IoT Technologies

iottechnologies.co.uk

WHITEPAPER

Heat: the silent killer of unmanned infrastructure sites

Why equipment enclosures overheat, how fast a cooling failure escalates, what sustained heat does to electronics and standby batteries, and what a live national monitoring estate shows.

IoT Technologies Ltd · August 2026

Bournemouth, United Kingdom · Member of Made in Britain, No. 5442

Heat is the most common environmental threat to unmanned equipment enclosures, and the least visible. A telecoms cabin or equipment cabinet that loses its cooling does not announce the fact. Nothing trips, nothing alarms, and the equipment inside carries on working, at first. What changes is the rate at which it is being consumed. Electronics age faster in proportion to temperature, standby batteries lose service life on a well-documented curve, and by the time the consequence surfaces as a failed rectifier or a battery string that cannot hold up a site through a mains outage, the thermal event that caused it is weeks or months in the past and invisible. Continuous equipment cabinet temperature monitoring is the only practical way to see these events while they are still cheap to fix, and this paper sets out the engineering case for it: why enclosures run hot, what happens in the hours after a cooling failure, what sustained heat does to electronics and batteries, and what a live national monitoring estate shows about how often it actually happens.

01

Why equipment enclosures run hot

An equipment enclosure is a box that converts electricity into heat. Almost every watt drawn by the equipment inside a telecoms cabin, a roadside cabinet or a plant-room panel ends up as heat within the enclosure, because that is what electronics do with electrical energy once the useful work of switching, amplifying and transmitting is done. A modestly loaded telecoms cabin dissipating two kilowatts continuously is releasing the heat of a domestic fan heater running flat out, twenty-four hours a day, inside a sealed insulated box a few metres on each side.

Getting that heat out is a design problem solved at commissioning and then largely forgotten. Small cabinets rely on passive ventilation: louvres, natural convection, sometimes a heat-exchanger panel. Larger cabins carry forced ventilation or a small air-conditioning unit. All of these were sized for the equipment load and the climate assumptions of the day they were installed, and both of those assumptions drift. Equipment gets added over the years, upgrades draw more power than the kit they replaced, and the UK climate has moved against the design margin. The Met Office's UK climate projections and the run of recent summers point the same way: hot spells are becoming more frequent, more intense and longer. An enclosure specified in 2005 against a design ambient of the low thirties now sees those temperatures routinely rather than exceptionally, and a passive ventilation scheme that once held a comfortable margin may now be operating at the edge of it every July.

The failure modes are mundane. Fans are the moving part in an otherwise solid-state system, and they fail: bearings wear, capacitors in the fan supply dry out, filters block with dust and pollen until airflow collapses even though the fan still spins. Air-conditioning units lose refrigerant or trip on their own protection. Louvres get painted over, blocked by stored materials, or sealed by well-meaning contractors dealing with a water-ingress complaint. Doors left ajar change the airflow pattern the design depended on. None of these events is dramatic and none of them stops the site working on the day it happens, which is precisely the problem. On an unmanned site there is nobody there to notice that the background hum of the fan has stopped.

02

What happens when the cooling fails: a worked example

It is worth putting numbers on what follows a ventilation failure, because the timescale surprises people. The model below is illustrative and its assumptions are stated so the reader can substitute their own; the physics it rests on is first-year thermodynamics and is not in dispute.

Take a representative telecoms equipment cabin: three metres by two and a half by two and a half, an internal volume of roughly 19 cubic metres, with a continuous internal heat load of 2 kW from rectifiers, transmission equipment and radio. The envelope is insulated panel with a conductance in the region of 30 watts per kelvin across the whole shell. With the ventilation running, the interior sits a few degrees above ambient and everything is unremarkable.

Now let the ventilation fail on a summer morning with the ambient at 28 degrees.

The air in the cabin has very little thermal mass, about 23 kilojoules per kelvin for 19 cubic metres, so in the first minutes after the failure the air temperature climbs steeply, at several degrees per minute, until the equipment racks, the steelwork and the wall linings begin to absorb heat and slow the rise. Blending the air with a plausible allowance for that structural thermal mass gives the timeline below.

Time after ventilation failure	Modelled internal temperature	Condition
t = 0 (failure)	~30°C	Normal operation, a few degrees above the 28°C ambient
t + 15 minutes	36–38°C	Air-led initial spike; nothing visible or audible outside
t + 1 hour	42–46°C	Above the rated environment of most batteries and much equipment
t + 3 hours	52–58°C	Electrolytic capacitor territory; thermal shutdowns plausible
Equilibrium (never reached)	~95°C	65–70°C above ambient; equipment fails long before this

The interior passes 40 degrees before the morning is out and continues climbing through the afternoon as the ambient peaks.

The equilibrium arithmetic is the sobering part. A 2 kW load against a 30 W/K envelope settles roughly 65 to 70 degrees above ambient. On a 28-degree day the cabin is therefore heading for an internal temperature in the mid-nineties. It never gets there in practice, because equipment protection trips or components fail long before the interior reaches equilibrium, but that is exactly the point: the destination is so far above any survivable operating temperature that the only question is what breaks first and when. There is no self-limiting mechanism a site owner can rely on, only the race between the temperature curve and the thermal ratings of whatever is inside.

Two features of this model deserve emphasis. First, the timescale is hours, not days. A fan failure at eight in the morning is an equipment-threatening condition by mid-afternoon on a warm day. A weekly or monthly site visit regime has no realistic prospect of catching it; even a daily drive-by would miss most of the damage window. Second, the early rise is unmistakable in data and invisible in person. A temperature trace that has run flat for months and then climbs fifteen degrees in ninety minutes is about as clear a signal as instrumentation ever produces, but from outside the cabin, at the gate of an unmanned compound, nothing whatsoever has changed.

03

What sustained heat does to electronics

The relationship between temperature and electronic reliability has been studied for as long as electronics have been fielded, and the traditional engineering shorthand comes from the MIL-HDBK-217 reliability-prediction tradition: component failure rates roughly double for every ten-degree rise in operating temperature. That rule deserves honest handling. It is a crude aggregate, derived by applying Arrhenius-type acceleration to whole component populations, and the reliability literature has criticised it extensively, because real failure mechanisms vary in their temperature sensitivity and some dominate at different temperature ranges. But the criticism is about precision, not direction. No serious treatment of electronic reliability finds that running hotter is neutral, and across component classes the direction of the relationship is never overturned: sustained elevated temperature shortens electronic service life, and the effect compounds.

Some mechanisms are particularly well documented. Aluminium electrolytic capacitors, present in nearly every power supply and rectifier, lose electrolyte at a rate that follows the Arrhenius relation closely, and manufacturers publish the consequence plainly in their application data: expected lifetime roughly halves for every ten degrees of additional operating temperature.

Capacitor operating temperature	Expected lifetime (manufacturers' halving rule, 10-year part rated at 45°C)
45°C	10 years
55°C	5 years
65°C	2.5 years
75°C	~15 months

A power supply designed around capacitors rated for ten years at 45 degrees becomes a five-year power supply at 55 and a thirty-month one at 65. Solder joints accumulate fatigue damage through thermal cycling, so an enclosure that swings hot every afternoon and cools every night is working its interconnects harder than a stable one even if the peaks stay inside absolute ratings. Semiconductor junctions leak more and age faster at elevated temperature, and equipment that thermally throttles to protect itself is degrading its service to do so.

The practical consequence for a site owner is not usually a dramatic same-day failure, although the worked example above shows that a bad enough thermal event can produce one. It is the quiet forward-loading of the failure curve. Equipment on a site that spent a summer running twenty degrees hotter than intended will fail earlier, and it will do so at some later date, apparently at random, with no visible connection to the thermal history that caused it. Without a temperature record, the site owner books the failure as bad luck and replaces like for like into the same hot enclosure.

04

What heat does to standby batteries

If electronics degrade quietly in heat, valve-regulated lead-acid batteries do so on a published schedule. VRLA strings are the standby power reserve on most unmanned infrastructure sites, and their chemistry is textbook Arrhenius: the internal corrosion and dry-out mechanisms that end a battery's life are chemical reactions whose rates rise exponentially with temperature. Battery manufacturers state the consequence directly in their application data, and the figures are strikingly consistent across the industry: design life is quoted at a rated temperature of 20 to 25 degrees, and halves for roughly every 8 to 10 degrees of sustained operation above it. This is not a worst-case lawyer's caveat; it is the manufacturers' own published expectation for their products, and battery monitoring standards bodies repeat it.

The float life quoted on the datasheet is therefore not a property of the battery. It is a property of the battery at 20 degrees, and an enclosure that spends its summers in the high twenties and low thirties is buying a materially shorter life than the datasheet number with the same purchase order.

An illustrative model, consistent with the cost basis used in our earlier paper on the economics of alert-driven maintenance, shows what that means in money. Take a standby string with an installed cost of £2,750, planned for seven years of service at its 25°C rated temperature, and apply the manufacturers' halving rule to a range of sustained average enclosure temperatures.

Average enclosure temperature	Expected string life	Annualised string cost	Heat premium per string per year
25°C (rated)	7.0 years	£393	—
30°C	~5.0 years	£550	+£157
33°C	~4.0 years	£688	+£295
35°C	~3.5 years	£786	+£393
40°C	~2.5 years	£1,100	+£707

An enclosure whose thermal problems hold the string around eight degrees above rating takes the planned seven years down to roughly four, adding close to £300 a year to the standby power cost of that one string, before counting the engineer visit to replace it. On the field-service cost basis published by scheduling providers such as Fieldmotion, a loaded engineer hour in the UK runs £55 to £60, and industry dispatch analyses from TSIA place a fully burdened service visit in the hundreds of pounds once travel, vehicle and overhead are counted, which sits consistently with UK field service engineer salary medians around £35,000 reported by ITJobsWatch. A battery replacement visit to a remote site consumes most of a day. Multiply across an estate of dozens or hundreds of sites and premature battery replacement driven by heat becomes one of the larger controllable line items in the maintenance budget.

The money is arguably the lesser risk. A battery string does one job: it holds the site up when the mains fails. Heat-aged strings lose capacity before they fail outright, and a string at 60 per cent of rated capacity will pass a float-voltage glance and then let the site down forty minutes into the outage it existed to survive. The battery failure discovered during a mains failure is the most expensive kind there is, and elevated temperature is the single most common reason batteries reach that state ahead of schedule.

05

What a live estate actually shows

The sections above are physics and manufacturers' data. It is reasonable to ask how often any of this actually happens on real infrastructure, and here we can speak from operational evidence. IoT Technologies monitors a critical national infrastructure estate of around a thousand unmanned sites, and over the last eighteen months that estate has returned approaching two billion individual data points. Each site report carries eight measurement channels: temperature, humidity, pressure, sound level, smoke state, power source, battery voltage and signal quality. Smoke state is a reported channel like any other in that list, carried on a best-endeavours basis for visibility; it is not, and is not offered as, a fire detection or life-safety function, which remains the province of dedicated certified systems.

Three qualitative observations from that estate bear directly on this paper. First, temperature leads the alert stream. Of all the conditions the estate is instrumented for, thermal excursions are the most frequently raised, ahead of power events and every other channel. Whatever intuition suggests about how rare cooling failures ought to be, the data from a large fleet of real enclosures says that heat is the environmental condition an operator will hear about most often.

Second, the pattern is seasonal in exactly the way the physics predicts. The summer months dominate the thermal alert record. Enclosures that hold their margins comfortably from October to April run out of headroom in July, which is precisely the signature of ventilation schemes sized close to their limit and losing ground to a warming climate. It also means thermal risk is forecastable: an estate's exposure is not uniform through the year, and the case for verifying cooling plant before summer, rather than after the first hot week, is written plainly in the seasonal shape of the alert data.

Third, door state matters thermally. A door left open, or opened and not fully secured after a visit, changes the thermal behaviour of the enclosure, and extended door-open duration shows up in the estate record as a thermal event in its own right, not merely a security one. The enclosure's design assumes a closed envelope; monitoring shows how often that assumption quietly stops being true after routine site work.

None of these observations required anyone to visit anything. They are the by-product of continuous telemetry from sites that would otherwise report their condition only through failure.

06

Reducing the risk: measurement before mitigation

The engineering responses to enclosure heat are well understood: restore or upgrade ventilation, clean and schedule filter changes, resize cooling for the current equipment load rather than the original one, shade or reorient sun-loaded cabinets, and specify batteries and electronics with the real thermal environment in mind. All of these work. The difficulty on an unmanned estate has never been the fixes; it is knowing which sites need them, and knowing when a site that was fine last month has stopped being fine.

That is a measurement problem, and it is the one continuous monitoring solves. A temperature channel reporting from inside every enclosure turns the invisible failure chain of this paper into a sequence of observable, actionable events. The fan failure announces itself as a rate-of-rise alert within the first hour, on the day it happens, rather than surfacing as an equipment failure months later. The enclosure running chronically warm identifies itself in the trend data, letting an operator rank an estate's sites by thermal stress and direct ventilation spend at the sites that are actually consuming their batteries and electronics early. The pre-summer question, which sites will struggle this July, stops being a matter of guesswork or blanket surveys and becomes a question the last twelve months of data already answers. And because the telemetry travels one way from site to dashboard, the monitoring layer adds visibility without adding an attack surface or interfering with the equipment it watches.

There is a wider efficiency in this worth naming. Fixing thermal problems from data rather than discovering them from failure removes exactly the site visits that achieve nothing: the inspection that finds everything fine, and the emergency call-out that arrives after the damage is done. Every removed visit saves the engineer time priced above and, for operators reporting scope 1 emissions, the fuel: on the government's published 2026 greenhouse-gas conversion factor of 2.57 kgCO_2e per litre of diesel, a single 200-mile round trip in a service vehicle at typical van consumption is roughly 75 to 80 kilograms of CO_2e that a monitored estate did not need to emit to learn what a sensor already knew.

Heat will remain the quiet tax on every unmanned enclosure that nobody measures. The physics is not negotiable, the manufacturers' derating curves are published, and a large live estate confirms that thermal excursions are the most common environmental event in the field. What is negotiable is whether an operator finds out about them in hours, from a dashboard, or in years, from the failure record. If parts of your estate are running hotter than anyone knows, a site survey and a temperature trace are the fastest way to find out.

Talk to us about instrumenting your enclosures before the next hot season does the survey for you.

Methodology note. The cabin thermal model in this paper, including the failure timeline table, is illustrative: it assumes a 19 m³ enclosure, a continuous 2 kW internal load and an envelope conductance of approximately 30 W/K, with a blended allowance for internal thermal mass; readers can substitute their own dimensions, loads and envelope values, and the qualitative conclusions are insensitive to the exact figures. The electronics failure-rate doubling rule per 10°C derives from the MIL-HDBK-217 reliability-prediction tradition and is a crude aggregate, cited here for direction rather than precision. The capacitor and battery tables apply manufacturers' published halving rules (electrolytic capacitor lifetime halving per 10°C; VRLA design life halving per 8–10°C of sustained operation above a 20–25°C rating, with 10°C used in the table as the conservative case) to a stated baseline and are illustrative models, not measurements. Cost figures appear only within labelled illustrative models: the £2,750 string cost and visit-cost basis are consistent with our paper on the economics of alert-driven maintenance; loaded engineer rates of £55–60 per hour are as published by Fieldmotion; dispatch cost ranges are as published by TSIA; the £35,000 UK field service engineer salary median is as reported by ITJobsWatch (January 2026); the diesel conversion factor of 2.57 $kgCO_2e$ per litre is from the DEFRA/DESNZ 2026 greenhouse gas reporting factors (published June 2026). Estate observations are qualitative and drawn from an anonymised critical national infrastructure estate; no client is identified and no per-site data is disclosed. Nothing in this paper is a service commitment, a life-safety claim or professional advice.