





ACTRIS
CCRES

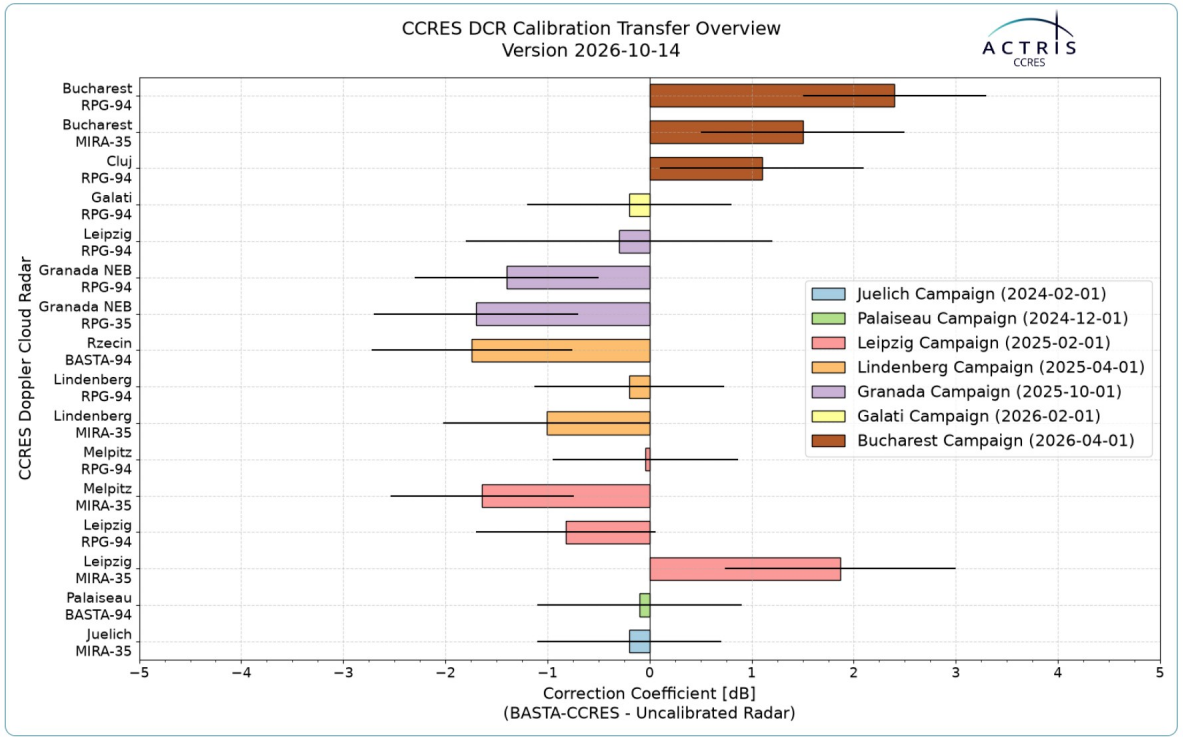
Monitoring cloud radar calibration with disdrometers

Thibault Delahaye[1], Jean-Charles Dupont[2], Jean-François Ribaud[2], Marc-Antoine Drouin[1], Martial Haeffelin[2]

[1] CNRS · [2] Institut Pierre-Simon Laplace

CCRES/CLU Autumn Workshop, Katowice (Poland), October 21, 2026

The network's calibration, and the years in between



N.B.: From the Granada campaign (October 2025) onwards, the coefficients and their uncertainties are *preliminary*: the analysis is well advanced, but not final.

A travelling reference

The CCRES BASTA is taken to a station and run beside the host radar. Comparing the two gives that radar one correction coefficient - an absolute level, tied to a single well-characterised instrument.

Seven campaigns so far

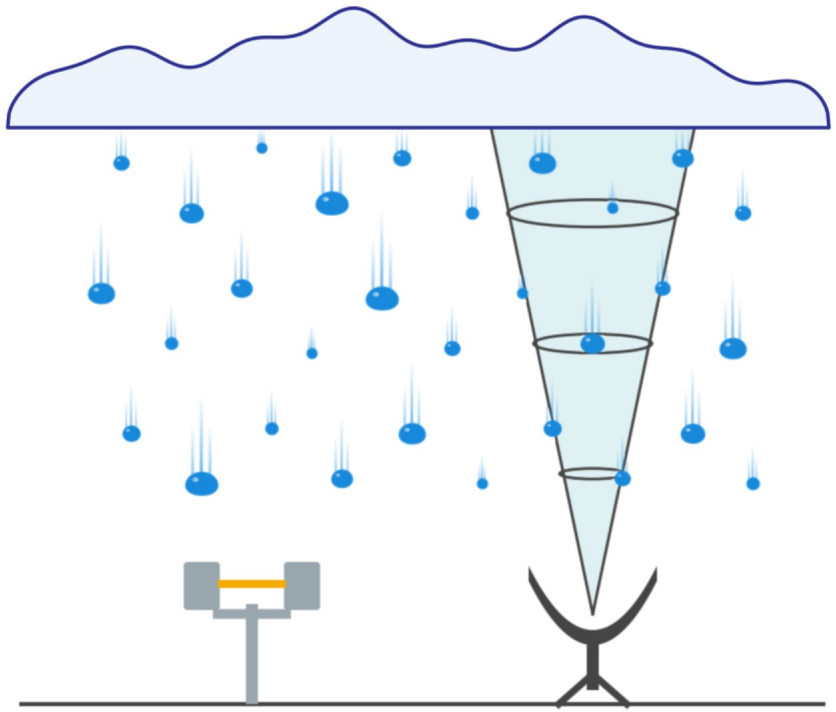
From Jülich in February 2024 to Bucharest in April 2026, sixteen radar comparisons, with corrections spread from -1.8 to +2.4 dB. This is the reference the network is calibrated on.

Four to five years before the next visit

A campaign lasts a few weeks, and at the scale of the network a station's turn comes round every four or five years. It fixes the level on the day it is measured, and nothing afterwards tells the station whether it still holds. **So what happens in between?** The stations are equiped with disdrometers, to follow the level remotely until the next campaign.

Method
Event selection
Results
What we need

Between campaigns: a reference already on site



disdrometer **cloud radar**

Drop size near the surface gives the reflectivity a radar should see. The comparison can be made whenever it rains - steady, stratiform rain is where it is worth trusting.

The level drifts after the visit

An ageing radome, a component replaced, work done on site: the correction measured during the campaign slowly stops being true, and nothing in the radar's own data says by how much.

Raindrops are a reference

A disdrometer measures the size of every drop reaching the surface - an instrument most stations already run. Reflectivity is the sixth moment of that size distribution, so those sizes give the reflectivity a radar should report just above it.

Every rain event is a check

Comparing the two, event after event, turns the level into something followed continuously: between campaigns, without a visit, and for as long as it rains.

What ΔZ is, and what it is not

$$\Delta Z\,(t) = Z_{DCR}\,(z_{min},\, t) - Z_{DD}\,(t)$$

in dB - zmin : lowest altitude where the radar data are validated, 120 to 300 m depending on the station

ZDD

What the disdrometer saw at the ground, converted into the reflectivity a radar would return. It owes nothing to the radar.

ZDCR

What the radar reported at zmin, at the same minute, with whatever calibration it has that day.

ΔZ - three contributions

- the radar's calibration error
- the error of our forward model
- the sampling mismatch between the two

Only the first is what we want, and it cannot be isolated - so the **value** of ΔZ is not a calibration figure. The other two do not change while the model and the instruments stay in place, so any **change** in ΔZ comes from the radar.

Processing chain

Steps 1 and 2 build *ZDD* from the disdrometer alone - step 3 is where it meets the radar.

preprocess

1 Drop size distribution

Drop counts per diameter and fall-speed class, every minute.

32 × 32 classes on an OTT Parsivel², 22 × 20 on a Thies LNM, resampled onto an exact 1-minute grid.

2 Forward model (*ZDD*)

Each drop's reflectivity from scattering theory, summed over the distribution.

Fall speed Gun & Kinzer (1949), axis ratio Beard & Chuang (1987) fitted by Andsager et al. (1999), T-matrix via *pytmatrix*. *ZDD* is modelled at 10, 24, 35 and 94 GHz.

process

3 Comparison

ZDD is subtracted from the radar profile: $\Delta Z = ZDCR - ZDD$.

At one range gate, set per station between 120 and 300 m, and at the frequency nearest the radar's own.

4 Event statistics

One ΔZ per rain event, over the minutes that pass every check.

Median and quartiles, the number of qualified minutes, and the *ZDCR*–*ZDD* regression.

Processing chain

Steps 1 and 2 build *ZDD* from the disdrometer alone - step 3 is where it meets the radar.

preprocess

1 Drop size distribution

Drop counts per diameter and fall-speed class, every minute.

32 × 32 classes on an OTT Parsivel², 22 × 20 on a Thies LNM, resampled onto an exact 1-minute grid.

2 Forward model (*ZDD*)

Each drop's reflectivity from scattering theory, summed over the distribution.

Fall speed Gun & Kinzer (1949), axis ratio Beard & Chuang (1987) fitted by Andsager et al. (1999), T-matrix via *pytmatrix*. *ZDD* is modelled at 10, 24, 35 and 94 GHz.

process

3 Comparison

ZDD is subtracted from the radar profile: $\Delta Z = ZDCR - ZDD$.

At one range gate, set per station between 120 and 300 m, and at the frequency nearest the radar's own.

4 Event statistics

One ΔZ per rain event, over the minutes that pass every check.

Median and quartiles, the number of qualified minutes, and the *ZDCR*–*ZDD* regression.

Workflow — from CloudNet files to a long-term series

Main inputs - one CloudNet file each per day

Cloud radar
reflectivity and Doppler

Disdrometer
drop counts by size and speed

Optional input
Weather station
4 extra filters and the rain-gauge check

preprocess resample to 1 minute · forward-model *ZDD* · merge the sources

preprocessed .nc one file per station-day - *ZDCR*, *ZDD* and the drop size distribution

process cut rain events · quality checks · ΔZ at the gate · event statistics

processed .nc one record per rain event: ΔZ, quartiles, flags

from the preprocessed and the processed files
quicklooks
a daily figure for every processed day, rain or not · one figure per selected rain event

from the processed files
long-term monitoring
every event of the record pooled, per station and radar

Code: github.com/ACTRIS-CCRES/ccres_disdrometer_processing · Doc: ccres-disdrometer-processing.readthedocs.io

Workflow — from CloudNet files to a long-term series

Main inputs - one CloudNet file each per day

Cloud radar
reflectivity and Doppler

Disdrometer
drop counts by size and speed

Optional input
Weather station
4 extra filters and the rain-gauge check

preprocess — resample to 1 minute · forward-model *ZDD* · merge the sources

preprocessed .nc — one file per station-day - *ZDCR*, *ZDD* and the drop size distribution

process — cut rain events · quality checks · ΔZ at the gate · event statistics

processed .nc — one record per rain event: ΔZ, quartiles, flags

from the preprocessed and the processed files
quicklooks
a daily figure for every processed day, rain or not · one figure per selected rain event

from the processed files
long-term monitoring
every event of the record pooled, per station and radar

Code: github.com/ACTRIS-CCRES/ccres_disdrometer_processing · Doc: ccres-disdrometer-processing.readthedocs.io

Workflow — from CloudNet files to a long-term series

Main inputs - one CloudNet file each per day

Cloud radar
reflectivity and Doppler

Disdrometer
drop counts by size and speed

Optional input
Weather station
4 extra filters and the rain-gauge check

preprocess — resample to 1 minute · forward-model *ZDD* · merge the sources

preprocessed .nc — one file per station-day - *ZDCR*, *ZDD* and the drop size distribution

process — cut rain events · quality checks · ΔZ at the gate · event statistics

processed .nc — one record per rain event: ΔZ, quartiles, flags

from the preprocessed and the processed files
quicklooks
a daily figure for every processed day, rain or not · one figure per selected rain event

from the processed files
long-term monitoring
every event of the record pooled, per station and radar

Code: github.com/ACTRIS-CCRES/ccres_disdrometer_processing · Doc: ccres-disdrometer-processing.readthedocs.io

Selecting the "good" rain event

Convective events are too intense and variable in space and time. The criteria below keep the long, steady, stratiform events.

Event definition

Criterion	Value
Minimum rain accumulation	3 mm
Minimum rain duration	180 min
Dry interval closing it	60 min

Built on the disdrometer's own rain rate, so these three criteria apply at every station. The next slide adds what a weather station brings.

Per-minute checks

Criterion	Value
Rain rate	< 3 mm/h
Fall speed vs Gunn & Kinzer	within 30%

The rate limit keeps attenuation and non-Rayleigh effects small. The fall speed is compared with the Gunn & Kinzer law for liquid drops, which rejects snow and melting particles as well as wind-driven and splashing drops.

An event enters the statistics only with **at least 50 qualified minutes** and a finite median ΔZ. The criteria trade the number of events against the reliability of each median.

Currently, identical thresholds at all configurations of the network; only the comparison range gate is station-specific.

Additional checks with a weather station

These come on top of the disdrometer checks of the previous slide.

Variable	Accepted range	Process excluded
Air temperature	> 2 °C	snow, sleet, melting particles
Wind speed	< 10 m/s	vertical and stable precipitation
Wind direction	±45° of the axis	manufacturer recommandation
Relative humidity	80 - 95%	evaporation below the gate, and fog

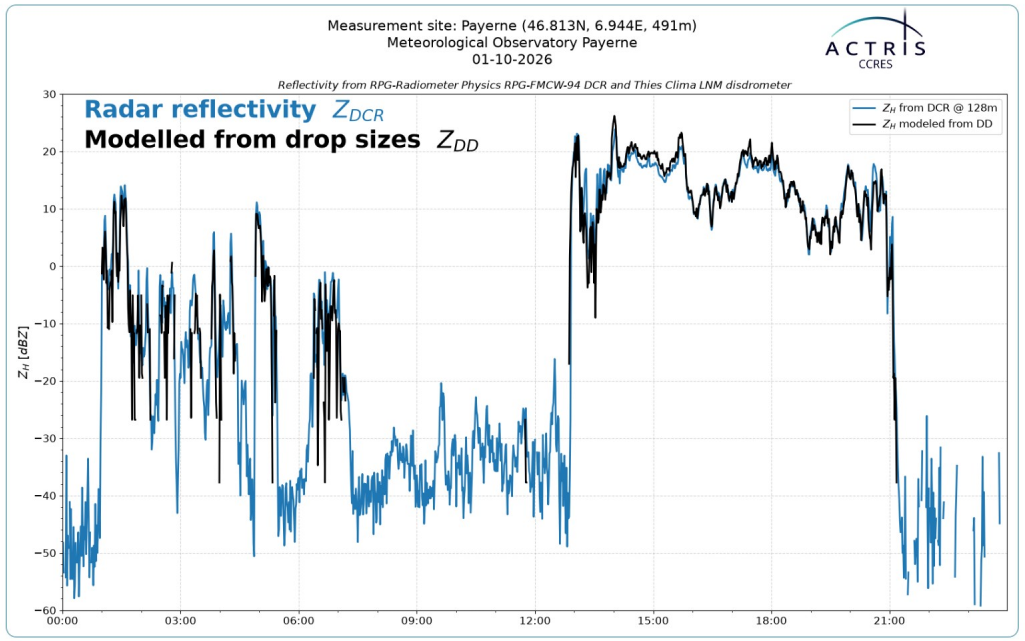
The strongest check: rain gauge against disdrometer

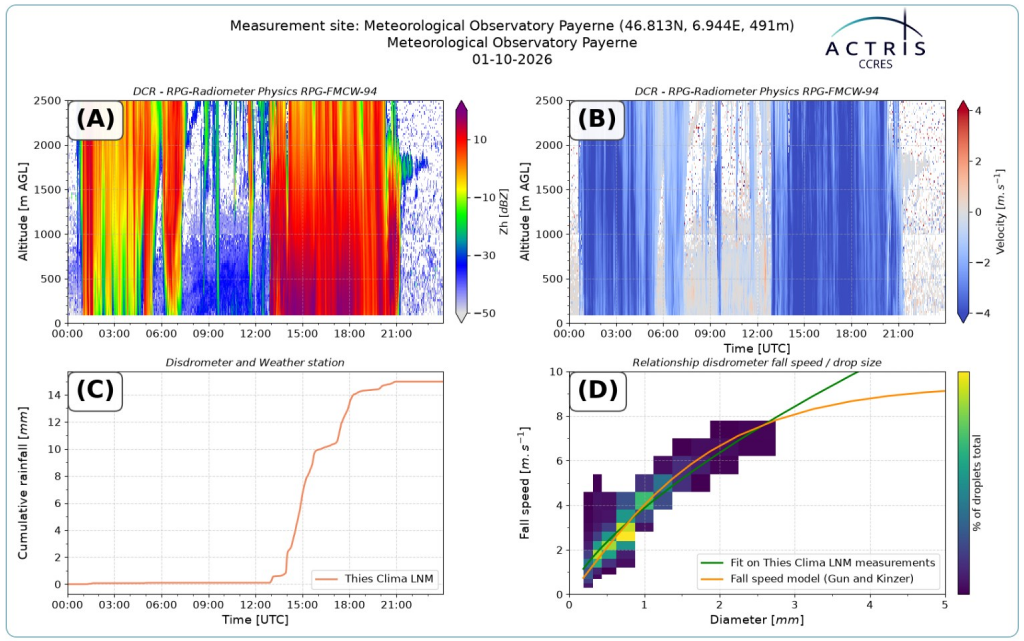
The two must agree on the accumulated rain to within 30% over the event - the only independent check we have on the reference itself. When they disagree, the problem is the disdrometer, not the radar, and the event is dropped.

Without a weather station — events come ONLY from the disdrometer rate and these five tests are not "computable". The series stays usable, its uncertainty less constrained.

Work in progress — Every station that has a weather station is now processed both ways, with and without, so we can measure what each filter actually removes and revise the values on that basis.

What comes out every day





Reflectivity overview

Z_{DCR} at the comparison gate and Z_{DD}, over the 24 hours of the day.

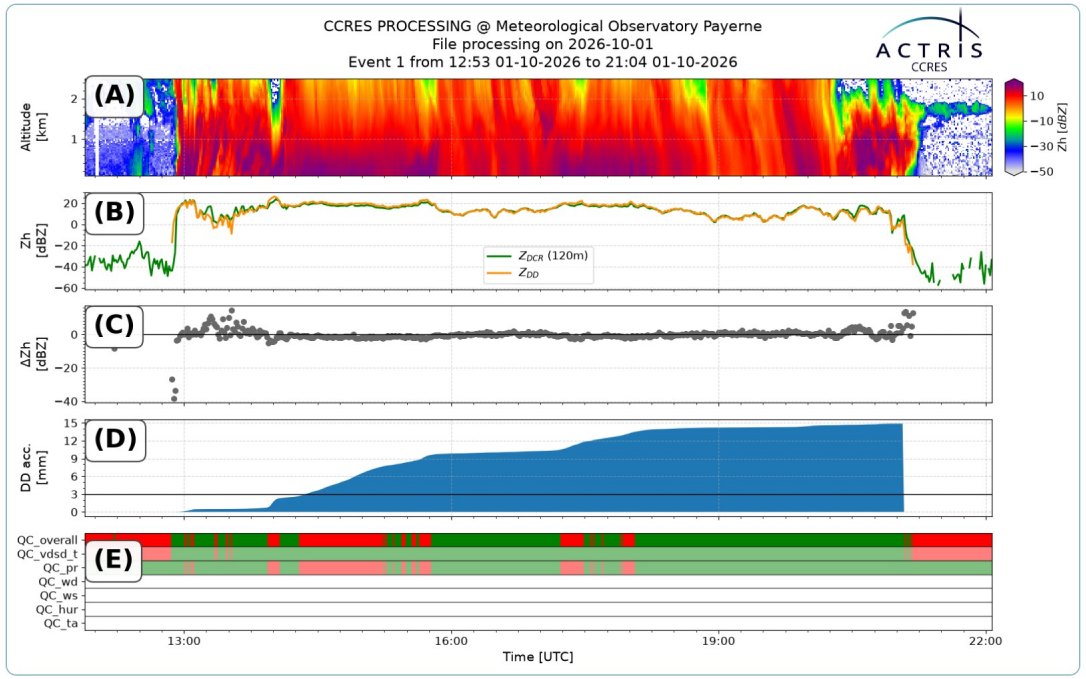
Instrument overview

(A) Radar reflectivity
(B) Doppler velocity profiles
(C) Disdrometer rain accumulation
(D) The measured fall speed against drop diameter

Example: Payerne, 1 October 2026. Both figures are produced for every processed station-day, beside the netCDF, and published at **ccres.aeris-data.fr/en/data-visualization/**

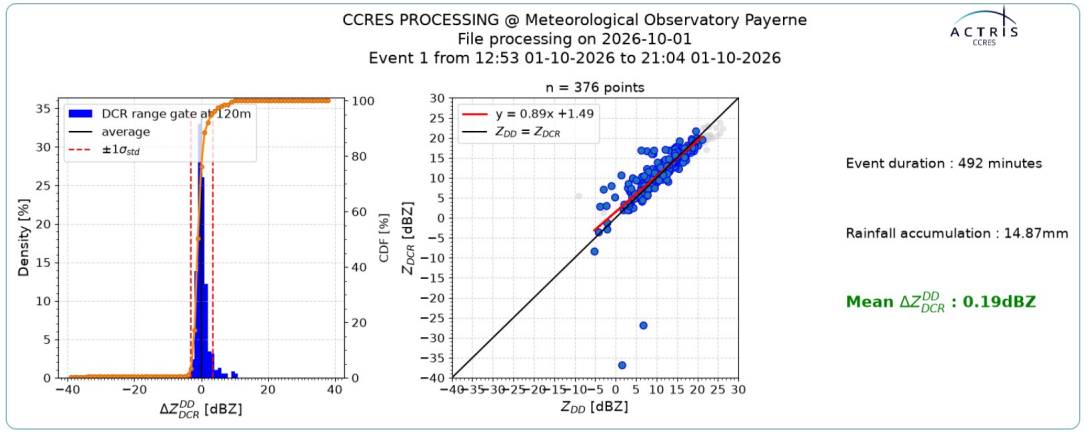
Method
Event selection
Results
What we need

What comes out for every rain event



Event detail

(A) The reflectivity profile
(B) Z*DCR* and Z*DD* at the gate
(C) The ΔZ difference
(D) The disdrometer accumulation
(E) The quality checks (one row per criterion).

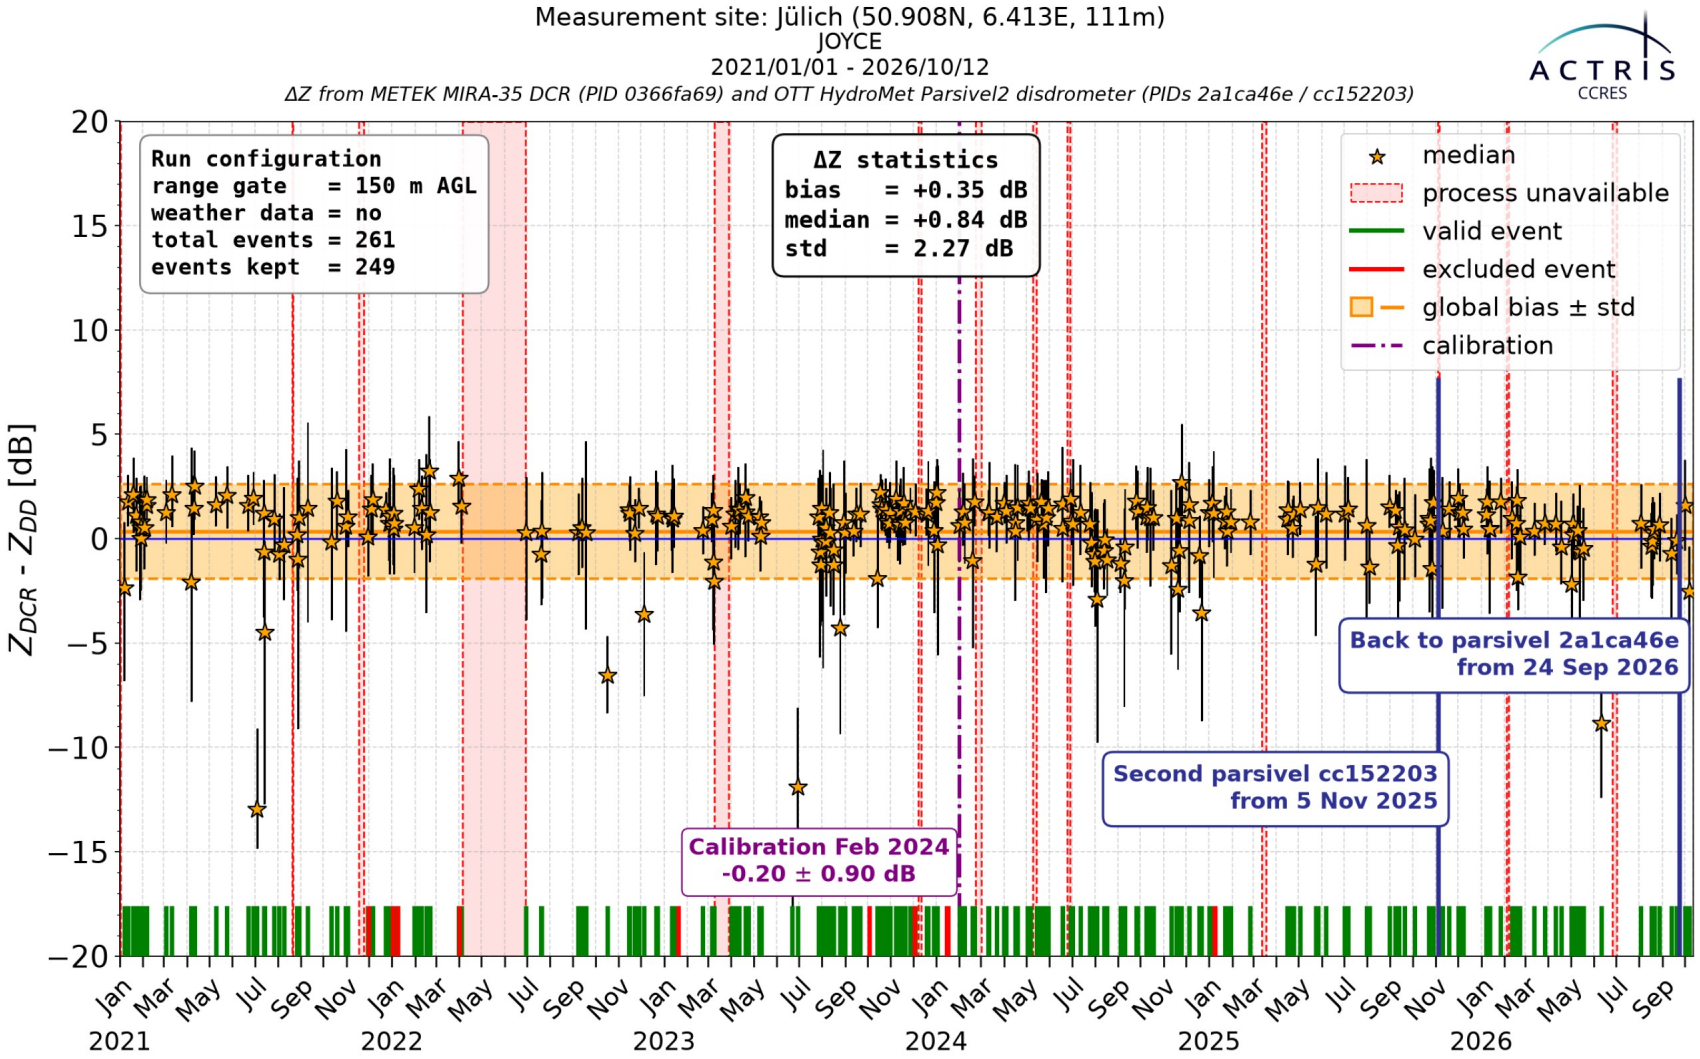


Event summary

The ΔZ distribution, the Z*DCR*-against-Z*DD* regression, and the statistics: rain duration, accumulated precipitation, mean ΔZ.

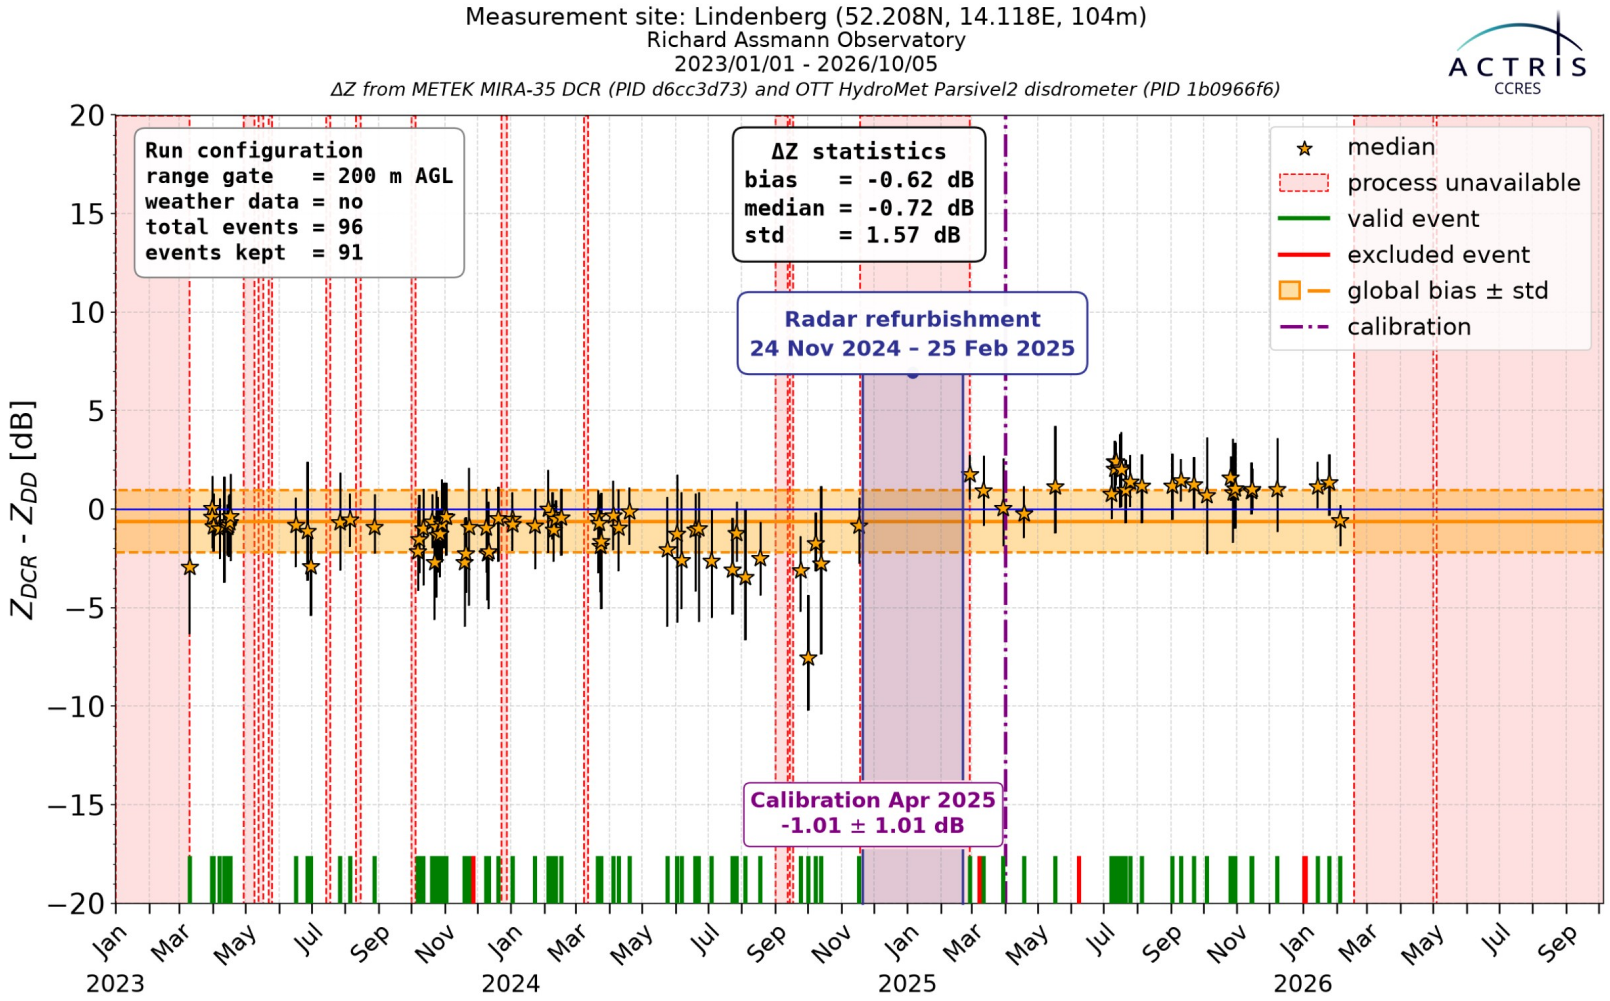
Browse them all at ccres.aeris-data.fr/en/data-visualization/

A stable series: Jülich, six years



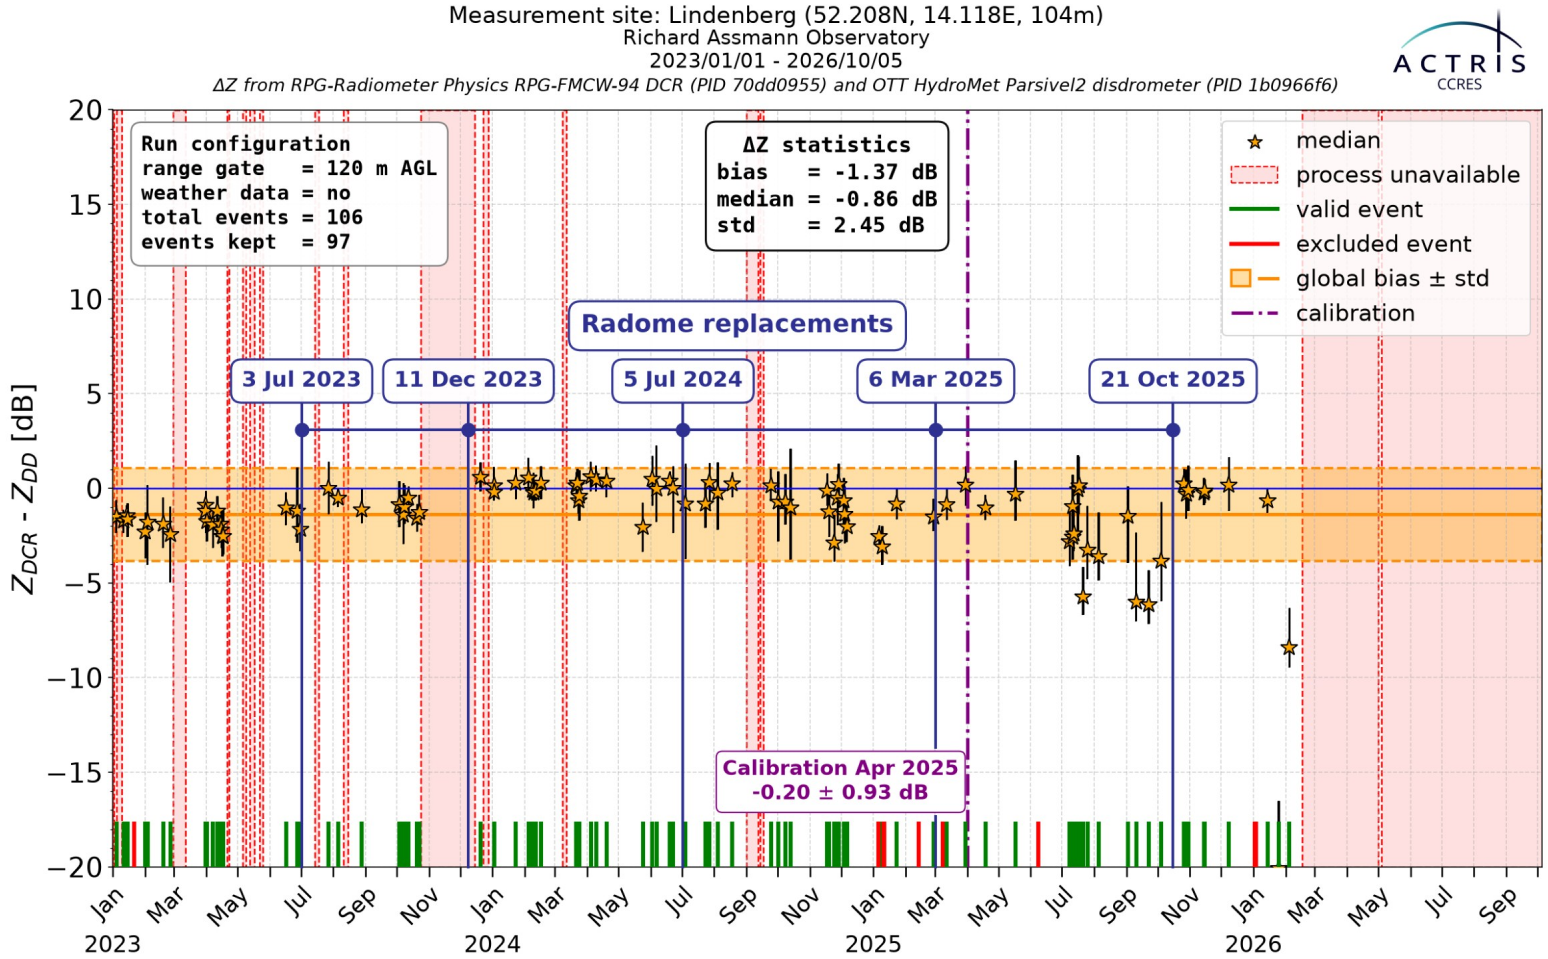
METEK MIRA-35 against two successive Parsivel², 239 qualified events out of 260: bias **+0.37 dB**, standard deviation **2.31 dB**. The two disdrometers agree to 0.2 dB across the November 2025 change.

Before and after a radar refurbishment



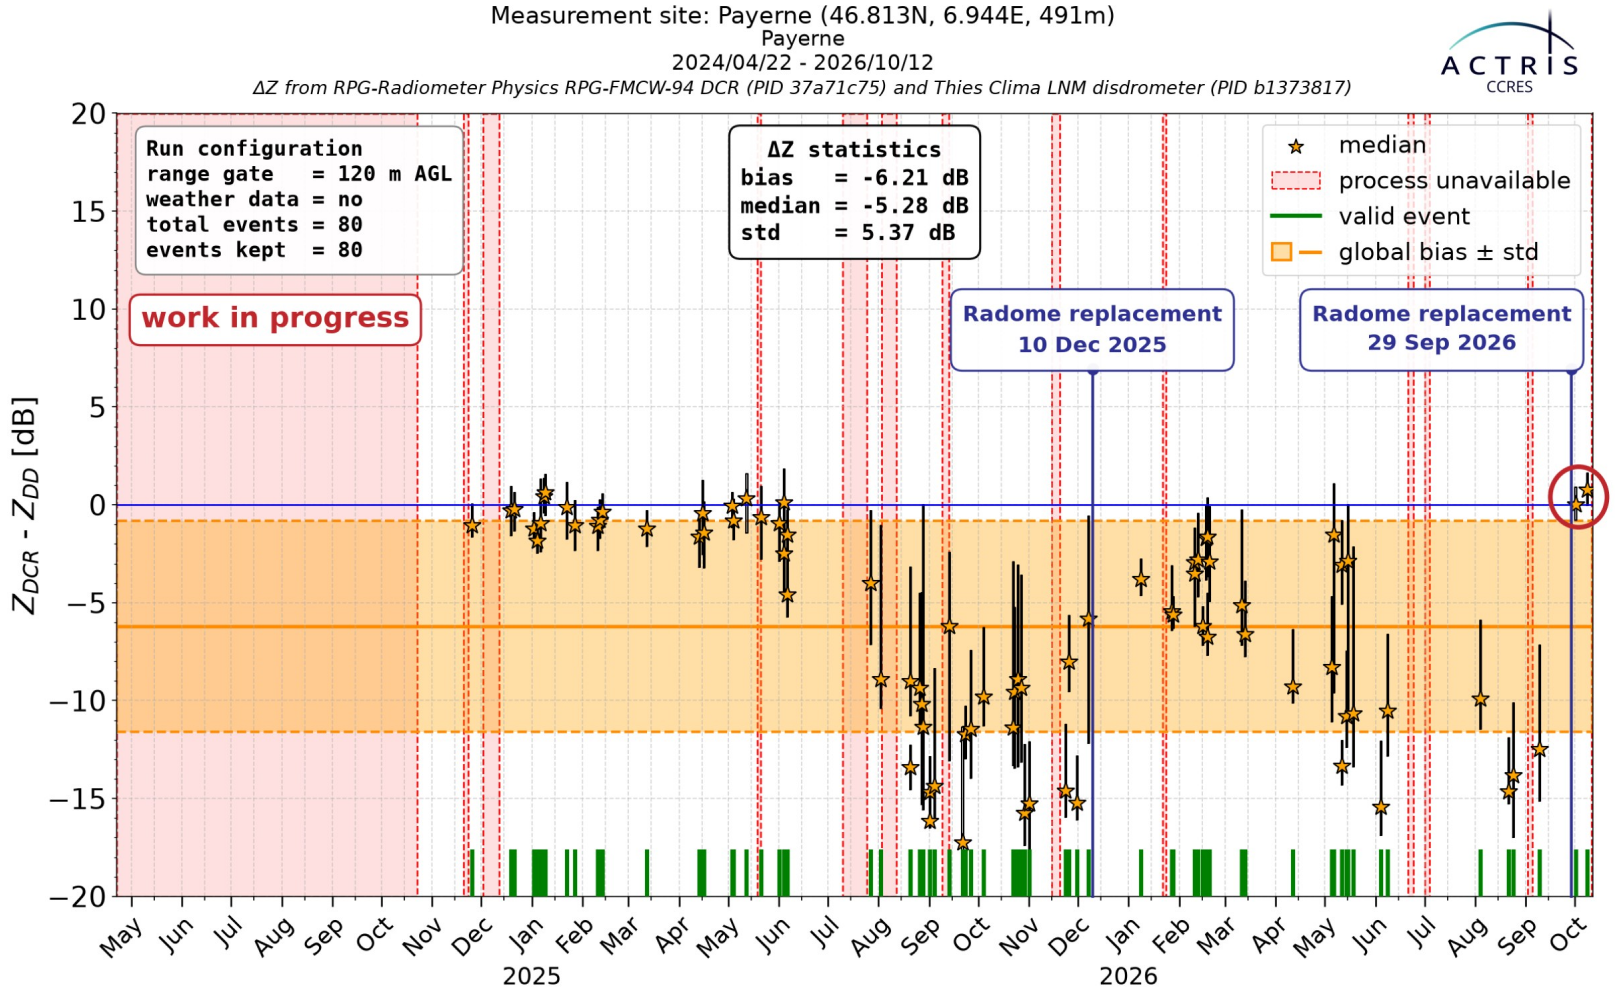
Lindenberg, METEK MIRA-35 against the Parsivel², 91 qualified events out of 96. Before the refurbishment the medians sit below the global bias, after it consistently above - a shift of roughly 2 dB, invisible in any single event. The refurbishment included a recalibration, so a step is expected; what the series adds is the confirmation that the radar held steady on either side of it.

The series sees the radome ageing



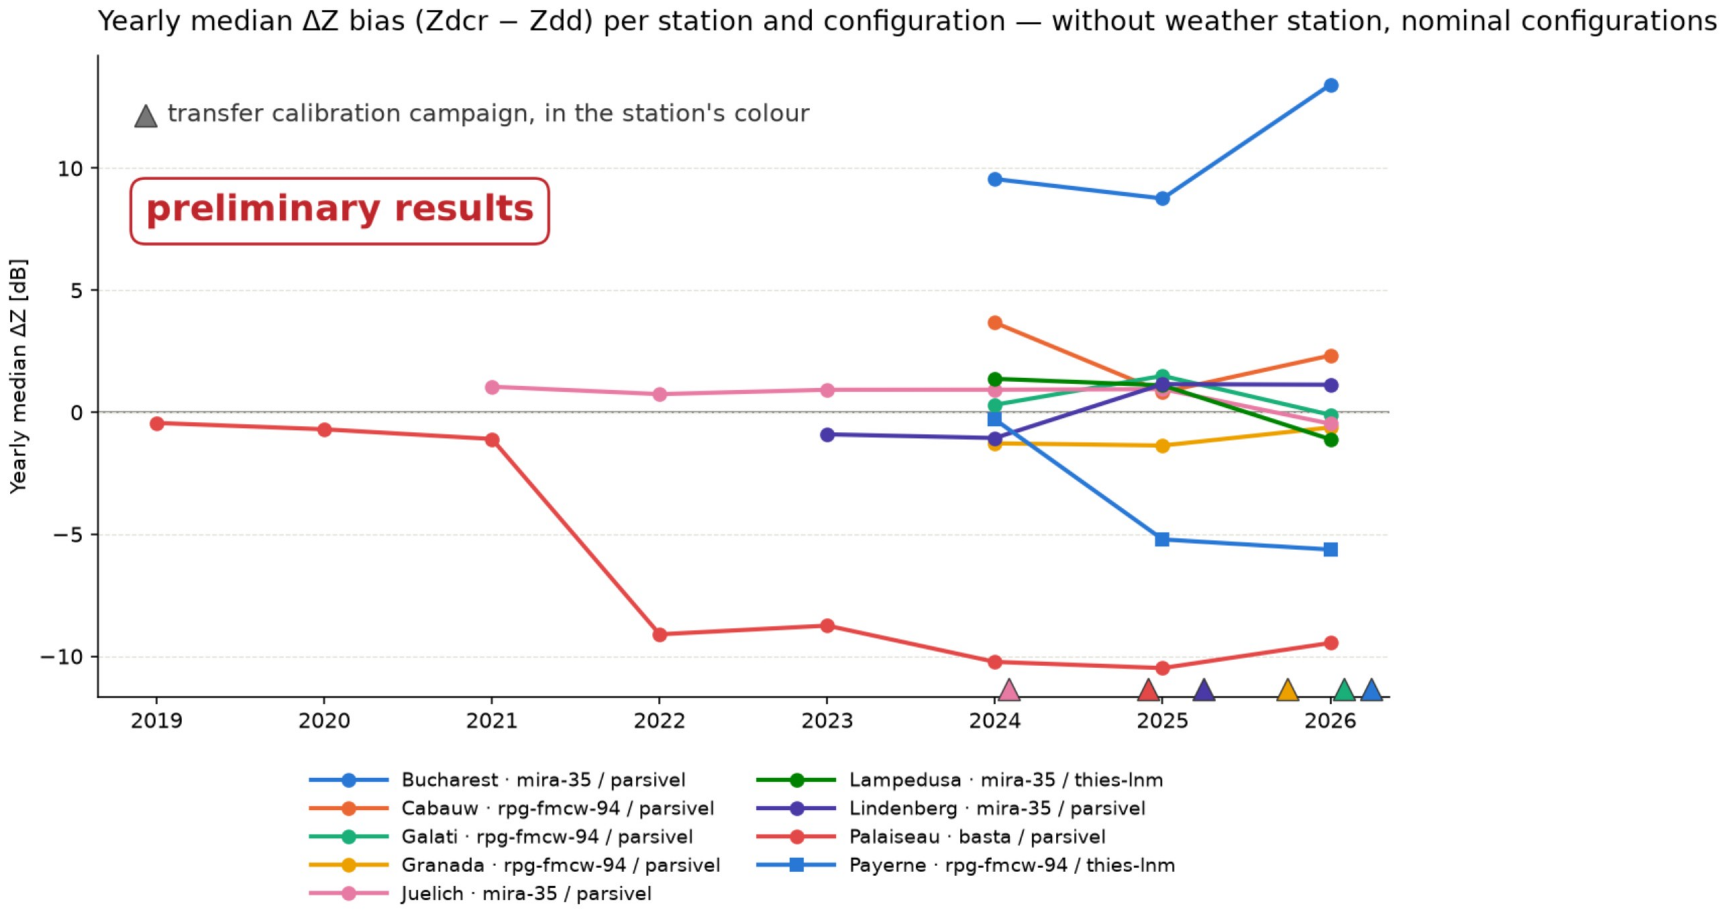
Lindenberg, RPG-FMCW-94. Five replacements since 2023. The first four leave the series flat - nothing had degraded yet. Before the fifth, the medians drift to -6 dB through the summer of 2025. The series does not measure the intervention, it measures the ageing that preceded it: a radome changed before it has degraded leaves no trace, and that is what a replacement on schedule looks like.

How fast it ages is not the same everywhere



Payerne, RPG-FMCW-94. The radome ages faster here than at Lindenberg (15 dB against 6 dB). Two technical failures sit inside the same series, however - a circuit breaker in the summer of 2025, the blower in the winter - so how much of the loss belongs to the radome alone is still being analysed. Either way, no single replacement interval fits every station: each one has to be read on its own series.

Yearly medians across the network



- DCR vs DD analysis at **6 stations** shows good consistency (bias < 2 dB)
- **2 stations** need further investigation to understand the discrepancy
- **Palaiseau:** stable since 2022, but bias likely due to liquid water attenuation

What helps most

Log what you do to the instruments

A date is enough: radome, calibration, repair, move, or a switch to a backup instrument. With it, a step in the series becomes a measurement, and we can tell you what it was worth.

The disdrometer and the weather station count too - and one entry per intervention, rather than several gathered into one note.

The series warns before a radome costs dB

At Lindenberg and at Payerne the loss built up over months, and the series shows it building - early enough to plan the replacement, and long before it would show up anywhere else.

Perspectives: evaluating the impact of the weather stations

A weather station adds four filters on temperature, wind and humidity, plus the rain-gauge cross-check. Fewer events survive, and the annual bias comes out better constrained. What each filter removes, and how far it moves the median, is what we are quantifying next.

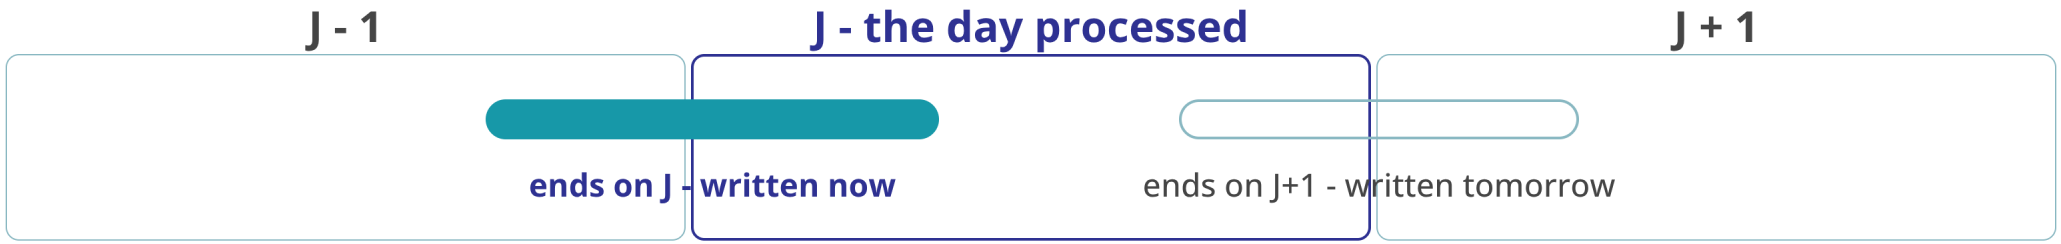
ACTRIS
CCRES

Thank you !

CCRES/CLU Autumn Workshop, Katowice (Poland), October 21, 2026

Backup — rain does not stop at midnight

Processing day J alone would cut every event crossing 00:00 in two, and both halves would then fail the 180-minute and 3 mm criteria. So the chain always reads **J-1, J and J+1** and concatenates them before any event is cut.



Three days are read

The three preprocessed files are concatenated in time before anything is cut.

Events stay whole

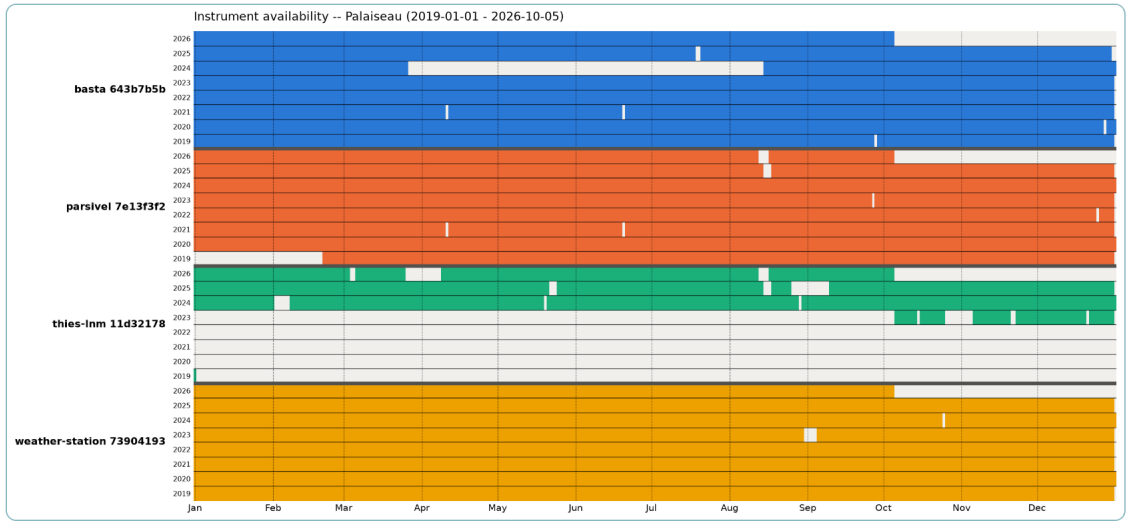
An event keeps its real duration and accumulation, so it is judged on what it was.

Written once

Only events ending on J are kept, so nothing is counted twice when J+1 runs.

Backup — knowing what exists before reading any bias

An interruption and a flat series look alike on a figure. Two calendars are produced for every site, so that a missing month is seen as a missing month. Example: Palaiseau, 2019-2026.





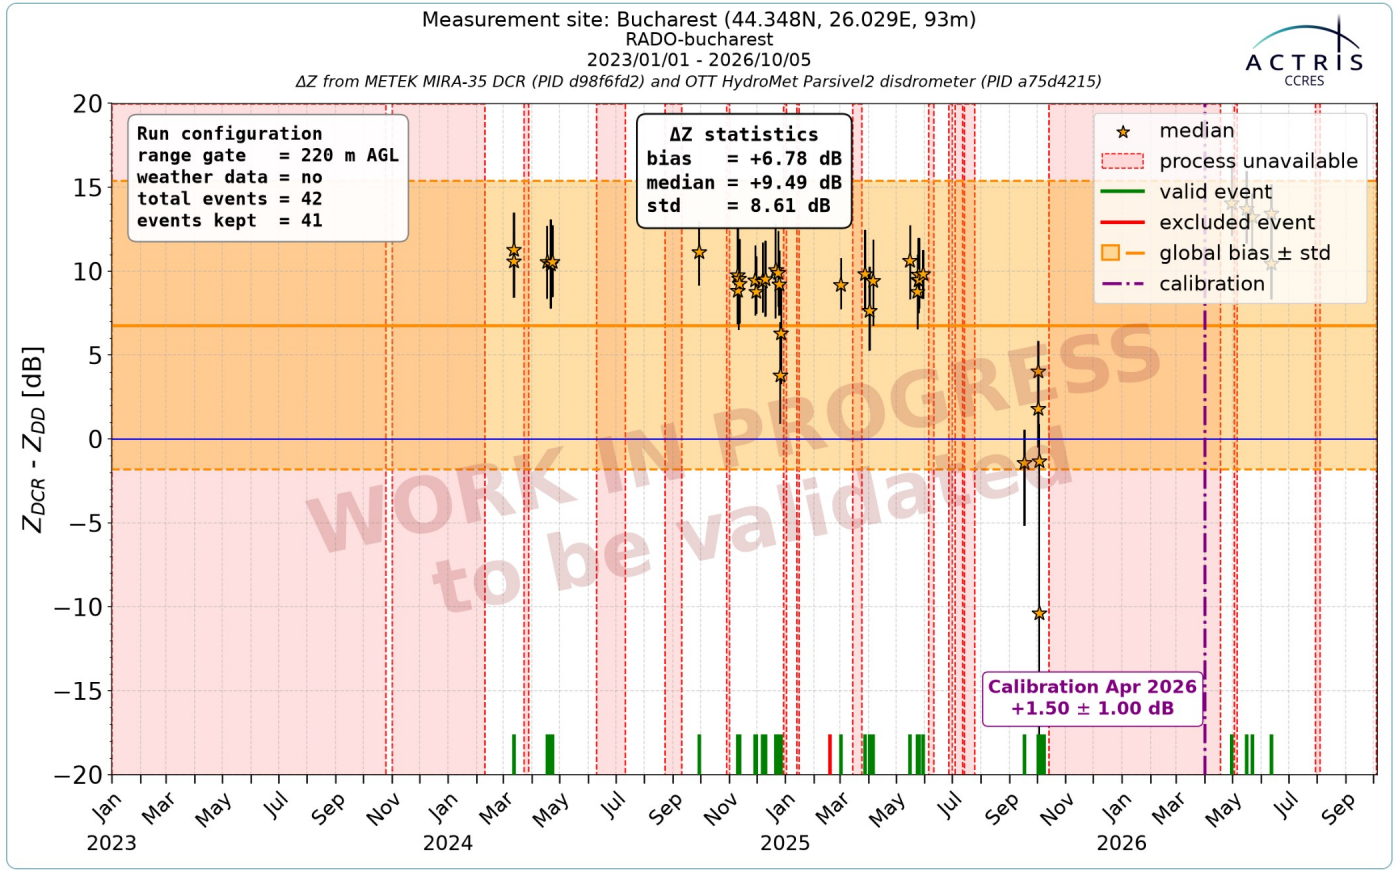
What the station delivered

One row per instrument and per year, built from the CloudNet files themselves. Each physical instrument keeps its own colour across the years, so an arrival, a replacement or a move stands out without reading the labels - here the Thies starting in late 2023.

What the chain produced

One row per radar–disdrometer pair, preprocessing and processing shown separately. This is what separates a day we never received from a day that failed to run - two problems with two different fixes.

Backup — Bucharest, where the record is too thin



Bucharest, METEK MIRA-35 against the Parsivel². 41 qualified events out of 42 in four years: bias **+6.78 dB**, standard deviation **8.61 dB**. Most of the record carries no processed file at all, and the handful of events of autumn 2025 sit some 10 dB below the earlier ones - with months between events, nothing says when that happened or whether it has held since. The April 2026 transfer calibration reads +1.50 ± 1.00, far from the ΔZ level: a gap this size is what we would want the continuous series to explain.