

Hive

Loading data to Hadoop

We will work with records of hourly temperatures at weather stations in the USA. The file with data is on the local filesystem:

```
/home/pascepel/fel_bigdata/data/teplota-usa.zip
```

We need to get this data on Hadoop (i. e., on HDFS). So:

- copy data to your user directory (I recommend to make a new directory for it);
- unzip;
- look into some lines of unzipped data, find the total number of lines (why are we doing it?);
- make a subdirectory on HDFS in your user directory (why do we need a subdirectory?);
- copy unzipped files from the local filesystem to HDFS (to the directory just created).

Starting of Hive shell

```
beeline -u "jdbc:hive2://hadoop-c1.ics.muni.cz:10000/default;principal=hive/hadoop-c1.ics.muni.cz@ICS.MUNI.CZ"
```

Create your database (if not created yet), give it the same name as your username. Then switch to your database (USE command).

1. Input data as a temporary, external table

1.1. Create an external table *temperature_ext* from files which you have loaded to HDFS.

- format *textfile*
- fields delimited by ","
- rows delimited by "\n"
- first line contains headers – to be skipped
- fields in the file:

Field	Type	Description
station	string	station code
month	int	month number
day	int	day in month number
hour	int	hour number (1–24)
degrees	int	temperature as round($10 \times ^\circ\text{F}$)
flag	string	code of data quality
latitude	double	GPS latitude (negative = southern, positive = northern)
longitude	double	GPS longitude (negative = western, positive = eastern)
elevation	double	position above the sea level in meters
state	string	US state code (incl. dependent territories)
name	string	station name

1.2. We can do SQL queries in the external table. Do some check:

- list some lines (records) of *temperature_ext* table, compare with input data;
- find total numbers of lines (records) and compare with input data (should not be exactly the same – why?);
- find number of lines with NULL values at the field *degrees* (should be only a small part of all lines).

2. Trasfer to the optimized table

2.1. Create an empty internal (managed) table *temperature* with other format and compression:

- format parquet
- compression Snappy (it's necessary to type it by uppercase: SNAPPY)
- we will save degrees as a decimal number, therefore use *double* type

2.2. Insert data from the table *temperature_ext* to the table *temperature*:

- transform degrees from $10\times^{\circ}\text{F}$ to $^{\circ}\text{C}$;
- all other fields transform as they are (no change).

2.3. Check the table *temperature*:

- Write out some rows.
- Find total number of records in the table *temperature* and compare it with total number of records in the table *temperature_ext*.

2.4. The table *temperature* is internal, so Hive is its owner.

- Look for it on HDFS under `/user/hive/warehouse/database_name.db` and find the volume (megabytes).
- Compare the volume with the volume of external table (data loaded on HDFS by you).

2.5. Drop the external table *temperature_ext*. Check that the table is no longer in your database but data are still on HDFS.

3. Partitioned table

3.1. Create an empty internal table *temperature_part*, which will be the same as *temperature* (fields, their types, format, compression) but will be partitioned by month. (Pay attention to the field order!)

3.2. Copy data from the table *temperature* to the table *teplota_part* and make a partitioning by month. Before it, you will have to enable partitioning by commands:

```
set hive.exec.dynamic.partition=true;
set hive.exec.dynamic.partition.mode=nonstrict;
```

3.3. Look for the table *temperature_part* on HDFS under `/user/hive/warehouse/database_name.db` and watch how it is divided by partitioning.

4. Hive querying

We will work with the table *temperature*.

4.1. How many unique stations are in table? (457)

4.2. Which station is the most northern? (*USW00027502, BARROW POST ROGERS AP*)

4.3. Which state has the highest number of unique stations? (*TX*)

4.4. For how many records from the North Dakota (ND) is temperature under -10°C ? (8 446)

4.5. Which state has the highest average temperature in summer months (6, 7, 8)? (*MH, 28.1*)

4.6. Which states have difference of longitudes between the most western and the most eastern station bigger than 8 degrees? (*AK, FM, MT, TX*)

4.7. For every station with elevation over 1 500 m find a difference between total average temperature on this station and total average temperature of the station's state.

station	difference
USW00003103	-9.732538702083291
USW00023225	-4.890214469033042
...	...