Instrumentation health audit

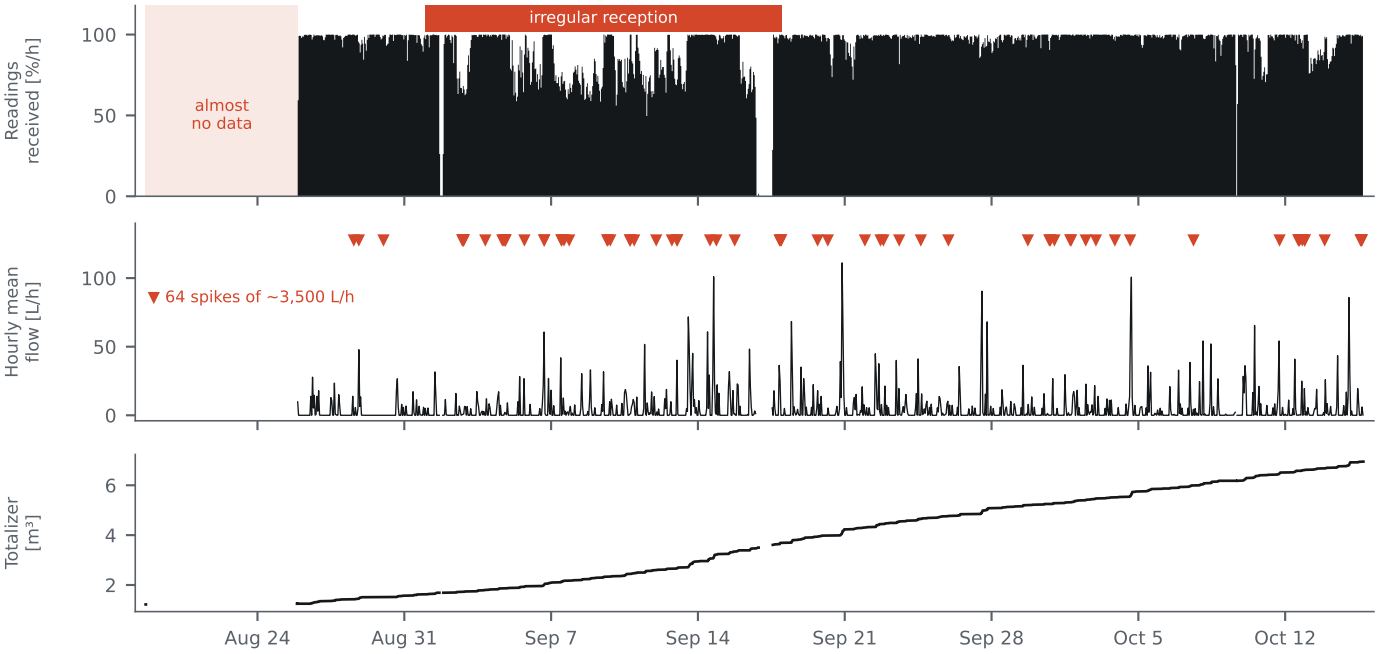
Sample report · Ultrasonic flow meter (RS-485) · real installation, anonymized data · Aug 18 – Oct 15, 2026 · 1 reading per second

93 %
of expected readings arrived (from Aug 25)

6
long outages, 25 h in total

64
~3,500 L/h spikes lasting seconds

0.98
flow–totalizer match, hours without gaps

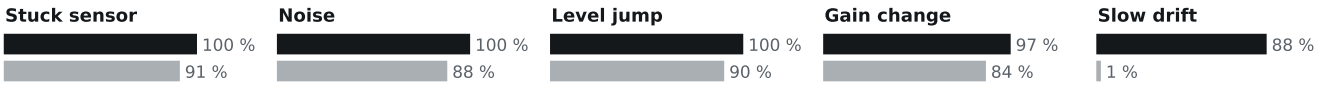


	Finding	What to do
1	Aug 18–25: almost nothing was logged, only a few minutes of the totalizers. Flow and temperature start on the 25th at 22:24.	Check the logger configuration at start-up.
2	Sep 1–17: irregular reception (most days 70–88 % of readings) and a ~16 h outage on the night of the 16th. On the worst days, up to 6 in 10 readings arrived in delayed batches, stored with the time they arrived instead of the time they were measured.	It is the acquisition system, not the meter: the totalizer kept counting. On those days the exact second of each reading can't be trusted.
3	64 spikes of 3,500–3,635 L/h that last a few seconds and start and end at zero flow. The forward totalizer does not count them, but 12 of the 16 steps of the reverse counter happen right as a spike starts.	They may be real backflow (water hammer, a failing check valve) or a meter artifact. Inspect the check valve and the meter's settings; meanwhile, keep those readings flagged apart.
4	Without those spikes, flow follows the totalizer hour by hour (r ≈ 0.98 in hours without data gaps).	The meter measures forward flow correctly; most problems are in the data chain.

With many sensors: which sensor is failing

We tested the audit on a real array of 9 strain gauges (24 h, 1 Hz) with 2,560 simulated faults, one at a time. A regression between sensors finds the fault and points to the sensor; a second check, which looks at whether that sensor stopped moving in step with its neighbours, confirms it.

1,639 / 1,639 times both checks pointed to the same sensor, and it was the right one, every time.



In the report, every faulty sensor is marked confirmed or probable. Slow drift is caught by regression alone, so it shows up as probable. With faults half this size, the right sensor still comes up 91 % of the time; drift drops to 58 %. Simulated faults on 24 h of one real array; every audit is checked against your own records.