RefrigSense
PREDICTIVE REFRIGERATION INTELLIGENCE

Validating Fault Detection Before the First Install

Certification of the RefrigSense detection engine across 200 simulated store-quarters

Public white paper · v1.0 · July 2026

Executive summary

We stress-tested the RefrigSense detection engine against 200 independent simulated stores — nine refrigeration units each, run for 90 days — with 1,400 equipment problems of seven kinds planted against a known answer key. None of the test scenarios were used during development, and the engine's configuration was locked and fingerprinted before testing began.

The engine produced a correct alert on 1,397 of the 1,400 planted problems, and caught 100% of the 1,047 problems that went on to cause a failure — every one of them, before the failure happened. Slow-developing faults were flagged one to four weeks ahead, and false alarms on healthy equipment ran about one per store every eight months.

These are simulation results, and we label them that way: they certify that the engine works as designed. Field pilots with technician-confirmed faults are underway to turn them into real-world claims.

1 · Why we test this way

Field data cannot certify a detection engine at this stage, for a simple reason: you cannot plant a known compressor failure in a customer's walk-in and wait to see whether the software notices. Simulation can. A simulated fleet allows faults to be injected with a precise answer key — what failed, on which unit, when it began, and how fast it progressed — so that every alert can be graded against ground truth. It also reaches a scale no pilot could: 200 store-quarters is roughly fifty years of single-store operation, compressed into one certification run.

The trade-off is honesty about what is being measured. Simulation certifies the engine's design against fault physics; only field results certify real-world performance. This paper is the first of those two steps, and every figure in it is labeled accordingly. It also reflects a product principle: the system proves itself before it reaches a customer's kitchen, not after.

2 · What was tested

The system under test is the RefrigSense deterministic detection layer — the physics-based rules and safety thresholds that monitor every installation from day one. It reads the vital signs each equipment class

exposes: refrigerant pressures and line temperatures on remote condensed systems, and compressor current, cabinet temperature and humidity, door activity, and leak detection on sealed units. The machine-learning anomaly layer that runs alongside it in production was not required to achieve the results below; it is expected to add earlier and subtler detection as it matures on each unit's own history.

3 · Methodology

Fleet. 200 independent virtual stores, each with nine refrigeration units — a representative mix of walk-in coolers and freezers, reach-ins, prep tables, and ice machines — operated for 90 simulated days: 162,000 unit-days in total.

Planted problems. 1,400 problems across seven fault classes: refrigerant leak, dirty condenser, compressor hard-start, ice-machine mineral scale, defrost failure, door left open, and unit dead — both slow-developing and sudden faults.

Randomization. Every store was drawn differently: equipment behavior, weather, fault onset timing, progression rate, and sensor noise.

Hold-out discipline. None of the 200 test scenarios were used while the detection rules were being developed.

Lock and fingerprint. The engine's configuration was frozen and fingerprinted (f8a8f61fb46d) before the test began; nothing was tuned after seeing results.

Grading. Every alert the engine produced was scored against the answer key — the full record of what was planted, where, and when.

A note on scope: the parameters and physics of the fleet simulator are proprietary. This paper describes the experiment and its grading; it does not describe the simulator's internals.

4 · Results

1,400	1,397	100%	1 / 8 mo
equipment problems planted	correctly alerted — on scenarios never seen before	of failures caught before they failed (1,047 of 1,047)	false alarms per store on healthy equipment

Of the 1,400 planted problems, 1,047 grew severe enough to cause a failure inside the 90-day window. The engine caught all 1,047 before the failure — 100%. The remaining 353 problems were still developing when the window closed, and 350 of those had already been flagged as well. With this many trials, the statistics guarantee a true catch rate of at least 96% for every individual fault type, and at least 99.6% overall, even after accounting for luck.

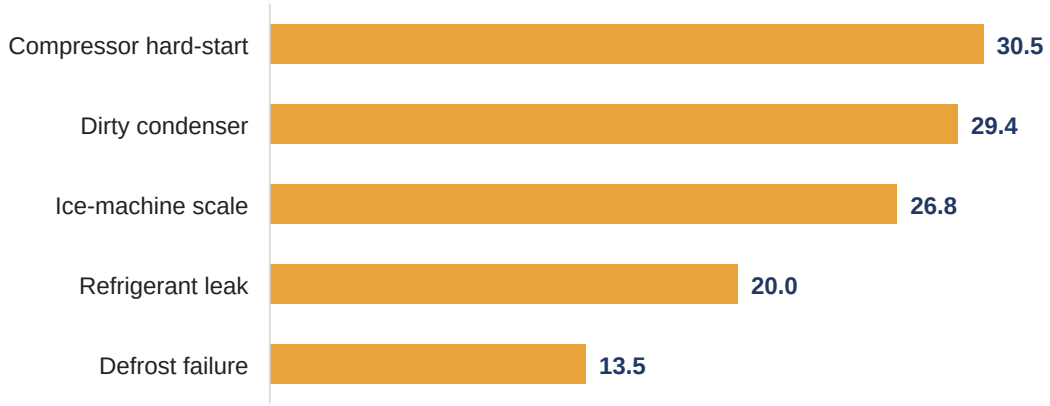

Average days of warning before failure, by fault type — certified test data

The problem	Caught	Average warning	Median warning	Usually
Compressor hard-start Start components wearing out — the compressor strains harder every start	118 of 118	30.5 days	29.5 days	26–34 days ahead
Dirty condenser Coils clogged with dust and grease, making the whole system work harder	110 of 110	29.4 days	29.0 days	25–34 days ahead
Ice machine scale Mineral buildup slowing every freeze cycle	101 of 101	26.8 days	27.0 days	24–29 days ahead
Refrigerant leak The system slowly losing its charge	128 of 128	20.0 days	19.0 days	16–23 days ahead
Defrost failure The evaporator icing up because defrost stopped working	190 of 190	13.5 days	13.0 days	11–16 days ahead
Door left open Door standing open, warm air pouring in	200 of 200	Same day	Same day	Alerted within minutes
Unit dead Compressor stopped or power lost	200 of 200	Same day	Same day	Alerted within hours

Averages and medians come straight from the certified 200-run dataset. The median sits within half a day of the average for every fault — a sign of clean, consistent results. “Usually” shows the middle half of cases.

5 · False alarms — two honest ways to count

Counted strictly, 88% of all alerts matched the exact planted problem; the remainder works out to about one non-matching alert per unit every 17 months. Most of those fired on equipment that genuinely had a problem — just a different one than the alert named, such as a door alert during a defrost failure. Alarms on perfectly healthy equipment: about one per store every eight months.

6 · What these numbers are — and are not

These numbers certify the detection engine against simulated fault physics — proof the design works, measured the hard way, on fresh test conditions it was never tuned on. They are not field performance.

Real-world claims come from pilot deployments, where technicians confirm each fault the system called; roughly fifty confirmed field faults is the point where the same statistics begin backing field numbers. Those pilots are underway now.

RefrigSense is a decision-support tool. It can help reduce downtime and cost by surfacing early indicators of developing faults; it does not guarantee detection of every failure, and the figures in this paper are certified test results, not warranted performance.

7 · About RefrigSense

RefrigSense Monitor is a predictive refrigeration monitoring platform for multi-site food operators. Battery-powered wireless sensors read each machine's vital signs, a site gateway with its own cellular connection carries them to the cloud, and the software flags developing faults weeks before a temperature alarm would sound. The company was founded by a facilities executive who ran refrigeration programs for an 800-location national chain for more than twenty years.

The live product can be explored at **refrigsense.com** — the demo requires no signup. The full technical validation report is available to qualified prospects under NDA.

Contact: drew@refrigsense.com · refrigsense.com