

zxweather

Database Structure

Document Number: DAZW-DB01

July 2012

This document provides an overview of the database structure used by zxweather.

Revision/Update Information:	This is a new manual.
Operating System:	Any
Software Version:	zxweather 0.1

© Copyright David Goodwin, 2012.

Use, reproduction and modification of this document is permitted subject to the terms of the GNU Free Documentation License, Version 1.3 or any later version published by the Free Software Foundation. See <http://www.gnu.org/copyleft/fdl.html> for full license text.

Contents

1	Introduction	1
1.1	Related Documents	1
2	Tables	3
2.1	sample	3
2.1.1	Applications	3
2.1.2	Columns	3
2.1.3	Indices	6
2.1.4	Triggers	6
2.2	live_data	6
2.2.1	Applications	6
2.2.2	Columns	7
3	Views	9
3.1	latest_record_number	9
3.2	Daily, Montly and Yearly Records	9
3.2.1	Columns	9
4	Functions and Trigger Functions	11
4.1	Functions	11
4.1.1	dew_point	11
4.1.2	wind_chill	11
4.1.3	apparent_temperature	12
4.2	Trigger Functions	12
4.2.1	compute_sample_values	12
4.2.2	live_data_update	13
5	Domains	15
5.1	rh_percentage	15
5.2	wind_direction	15
6	Notification Channels	17
6.1	live_data_updated	17
6.1.1	Applications	17
6.2	new_sample	17
6.2.1	Applications	17
6.3	update_complete	17
6.3.1	Applications	18

Introduction

zxweather uses a PostgreSQL database to store weather data and for some communication between components. No other RDMBS is supported.

This document is divided up into five chapters which define the various parts of the database:

- Chapter 2: Tables
- Chapter 3: Views
- Chapter 4: Functions and Trigger Functions
- Chapter 5: Domains
- Chapter 6: Notification Channels

1.1 Related Documents

DAZW-IG01 zxweather Installation Reference, version 0.1
DAZW-WG01 zxweather WH1080 Utilities Users Guide, version 0.1

This chapter documents the database tables used by zxweather.

2.1 sample

This table contains weather samples from the weather station. Most columns come direct from the weather station without any processing but some (such as the rainfall column) are calculated by a trigger on insert.

2.1.1 Applications

The following applications use this table:

- wh1080 (database update utility)
- wh1080d/wh1080svc (database update daemon/service)
- web interface

2.1.2 Columns

sample_id

This is the primary key.

sample_interval

This is the number of minutes since the previous sample on the weather station. It is used for calculating timestamps. It is not used directly for anything else.

record_number

The history slot on the weather station this record was downloaded from. Its range is 0-4079. It is only used to track which records have already been downloaded from the weather station when performing updates and for calculating timestamps.

download_timestamp

When the record was downloaded from the weather station. It is not used directly for anything.

last_in_batch

If this record was the final record in a batch of records downloaded from the weather station. This records timestamp is the one used to calculate all other records downloaded in the batch.

invalid_data

If the record on the weather station had the invalid data status flag set. The WH1080 documentation refers to this as both "rain counter value is not valid" and "no sensor data received"

time_stamp

The approximate time this sample was taken by the weather station.

If the `last_in_batch` column is `true` then the timestamp value was determined by waiting for the live data record on the weather station to change. The live data records `sample_time` field is then used to compute the timestamp for this record.

If the `last_in_batch` column is `false` then the records timestamp value was calculated based on the its `record_number`, the `sample_interval` of the previous record and the `time_stamp` of the previous record.

The value in this column may be out by as much as a minute.

indoor_relative_humidity

The relative humidity at the base station. Its type is the `rh_percentage` domain type (see section 5.1). Values are 0-99%.

indoor_temperature

The temperature at the base station in degrees Celsius.

relative_humidity

The relative humidity at the outdoor unit. Its type is the `rh_percentage` domain type (see section 5.1). Values are 0-99%.

temperature

The temperature at the outdoor unit in degrees Celsius.

dew_point

Dew point. This column's value is calculated on insert by a trigger using the `dew_point` function (Section 4.1.1). It is calculated from the `temperature` column.

wind_chill

Wind chill. This column's value is calculated on insert by a trigger using the `wind_chill` function (Section 4.1.2). It is calculated from the `temperature` column.

apparent_temperature

Apparent temperature. This column's value is calculated on insert from the `temperature` column by a trigger using the `apparent_temperature` function (Section 4.1.3).

absolute_pressure

The absolute pressure in hectopascals (hPa, same as the millibar).

relative_pressure

Calculated relative pressure in hectopascals (hPa, same as the millibar).

This column is for future use. Its value is not currently calculated.

average_wind_speed

Average wind speed in metres per second.

gust_wind_speed

Gust wind speed in metres per second.

wind_direction

Wind direction. Its type is the `wind_direction` domain type (see section 5.2 for valid values).

rainfall

The rainfall for this record. Its value is calculated based on the `total_rain` and `rain_overflow` values. Rain is measured in millimetres. Rainfall is in multiples of 0.3mm.

total_rain

Total rainfall since it last overflowed. This is used only for calculating the values in the `rainfall` column. Values are a multiple of 0.3. Unit is the millimetre.

rain_overflow

If an overflow in the `total_rain` counter has occurred. Used for calculating values in the `rainfall` column.

2.1.3 Indices

This table has the following indices:

- `idx_time_stamp` on `time_stamp` column.

2.1.4 Triggers

The `compute_sample_values` trigger function is executed for each row inserted. This trigger calculates values for columns such as `dew_point`. See section 4.2.1 for more details on this trigger function.

2.2 live_data

This table stores live data from the weather station. It only ever contains a single record and as such has no primary key. It is updated by the `wh1080d` program whenever there is new live data. Its set of fields is a subset of those in the sample table.

It has a single trigger on INSERT, DELETE and UPDATE. Inserts and Deletes are blocked, Updates cause a number of fields to be recalculated.

This table does not provide access to live rainfall data. Calculating this would be difficult as it requires access to the previous samples rainfall. The `wh1080d` application updates live data before downloading and inserting any new history records so the previous samples rainfall data is not available from the `live_data` update trigger or within the `wh1080d` application.

2.2.1 Applications

The following applications use this table:

- `wh1080d/wh1080svc` (database update daemon/service)
- web interface

2.2.2 Columns

This table has no primary key. All columns except wind direction and invalid data allow nulls.

download_timestamp

When this data was downloaded from the weather station.

invalid_data

If the record on the weather station had the invalid data status flag set. The WH1080 documentation refers to this as both "rain counter is not valid" and "no sensor data received". This column must not be null.

indoor_relative_humidity

The relative humidity at the base station. Values are 0-99%.

indoor_temperature

The temperature at the base station in degrees Celsius.

relative_humidity

The relative humidity at the outdoor unit. Values are 0-99%.

temperature

The temperature at the outdoor unit in degrees Celsius.

dew_point

Dew point. This columns value is calculated on update by a trigger using the `dew_point` function (Section 4.1.1). It is calculated from the `temperature` column.

wind_chill

Wind chill. This columns value is calculated on update by a trigger using the `wind_chill` function (Section 4.1.2). It is calculated from the `temperature` column.

apparent_temperature

Apparent temperature. This column's value is calculated on update from the `temperature` column by a trigger using the `apparent_temperature` function (Section 4.1.3).

absolute_pressure

The absolute pressure in hectopascals (hPa, same as the millibar).

relative_pressure

Calculated relative pressure in hectopascals (hPa, same as the millibar).

This column is for future use. Its value is not currently calculated.

average_wind_speed

Average wind speed in metres per second.

gust_wind_speed

Gust wind speed in metres per second.

wind_direction

Wind direction. Its type is the `wind_direction` domain type (see section 5.2 for valid values).

3.1 latest_record_number

This view is used by the wh1080 database update utilities to determine what the most recent record in the database is.

3.2 Daily, Montly and Yearly Records

These views provide minimum and maximum weather conditions (temperature, humidity, etc) for each day, month or year in the database. The views are:

- daily_records
- monthly_records
- yearly_records

The PostgreSQL query planner seems to have difficulty with the queries behind these views when more than one column is included in a where clause. This has been observed in PostgreSQL version 9.1.

3.2.1 Columns

- date_stamp: The date the record is for (daily_records and monthly_records only).
- year_stamp: The year the record is for. Present only in the yearly_records view.
- total_rainfall: Total rainfall during this day, month or year.
- max_gust_wind_speed: Maximum gust wind speed for the period.
- max_gust_wind_speed_ts: Most recent timestamp for the maximum gust wind speed record.
- max_average_wind_speed: Maximum average wind speed for the period.
- max_average_wind_speed_ts: Most recent timestamp for the maximum wind speed record.
- min_absolute_pressure: Minimum absolute pressure during the period.
- min_absolute_pressure_ts: Most recent timestamp for the minimum absolute pressure record.

- `max_absolute_pressure`: Maximum absolute pressure during the period.
- `max_absolute_pressure_ts`: Most recent timestamp for the maximum absolute pressure record.
- `min_apparent_temperature`: Minimum apparent temperature during the period.
- `min_apparent_temperature_ts`: Most recent timestamp for the minimum apparent temperature record.
- `max_apparent_temperature`: Maximum apparent temperature during the period.
- `max_apparent_temperature_ts`: Most recent timestamp for the maximum apparent temperature record.
- `min_wind_chill`: Minimum wind chill during the period.
- `min_wind_chill_ts`: Most recent timestamp for the minimum wind chill record.
- `max_wind_chill`: Maximum wind chill during the period.
- `max_wind_chill_ts`: Most recent timestamp for the maximum wind chill record.
- `min_dew_point`: Minimum dewpoint during the period.
- `min_dew_point_ts`: Most recent timestamp for the minimum dewpoint record.
- `max_dew_point`: Maximum dewpoint during the period.
- `max_dew_point_ts`: Most recent timestamp for the maximum dewpoint record.
- `min_temperature`: Minimum outdoor temperature during the period.
- `min_temperature_ts`: Most recent timestamp for the minimum outdoor temperature record.
- `max_temperature`: Maximum outdoor temperature during the period.
- `max_temperature_ts`: Most recent timestamp for the maximum outdoor temperature record.
- `min_humidity`: Minimum outdoor relative humidity during the period.
- `min_humidity_ts`: Most recent timestamp for the minimum outdoor relative humidity record.
- `max_humidity`: Maximum outdoor relative humidity during the period.
- `max_humidity_t`: Most recent timestamp for the maximum outdoor relative humidity record.

Functions and Trigger Functions

The database contains a number of stored procedures for computing values that don't come directly from the weather station. This chapter gives an overview of those functions.

4.1 Functions

4.1.1 dew_point

Calculates the approximate dew point given temperature and relative humidity. The calculation is based on the August-Roche-Magnus approximation for the saturation vapour pressure of water in air as a function of temperature. It is valid for input temperatures 0 to 60 °C and dew points 0 to 50 °C

The formula is:

$$T_d = \frac{b \gamma(T, RH)}{a - \gamma(T, RH)}$$

where

$$\gamma(T, RH) = \frac{a T}{b + T} + \ln(RH/100)$$

where the temperatures are in degrees Celsius The constants are $a = 17.271$ and $b = 237.7$ °C

This function is *immutable*.

See the article on Wikipedia for more details: <http://en.wikipedia.org/wiki/Dewpoint>

Parameters

- temperature, real
- relative_humidity, integer

4.1.2 wind_chill

Calculates the North American wind chill using the following formula:

$$T_{wc} = 13.12 + 0.6215T_a - 11.37V^{+0.16} + 0.3965T_aV^{+0.16}$$

where T_{wc} is the wind chill index in degrees Celsius, T_a is the air temperature in degrees Celsius and V is the wind speed at 10 metres in kilometres per hour.

The Wind chill temperature is only defined for air temperatures below 10 °C and wind speeds above 4.8 km/h.

This function is *immutable*.

See the article on Wikipedia for more details: http://en.wikipedia.org/wiki/Wind_chill

Parameters

- temperature, real
- wind_speed, real

4.1.3 apparent_temperature

Calculates the Apparent Temperature using the formula used by the Australian Bureau of Meteorology.

The formula used is:

$$AT = Ta + 0.33e - 0.7ws - 4.00$$

Where Ta is the dry bulb temperature (°C), e is the water vapour pressure in hPa and ws is the Wind speed (m/s) at an elevation of 10 metres.

The vapour pressure is calculated from the temperature and relative humidity using the formula:

$$e = rh/100 * 6.105 * 2.718281828^{17.27Ta/(237.7+Ta)}$$

where Ta is the dry bulb temperature (°C), rh is the relative humidity (%).

This function is *immutable*.

See the Australian Bureau of Meteorology website for more details:

http://www.bom.gov.au/info/thermal_stress/

Parameters

- temperature, real
- wind_speed, real
- relative_humidity, real

4.2 Trigger Functions

4.2.1 compute_sample_values

Computes values for the following fields when a new record is inserted:

- dew_point

- wind_chill
- apparent_temperature
- rainfall

Rainfall calculation

The rainfall value is calculated using the following query:

```
1 select into NEW.rainfall
2   case when NEW.total_rain - prev.total_rain >= 0 then
3     NEW.total_rain - prev.total_rain
4   else
5     NEW.total_rain + (19660.8 - prev.total_rain)
6   end as rainfall
7 from sample prev
8 — find the previous sample:
9 where time_stamp = (select max(time_stamp)
10                from sample ins
11                where ins.time_stamp < NEW.time_stamp)
```

The value 19660.8 is the maximum rainfall accumulator value (65536 * 0.3mm). Rainfall values are calculated based on the total rainfall value of the new record and the previous record in the database.

4.2.2 live_data_update

Computes values for the live_data table updates. If the operation is an insert or delete it is blocked.

It computes values for the following fields:

- dew_point
- wind_chill
- apparent_temperature

It also generates a notification on channel live_data_updated,

It is only used on the live_data table.

Domains

Domains are used on a number of columns to simplify constraints, etc. This chapter gives an overview of those domains and their valid values.

5.1 rh_percentage

Used for columns containing relative humidity percentages. Valid values are 1 to 99 (inclusive). It is an `integer` field and does not allow nulls.

5.2 wind_direction

Used for columns containing wind directions. Valid values are:

- N
- NNE
- NE
- ENE
- E
- ESE
- SE
- SSE
- S
- SSW
- SW
- WSW
- W
- WNW
- NW

- NNW
- INV (used for invalid values).

its type is `character varying(3)`

Notification Channels

zxweather uses the PostgreSQL NOTIFY feature to notify components of the system that new data has arrived. This chapter covers the notification channels used by zxweather and their meaning.

6.1 live_data_updated

A notification with no payload is broadcast on this channel when data in the live_data table has been updated. The notification event is sent automatically from the live_data_update trigger function.

Clients interested in displaying live data should listen on this channel rather than polling the database.

6.1.1 Applications

The following applications use this channel:

- zxweather Desktop (GUI) Client

6.2 new_sample

A notification with no payload is broadcast on this channel when a new sample has been inserted into the database. This notification is sent automatically from the compute_sample_values trigger function.

Clients interested in displaying history samples should listen on this channel rather than polling the database.

6.2.1 Applications

The following applications use this channel:

- zxweather Database Replicator

6.3 update_complete

A notification with no payload is broadcast on this channel by the zxweather WH1080 Update Daemon when it has finished updating the live_data table and inserting any new records into the sample table.

This channel does not specify what has been updated. Listen on the `live_data_updated` and `new_sample` channels as well to determine what has been updated.

If there are new history samples then the notification sequence will be:

- `live_data_updated`
- `new_sample`
- `update_complete`

If there are no new history samples then there will be no notification on the `new_sample` channel.

6.3.1 Applications

The following applications use this channel:

- `wh1080d` (WH1080 Update Daemon) - issues notifications when it has finished doing work.
- `zxweather` Database Replicator - waits for a notification on this channel before sending data to the remote database. Other notification channels (`new_sample` and `live_data_updated`) are used to determine what to send

