

Smart meter consumption data: Technical documentation

Creation date	2021-08-23
Edition	03
Author	Ellen Webborn, Jessica Few
Project	Smart Energy Research Lab (SERL)
Organisation	University College London (UCL)

Table of Contents

Introduction	2
Data collection	3
Smart meter read types	3
Changes since the previous edition	4
Smart meter data tables.....	4
Half-hourly data	4
Daily data	6
Error flags	8
Flags for individual reads	8

'Max reads' (error flag -1).....	9
Very high reads (error flag -2)	9
Readings in the wrong units (error flag -4)	10
Zero Reads	11
Flags for the daily and half-hourly sum match	11
Data quality summary tables.....	11
Read-type data quality summary	12

Introduction

This document describes the half-hourly and daily datasets available to researchers with secure access to the SERL Observatory datasets, along with two data quality summary tables. The half-hourly data is provided in monthly files and the daily data is provided in yearly files. The data were collected from the earliest date available. The datasets described in this document are:

- *serl_smart_meter_daily_{YYYY}_edition{number}.csv*: yearly files of daily electricity and gas readings with some additional derived columns.
- *serl_smart_meter_half_hourly_{YYYY}_{mm}_edition{number}.csv*: monthly files of half-hourly electricity and gas readings with some additional derived columns (note that this dataset also includes reactive readings and export readings where available). Note that the monthly cut off follows UTC time.
- *serl_smart_meter_rt_summary_edition{number}.csv*: data quality summary for each read type for each participant (such as number of errors found by type) and basic read statistics (such as mean and maximum)

where *{number}* is the number of the data release, e.g. "02" (note that the first data release files are labelled with the release date rather than edition number), *{YYYY}* and *{mm}* are the year and month that files relate to respectively.

This document is structured as follows: we start with some basic information about how the data were collected and the different types of reading available, then we describe the two types of smart meter data tables (daily - multiple files each containing a year of daily data, and half-hourly - multiple files each containing a month of half-hourly data). Next we define the different types of error flag created, and finally we describe the data quality summary table (at the read-type level). Note that the participant summary table documentation was previously given in this file, but has been moved to its own document.

Data collection

Half-hourly and daily smart meter readings are stored on the smart meter, and accessed by the Smart Energy Research Lab (SERL) as follows. The University of Essex (UK Data Archive) uses a DCC adaptor service provided by CGI to communicate with the DCC, which acts as a pipe to communicate the smart meter readings to CGI, who send the readings to the UK Data Archive. This happens every night to collect data from the previous day. Not all properties have a gas meter we are able to access - for example, if there is only a DCC-enrolled electricity meter but not a DCC-enrolled gas meter, or if the property does not have mains gas. Check the EPC and survey data to identify properties which have gas central heating but no SERL gas data if this may affect your analysis.

Smart meter read types

Smart meter read types are defined by two variables: 'deviceType' and 'readType'. Together they combine to define the type of smart meter data. The full list of smart meter data types are shown in the table below. Note that 'GPF' stands for 'Gas Proxy Function' (a proxy for the gas meter) and 'ESME' stands for 'Electricity Smart Metering Equipment' (the electricity meter).

Table 1: Smart meter data types, defined by 'deviceType' and 'readType'.

<i>deviceType</i>	<i>readType</i>	<i>Units</i>	<i>Description</i>
ESME	DL	Wh	Daily active electricity import
GPF	DL	m ³	Daily gas import
ESME	AI	Wh	Half-hourly active electricity import
ESME	RI	varh	Half-hourly reactive electricity import
GPF	AI	m ³	Half-hourly gas import

<i>deviceType</i>	<i>readType</i>	<i>Units</i>	<i>Description</i>
ESME	AE	Wh	Half-hourly active electricity export
ESME	RE	varh	Half-hourly reactive electricity export

Changes since the previous edition

In addition to the inclusion of participants from wave 3 recruitment and the extension of data to 31st May 2021, a number of changes have been made to the processing of the raw data. The main changes are:

- Previously we provided half-hourly data as one file, however this file was becoming too large to handle due to the increasing time-span of the data and the increase in participant numbers. Going forward we will provide the half-hourly data in monthly files. Similarly we previously provided the daily data as one file, but going forward we will provide the daily data in yearly files.
- We have removed the median and standard deviation of each read type in the participant summary.
- We previously documented the participant-level summary in this file, this has been moved to a separate file for increased clarity.

We continue to analyse and improve the SERL datasets and we aim to strike a balance between improving the data and minimising inconvenience to researcher caused by changes. We welcome feedback to improve our processes.

Smart meter data tables

Half-hourly data

The half-hourly data tables have 17 columns. The fields are described in the table below. While the original data have not been modified, additional columns have been added to flag potential errors (see the Error Flags section below) and convert between units. This processing was done using R version 4.0.1 (2020-06-06). Code used for processing will be made available shortly on the SERL Github repository github.com/smartEnergyResearchLab.

Note that if no data were returned for any meter at a particular time then that row will be missing from the dataset rather than being an empty row. The 'class' field is the R class (e.g. R calls a Boolean a logical).

Table 2: Half-hourly data fields. Error flags are defined in the Error Flags section below.

<i>Field</i>	<i>Description</i>	<i>Units</i>	<i>Class</i>	<i>Example value</i>	<i>Variable type</i>
PUPRN	Pseudonymised participant identifier	NA	character	1VUXXXF1	Assigned
Read_date_effective_local	Date of read (same as date of Read_date_time_local unless read taken at midnight, then the previous day since data pertains to the previous day)	%Y-%m-%d	Date	2019-11-01	Derived
Read_date_time_local	Time read taken (local time: GMT or BST)	%Y-%m-%d %H:%M:%S tz	character	2020-07-02 00:03:30 BST	Derived
Read_date_time_UTC	Time read taken in UTC	%Y-%m-%d %H:%M:%S	POSIXct, POSIXt	2020-07-02 00:02:30	Primary
HH	Half-hour identifier between 1 and 48 (NA if not exactly on the hour or half-hour)	NA	integer	48	Primary
Valid_read_time	FALSE if read time is not on the hour or half hour, otherwise TRUE	NA	logical	TRUE	Derived
Elec_act_imp_flag	Half-hourly electricity active import error flag	NA	numeric	-2	Derived
Elec_react_imp_flag	Half-hourly electricity reactive import error flag	NA	numeric	1	Derived
Elec_act_exp_flag	Half-hourly electricity active export error flag	NA	numeric	-4	Derived
Elec_react_exp_flag	Half-hourly electricity reactive export error flag	NA	numeric	2	Derived
Gas_flag	Half-hourly gas import error flag	NA	numeric	0	Derived
Elec_act_imp_hh_Wh	Half-hourly electricity active import read	Wh	integer	109	Primary
Elec_react_imp_hh_varh	Half-hourly electricity reactive import read	varh	integer	15	Primary
Elec_act_exp_hh_Wh	Half-hourly electricity active export read	Wh	integer	65	Primary

<i>Field</i>	<i>Description</i>	<i>Units</i>	<i>Class</i>	<i>Example value</i>	<i>Variable type</i>
Elec_react_exp_hh_varh	Half-hourly electricity reactive export read	varh	integer	14	Primary
Gas_hh_m3	Half-hourly gas import read	m ³	numeric	0.244	Primary
Gas_hh_Wh	Half-hourly gas import read in Wh using standard conversion, assuming calorific value = 39.5	Wh	numeric	2737.835	Derived

Daily data

The daily data tables have 15 columns. The fields are described in the table below. While original data has not been modified, additional columns have been added to flag potential errors (see Error Flags section below) and to convert between units. This processing was done using R version 4.0.1 (2020-06-06). Code used for processing is available on the SERL Github repository github.com/smartEnergyResearchLab.

Note that if no data were returned for any meter on a particular day then that row will be missing from the dataset rather than being empty. The exception is if there were no daily reads but the right number of half-hourly reads (48 unless the clocks changed) to sum to the daily total. In these instances the rows have been added to allow for easy imputation of a missing daily read with the sum of the half-hourly reads. For comparison between daily reads and daily sums, half-hourly data requires conversion (provided) to local time. A csv file is provided (*bst_dates_to_2024.csv*) which lists the start and end dates of British Summer Time (BST) for reference so the number of half-hours expected on each day is clear. Researchers may wish to avoid clock change days in their analysis, or take them into consideration.

The daily data includes fields added for comparison between the daily readings and half-hourly readings for the same days. As described in Table 3 there are columns for the sum of the half-hourly readings if there were 48 valid reads taken at the right times (or 46 or 50 reads if the clocks changed), the difference between the half-hourly sum and the daily read (if both exist and are valid), and a 'sum_match' column coded to state the condition of this match. The sum match codes are defined in Table 6 in the Error Flags section. Note that reads are considered invalid if taken at the wrong time (not at midnight for daily reads; not on the hour or half hour for half-hourly reads).

Table 3: Daily data fields. See the Error Flags section below for definitions of the error flags used in this table.

<i>Field</i>	<i>Description</i>	<i>Units</i>	<i>Class</i>	<i>Example value</i>	<i>Variable Type</i>
PUPRN	Pseudonymised participant identifier	NA	character	1VUXXF1	Assigned
Read_date_effective_local	Date that read relates to (in local time): previous day, unless after midday (then same day)	%Y-%m-%d	Date	2019-11-01	Derived
Read_date_time_local	Time and date of read (local time). Time not stated if at midnight	%Y-%m-%d	POSIXct, POSIXt	2019-11-02	Primary
Valid_read_time	TRUE if reading was at midnight, otherwise FALSE	NA	logical	TRUE	Derived
Elec_act_imp_flag	Daily electricity active import error flag	NA	numeric	-2	Derived
Elec_sum_match	Error code for whether the sum of half-hourly electricity active import matches the daily electricity read	NA	numeric	1	Derived
Gas_flag	Daily gas import error flag	NA	numeric	2	Derived
Gas_sum_match	Error code for whether the sum of half-hourly gas import matches the daily gas read	NA	numeric	0	Derived
Elec_act_imp_d_Wh	Daily electricity active import read	Wh	integer	5839	Primary
Unit_correct_elec_act_imp_d_Wh	Daily electricity active import read corrected from kWh to Wh where kWh reporting is suspected, otherwise equals Elec_act_imp_d_Wh	Wh	integer	5839	Derived
Elec_act_imp_hh_sum_Wh	Sum of half-hourly electricity active import reads for this date (NA if there were not 48* valid reads). *46 required when the clocks go forward, 50 when the clocks go back.	Wh	integer	5742	Derived
Gas_d_m3	Daily gas read	m ³	numeric	8.214	Primary

<i>Field</i>	<i>Description</i>	<i>Units</i>	<i>Class</i>	<i>Example value</i>	<i>Variable Type</i>
Gas_hh_sum_m3	Sum of half-hourly gas reads for this date (NA if there were not 48* valid reads). *46 required when the clocks go forward, 50 when the clocks go back.	m ³	numeric	8.763	Derived
Gas_d_kWh	Daily gas import read in kWh using standard conversion, assuming calorific value = 39.5	kWh	numeric	92.16628	Derived

Error flags

Flags for individual reads

Table 4 shows the meaning of each error flag value. These flags are used for both daily and half-hourly reads. The data quality report gives details about the number of each error found within the data along with other descriptive statistics. Note that the first data release (2020-08) had slightly different error flags (no flag 3 and a flag for invalid read times). Importantly, the error flags are now split so that **these flags relate only to the quality of the read** (missing, too high etc.) and **not the validity of the read time, unless stated otherwise**. The 'Valid_read_time' variable indicates whether the read was at the correct time (or not). Therefore to filter on valid reads at valid times, use both the relevant error flag and the Valid_read_time flag.

Table 4: Error flags and their meanings.

<i>Flag</i>	<i>Meaning</i>	<i>Details</i>
3	Ignore	Invalid read time and no read - row exists for a different read type so ignore
2	No meter	The gas (or very rarely electricity) meter does not exist in the DCC inventory
1	Valid	The read exists and does not meet any of the other error flag criteria, thus presumed valid (although may not have a valid read time - check separately)
0	Missing	The read should exist but is missing

<i>Flag</i>	<i>Meaning</i>	<i>Details</i>
-1	Max read	The read is (presumably) the largest storable number on the meter - details below
-2	Very high but not max	The read is higher than plausible but not a 'Max read' - see section below for thresholds
-3	Negative	The read is negative (none found)
-4	Elec in kWh	The electricity read was reported in kWh rather than Wh - details below
-5	Valid read, invalid read time	Originally flagged as valid (1) but with valid_read_time = FALSE, therefore we cannot say over what time period the data has been recorded.

'Max reads' (error flag -1)

We discovered that multiple participants have some electricity readings as 16777215 Wh and multiple gas readings of 16777.215 m³. These numbers are all 1s in binary which implies they are the maximum read the (32-bit) meter can store, and likely due to some technical error. We call this type of error 'Max read'. Note that in a few cases the number is the max read stored in 64 bits – these are replaced by the 32-bit maximum to save memory.

Very high reads (error flag -2)

For gas and active electricity reads we flag if the reading is larger than we deem plausible, attempting to be cautious with our definition of 'plausible'. The following table shows our definitions of what constitutes a 'very high' read (so long as the read is not high enough to be a 'max read'). For reactive reads we did not have an informed maximum, so set the somewhat arbitrary cutoff as one was required for the code.

Table 5: Limits used to determine if a read is high enough for a 'very high read' flag (-2).

<i>deviceType</i>	<i>readType</i>	<i>High Read Limit</i>	<i>Units</i>	<i>Assumptions</i>
ESME	DL	1,152,000	Wh	Max 200A fuse, 240V
ESME	AI	24,000	Wh	Max 200A fuse, 240V
ESME	AE	2,000	Wh	4kW max PV capacity

<i>deviceType</i>	<i>readType</i>	<i>High Read Limit</i>	<i>Units</i>	<i>Assumptions</i>
ESME	RI	10,000	varh	Arbitrary cutoff
ESME	RE	10,000	varh	Arbitrary cutoff
GPF	DL	384	m ³	16m ³ /hr max capacity
GPF	AI	8	m ³	16m ³ /hr max capacity

Readings in the wrong units (error flag -4)

We also discovered that all daily electricity active import readings were all extremely low for some participants. At the time of the initial investigation (August 2020) there were no participants with a maximum read between 85 and 2285 Wh; therefore we determined that any electricity active import (or export) meter with its highest reading less than 100 was erroneously reporting in kWh rather than in Wh as specified by the Smart Energy Code. This was verified by comparing the daily readings with the sum of the half-hours for the same day. Any readings that were deemed to be valid according to all other criteria were given the 'Elec in kWh' error flag and their data were multiplied by 1000 in the 'unit-corrected' column. Note that we set a limit of at least 30 readings in order to define a meter as recording in kWh.

Subsequent investigations in March 2021 revealed that some meters were correctly recording daily electricity readings in Wh up until the date of a meter replacement, at which the daily reads became approximately 1/1000th of the sum of the half-hourly reads. By 'approximately' we mean that after dividing the half-hourly sum by 1000 and rounding down, the result is within 1 of the daily read. For example, half-hourly sums between 5000 and 599 would be classed as approximately 1000 times bigger than a daily read between 4 and 6 (in order to handle rounding errors and slight read mismatch). In such cases, if a household has at least 5 rows with the daily and half-hourly sums in this situation (excluding daily reads of 0 which may be a different kind of error), all of such cases for the household are flagged as a unit error as above. Without checking all cases manually, a very small number of reads will be mis-flagged as unit errors or unit errors will be missed.

It is possible that some half-hourly sums are approximately 1/100th of the daily read, indicating readings in tens of Wh. This has been found to affect just a handful of meters at present, and it is left to researchers to decide how best to deal with daily and half-hourly sum mismatches in general. In most cases we believe that the sum of half-hours is more reliable than a daily read, but it depends on the individual meter, and not all days have complete valid half-hourly reads to sum. Note that if there are not the correct number of valid half-hourly reads taken at the right times then the half-hourly sum variable will be NA.

Zero Reads

Zero reads are not yet flagged as invalid although this may change in subsequent data releases. It has been found that some daily reads default to zero during British Summer Time (BST) which could be an obvious error to flag, but the validity of other zero reads is less clear. Considering mismatches between half-hourly sums and daily reads is advised when performing data quality analysis.

Flags for the daily and half-hourly sum match

The daily data table contains fields called "Elec_sum_match" and "Gas_sum_match" which give an error flag for how the sum of the half-hourly reads for that day compares with the daily read. They are described in the table below. It isn't always possible to compare the sum and the daily readings; if any were reported in the wrong units, if any reads were taken at the wrong time, or if any of the reads were missing. By 'any' we mean the daily read and the 48 half-hourly reads (46 or 50 on clock change days).

Table 6: Electricity and gas 'sum_match' values and their meanings.

Code	Meaning
3	Daily electricity read originally recorded in kWh, match not tested due to rounding issues.
2	No meter: the (gas) meter does not exist in the DCC inventory
1	Daily read and half-hourly sum match: for electricity difference ≤ 1 Wh, for gas difference ≤ 1 L
0	Comparison not possible: do not have 48 valid half-hourly reads or daily read isn't valid
-1	Daily read and half-hourly sum are similar but don't match: for electricity $1 < \text{difference} \leq 10$ Wh, for gas $1 < \text{difference} \leq 10$ L
-2	Daily read and half-hourly read are neither similar nor match but are both valid

Data quality summary tables

There are two tables that give a summary of the data quality at 1) the individual read type level and 2) the participant level. The second is described in the participant summary documentation.

Read-type data quality summary

A read type is a combination of device type and schedule type, such as daily active electricity import or half-hourly reactive electricity export. There are 24 columns in the read-type data quality table. Each read type for each participant has its own row, which gives the number of readings with each error code, the start and end dates of the schedule (when we expect the earliest and latest readings to be), and some calculated columns for the percentage missing and valid etc. The details are given in the table below. Note that reads with error code 3 are not included in this table as they are an artifact of the data structure rather than genuine reads (see above for details).

Table 7: Read-type data quality summary: data fields. Error codes are explained above in the Error Flags section. Note that statistics for valid reads exclude valid reads recorded at the wrong time.

<i>Field</i>	<i>Description</i>	<i>Units</i>	<i>Class</i>	<i>Example value</i>
PUPRN	Pseudonymised participant identifier	NA	character	1VUXXXF1
deviceType	Device type: gas (GPF) or electricity (ESME) meter	NA	character	GPF
readType	Defined in Table 1	NA	character	AI
theoreticalStart	Earliest possible reading for the schedule	%Y-%m-%d	Date	2019-11-01
theoreticalEnd	Latest possible reading for the schedule	%Y-%m-%d	Date	2020-02-29
firstValidReadDate	Earliest date with a valid read (error flag 1 and Valid_read_time = TRUE)	%Y-%m-%d	Date	2018-11-26
lastValidReadDate	Latest date with a valid read (error flag 1 and Valid_read_time = TRUE)	%Y-%m-%d	Date	2020-05-31
daysRange	Schedule length = scheduleEnd - scheduleStart + 1	NA	numeric	100
maxPossReads	Maximum possible reads available (= daysRange for daily data, = 48 * daysRange for half-hourly)	NA	numeric	4800

<i>Field</i>	<i>Description</i>	<i>Units</i>	<i>Class</i>	<i>Example value</i>
percValid	Percentage of possible reads that are valid (error flag 1 and Valid_read_time = TRUE) rounded to 2 decimal places	NA	numeric	95.02
percValidOrUnitError	Percentage of possible reads that are valid or have a unit error (Valid_read_time = TRUE and error flag 1, or -4) rounded to 2 decimal places	NA	numeric	96.98
percMissing	Percentage of possible reads that are missing (error flag 0) rounded to 2 decimal places	NA	numeric	2.13
percError	Percentage of possible reads that are erroneous (error flag -1, -2, -3, -4 or Valid_read_time = FALSE) rounded to 2 decimal places	NA	numeric	3.04
valid	Number of valid readings taken at the right time (error flag 1 and Valid_read_time = TRUE)	NA	integer	96
validOrHHsumValid	Number of valid readings taken at the right time (error flag 1 and Valid_read_time = TRUE), or days with a valid half hourly sum (all expected half hourly readings have error flag 1)	NA	integer	96
validWrongTime	Number of valid readings taken at the incorrect time (error flag 1 and Valid_read_time = FALSE)	NA	integer	7
wrongUnits	Number of readings with electricity recorded in kWh (error flag -4)	NA	numeric	3
missing	Number of missing readings (error flag 0)	NA	numeric	27
maxRead	Number of readings with the 'Max Read' error (flag -1)	NA	integer	4
highRead	Number of readings between the 'very high' and the 'max read' thresholds (error flag -2)	NA	integer	2
negative	Number of negative readings with (error flag -3)	NA	numeric	0

<i>Field</i>	<i>Description</i>	<i>Units</i>	<i>Class</i>	<i>Example value</i>
minValidRead	Minimum read of the valid reads (after unit-correction if necessary)	Wh (elec), m ³ (gas)	numeric	0
maxValidRead	Maximum read of the valid reads (after unit-correction if necessary)	Wh (elec), m ³ (gas)	numeric	302
meanValidRead	Mean of the valid reads (after unit-correction if necessary), 2 decimal places	Wh (elec), m ³ (gas)	numeric	43.21