_gwh), net_gwh = sum(net_tot_gwh))
```

DuckDB

DuckDB 1.1 or later reads the GeoParquet geometry directly; on older versions wrap `p.geometry` in `ST_GeomFromWKB()`. A `bbox` filter with literal bounds lets DuckDB skip the row groups outside your area; a join condition does not. Taking the centroid in longitude and latitude moves it by far less than a meter for a 1 km cell, so the sums match the recipes above.

```
INSTALL spatial; LOAD spatial;
INSTALL httpfs;  LOAD httpfs;

CREATE TABLE areas AS SELECT * FROM ST_Read('my_areas.gpkg');  -- in
longitude/latitude

-- the bounds of your areas, as literals (west, south, east, north), so only the
overlapping row groups are read
WITH pixels AS (
  SELECT * FROM
read_parquet('https://data.demandgrid.net/releases/2026.10.4/exports/gb/gb_electricity_pi
  WHERE bbox.xmin <= -1.3 AND bbox.xmax >= -1.8 AND bbox.ymin <= 53.9 AND bbox.ymax >=
53.7
)
SELECT a.area_name,
       SUM(p.elec_tot_gwh) AS gross_gwh,
       SUM(p.net_tot_gwh)  AS net_gwh
FROM pixels p
JOIN areas a ON ST_Within(ST_Centroid(p.geometry), a.geom)
GROUP BY a.area_name;
```

Comparing places

Before you put two places side by side, especially across countries:

- Use the same basis on both sides.
- Check the `year` convention: calendar years, or financial years in Australia.
- Check the region metadata for whether each area's total is an official figure or an allocated share of a coarser one (see "Reconciling" above), and `training_tier` for whether local data trained the model there.
- Check `btm_covered`, not the country-level `no_btm_model`, before comparing gross minus net.
- Use the residential columns, or set `elec_facility_gwh` aside, when one large site would swamp a comparison of ordinary demand.

Known issues in this release

Known in release 2026.10.4:

- **Great Britain.** Northern Ireland has no official figure for 2024 yet, so its cells are blank that year and its facility points stop at 2023.
- **France, Martinique, South Korea.** The region metadata file's `basis_note` says "published gross == net". That holds for residential and non-residential, which have no rooftop-solar model there, but not for large facilities, whose own generation makes gross higher than net.
- **Australia.** Sixty mines have `estimate_status = no_estimate` and add nothing to the raster.
- **Australia.** Where a shaping source starts partway through the series (Ausgrid residential figures from 2019), a council's residential split can step that year; check `shape_source` by year before charting a trend.
- **Everywhere.** The facility points do not add up to the raster wherever a region's sites are scaled to fit its official figure, and only the sites held below half their estimate carry `scaled_to_fit`; use the mapped columns for any comparison with the raster.

Anything found after publication is listed on the Changelog.

Versions, citing and license

Every release has an id such as `2026.10.4` and stays published unchanged at its own URL; the id is in the file path and in the metadata JSON, not inside the tables or rasters. Pin the id in anything you publish: a later release can change values everywhere, and the Changelog says what moved and why.

The data is licensed CC BY 4.0. Cite it as:

Cite this release

Ben Moscona, Stephen Pantano, Saul Griffith (Rewiring the World), 2026. Downscaled 1 km Electricity Consumption, release 2026.10.4. CC BY 4.0. demandgrid.net

Copy citation

Data license: CC BY 4.0 (attribution required). Pin the release id in anything you publish.

Add the date you accessed it where your style requires one, and use the release id in the URL of the files you actually downloaded.

Attribution also covers the sources named in each country's metadata: the official statistics in `level_source_citation` (for Great Britain, DESNZ subnational electricity consumption statistics, Open Government Licence v3.0), Climate TRACE (CC BY 4.0) behind most facility estimates, and, where a facility value was sized from an OpenStreetMap power plant (Great Britain, France, Australia), the credit line © OpenStreetMap contributors. Each country's `upstream_attributions` lists what applies.