



# Advanced Diagnostic Scenarios

MODULE M40

MASTER TECH

PREREQ D29

## Short Version

Everything in the D track assumed one fault at a time. One fault bends the seven readings into one recognizable shape, and you learned to name the shape. This module is about the calls where that skill betrays you: two faults living on the same system, each one bending the numbers the other way, so the combined picture matches nothing on your card. A dirty condenser props up the head pressure that an undercharge is dragging down, and the gauge reads textbook. A tripped internal overload reads like two open windings on a compressor whose real problem is a five dollar fan capacitor. Pattern matching fails here by design, because the pattern you memorized only exists when one fault has the system to itself.

The way through is not a bigger pattern library. It is discrimination: find the reading that contradicts the story, change exactly one variable, restabilize, and read again. This module works four full compound fault cases, one for each of the four failure patterns from D22, then covers the faults that refuse to be present when you are (intermittents), the structured teardown of a call that came back (the callback autopsy), and the master level judgment call of when to keep grinding versus when to time-box, document, and hand off. The NIST fault sensitivity data runs underneath all of it: which faults shout in the readings, and which ones hide inside numbers that look normal.

## Key Values

ITEM	VALUE	WHY IT MATTERS
Healthy 3-ton baseline (100 F ambient, 78 F return, R-410A TXV)	Suction 130 psig (45 F), head 390 psig (115 F), SH 10, SC 10, split 20 F, amps near 75 percent RLA	Every compound case in this module is this picture, bent twice
Healthy condensing over ambient	15 to 30 F	A head reading inside this band can still be two faults canceling out
R-410A anchors used here	102 psig = 32 F, 108 = 35 F, 130 = 45 F, 317 = 100 F, 340 = 105 F, 365 = 110 F, 390 = 115 F, 445 = 125 F, 475 = 130 F	Convert every pressure before you reason about it
NIST: subcooling at 30 percent undercharge	Down 87.7 percent	The loudest single fault alarm in the lab data
NIST: undercharge before 5 percent COP loss	About 25 percent	Slight undercharge hides inside normal looking efficiency

ITEM	VALUE	WHY IT MATTERS
NIST: liquid line restriction penalty threshold	No real penalty until about 48 percent	The quietest fault in the dataset
NIST: low indoor airflow	About 10 percent COP loss at 30 percent restriction	Second worst fault per percent severity
NIST: field charge statistics	Over 60 percent of 55,000 units wrong charge; 95 percent failed at least one diagnostic	The system you inherit on a callback probably has a preexisting fault
Capacitor replacement threshold	Beyond minus 6 percent of rated microfarads	A cap can be legitimately bad AND not be the root cause
Internal overload signature	OL from C to R and C to S with S to R intact	Tripped overload, not open windings; the sum check proves it
Compressor winding sum check	C to R plus C to S equals S to R	Both windings cannot open while their series path reads perfect
Megohm floor for a strong motor	100 megohms or better at 500 V DC, under pressure, never vacuum	Clears a cooled compressor for service
TOD trip / IPR open (Copeland scroll)	290 F discharge / 550 to 625 psid	The protections that fake compressor death
Static budgets on a 0.5 in WC system	Return about 0.10, filter about 0.10, wet coil 0.20 to 0.30 published, supply about 0.10; trouble past 0.8 TESP	The four-port map is the discriminator for half the cases here
Airflow target	400 CFM per ton nominal, 350 floor in dry climates	CFM per ton is the verdict number after an airflow repair
Post-defrost stabilization	About 60 minutes	Readings taken inside the window are intermittent fault bait
Diagnosis time-box	45 minutes without a falsifiable hypothesis: restart the funnel. 90 minutes: phone a friend	Grinding past the box wastes the day and invites a guess

## Field Checklist

- Converted every pressure to saturation temperature before reasoning about it
- All seven readings on paper before naming any fault
- Asked of the full set: does every number tell the same story, or do two numbers contradict

- On contradiction: identified the one variable to change, changed only it, restabilized 10 to 15 minutes, re-read everything
- Cleared the cheap masks first: condenser coil condition, filter, blower wheel, before judging charge
- Static map run any time airflow is in question, all four ports, drops summed against TESP
- Capacitor measured AND the question asked out loud: what killed it
- No compressor condemned without the D26 sequence complete and written down
- On an intermittent: customer interview bounded WHEN before any tool came out
- Logger, recording meter, or alert history capture deployed when the fault would not appear live
- On a callback: first visit reconstructed from invoice and photos before touching the unit
- Measured versus assumed list written for the first visit; broken assumption named
- Time-box honored: hypothesis review at 45 minutes, phone a friend at 90
- Handoff packet written before escalation: readings, statics, model data, what is ruled out and by which number

## Full Breakdown

---

### Two Faults, One Set of Gauges

The D24 fault signature card works because of an unstated assumption: one fault at a time. Low charge drops both pressures, starves superheat high, empties subcooling. Low airflow drops suction, floods superheat low, leaves subcooling alone. Each fault is a vector, pushing specific readings in specific directions, and you learned to read the direction.

Two faults superpose. The vectors add. And the cruel arithmetic is that real world fault pairs are not random: the second fault is usually caused by, or causes, conditions that push readings opposite to the first. A dirty condenser raises head; the undercharge it helped create by overworking a marginal braze joint lowers head. Net: head looks fine. A starving blower drops suction toward icing; a tech who responds by adding charge raises it back. Net: suction looks fine on the recheck and the coil ices anyway.

So the master level question is never only "which fault matches this picture." It is "does every reading in this set agree." Seven readings that all tell one story are a single fault, and the D track already taught you those. One reading that contradicts the story is not noise to round off. It is the second fault waving at you. The cases below are built entirely around that one habit: find the contradiction, then find the single measurement that splits the two possible explanations.

A compound fault is not exotic. NIST field data found over 60 percent of 55,000 surveyed units running a wrong charge and 95 percent failing at least one installation diagnostic. The unit you walk up to on any given call likely has a preexisting, tolerated fault already on board. Whatever broke today is fault number two. You have been diagnosing compound systems your whole career; the D track just kept the second fault small enough to ignore. This module is for the days it is not.

## What NIST Says Moves First, and What Hides

NIST instrumented an R-410A TXV heat pump and dialed in faults one at a time at known severities. Two findings matter for compound work.

First, faults are not equally loud. Subcooling collapsed 87.7 percent at 30 percent undercharge, the most sensitive single indicator in the dataset, and it moves long before capacity does: the unit gives up only about 14 percent capacity and 9 percent COP at that same 30 percent undercharge, and you can pull roughly 25 percent of the charge before efficiency drops even 5 percent. Translation: subcooling screams while the customer feels almost nothing. The thermostat is deaf to a fault your gauges catch instantly.

Second, some faults are nearly silent in the readings that techs trust most. Overcharge barely touches capacity until past 18 percent; its loudest witness is discharge temperature, a reading most techs never take. A liquid line restriction produces no meaningful penalty until roughly 48 percent blocked, which means a drier can spend two summers slowly plugging before the first complaint, then present as a sudden TXV failure. Meanwhile compressor and reversing valve internal leakage, the fault with almost no gauge signature at low severity, was the worst performer per percent of severity in the whole study, costing about 70 watts per percent at 47 F. That is exactly why D29 gave you the 2 degree test: the worst fault economically is the one the manifold alone cannot see.

Hold both findings together and the compound fault problem comes into focus. The loud fault (undercharge via subcooling) can be masked by a second fault pushing the same reading back toward normal. The quiet faults (restriction, internal leakage, creeping overcharge) hide inside a readings set that pattern matches as healthy. Either way, the defense is the same: the full set, every time, and a refusal to average away a contradiction.

### Case Study 1: The Head Pressure That Lied

Failure pattern recalled: D22 pattern one, misreading charge.

A 3-ton R-410A TXV split, 100 F afternoon, 78 F return. Complaint: cooling but slowly losing ground by late afternoon. Your seven readings:

READING	VALUE	CONVERTS TO
Suction	108 psig	35 F coil
Suction line temp	57 F	Superheat 22
Head	390 psig	115 F condensing, 15 over ambient
Liquid line temp	111 F	Subcooling 4
Split	78 return, 64 supply	14 F
Compressor amps	About 78 percent RLA	Unremarkable

The pattern matching tech reads it in two seconds: head 390 at 100 F ambient is 15 over, dead center of the band, "charge is fine by head." Or, if he favors the low side: superheat 22 and subcooling 4, "low on charge," and he starts adding refrigerant on the spot.

Both miss the contradiction. Superheat 22 with subcooling 4 is the D24 low charge picture, but true low charge drags head DOWN. The D24 worked example put a genuinely undercharged 3-ton at 340 psig, 105 F, barely 5 over ambient. This system says "empty" on the low side and "perfectly fed" on the high side at the same time. Those two statements cannot both be true of charge alone. Something is propping head up exactly as hard as the missing charge is pulling it down.

The discriminating move costs nothing: walk to the condenser and look. The coil is matted with cottonwood and dirt on the intake side, invisible from three feet away. Clean it, restabilize 15 minutes, re-read: suction 102 psig (32 F), suction line 56 F, superheat 24, head 340 psig (105 F, now only 5 over), liquid line 102 F, subcooling 3, split 12. There it is: the exact D24 low charge signature, finally allowed to show itself once the second fault stopped editing the head pressure.

Now the wrong repair becomes obvious in hindsight. The tech who charged to subcooling 10 without cleaning the coil left the customer with a full charge behind a blanket: head at 445 psig (125 F) on the next 115 F day, amps climbing toward RLA, and a high pressure trip waiting for the hottest week of the year. He fixed the loud fault and weaponized the quiet one.

The actual repair: clean the condenser, then treat the low charge as what it is, a leak call. Four out of five leaks live in the indoor coil, so the A-coil gets inspected first, the leak found and repaired, the system evacuated and the charge weighed in. Final verification readings: 130 psig suction (45 F), 390 head (115 F), superheat 10, subcooling 10, split 20, amps near 75 percent RLA. Both faults found, both fixed, and the verification set proves it because every number agrees again.

## Case Study 2: The Capacitor That Kept Dying

Failure pattern recalled: D22 pattern two, replacing capacitors without finding what killed them.

A 4-ton system, PSC blower, three weeks after another company replaced the blower run capacitor for "weak airflow and ice." The customer called back because the coil iced again. As found, coil thawed, system running:

READING	VALUE	CONVERTS TO
Suction	108 psig	35 F coil
Suction line temp	39 F	Superheat 4
Head	340 psig	105 F condensing
Liquid line temp	95 F	Subcooling 10
Split	78 return, 52 supply	26 F, with lazy register flow
Blower capacitor	Rated 10 microfarads, reads 9.0	Minus 10 percent, past the minus 6 threshold

Superheat 4 with subcooling 10 is the airflow signature, full stop. The refrigerant side is healthy; the heat is missing. And the capacitor really is bad, again, three weeks young. Here is where the pattern matching tech books a parts run: "bad batch of caps," or "motor is wearing out, replace cap and motor." A part that measured out of spec went in his hand and the diagnosis ended there. That is pattern two exactly: the capacitor is the corpse, not the killer.

The discriminating measurement is the D25 four-port static map, which the first tech never ran:

PORT	READING (IN WC)	BUDGET
A, return path before filter	0.38	About 0.10
B minus A, filter	0.12	About 0.10
C minus D, wet coil	0.26	0.20 to 0.30 published
D, supply path	0.16	About 0.10
TESP (sum)	0.92	System rated 0.50

The map convicts the return path: 0.38 against a 0.10 budget. In the attic, the return flex is crushed nearly flat under a stored ladder. Now the kill chain reads itself backward: crushed return starves the blower, low airflow means less air washing over the motor, the motor runs hot every cycle, and the run capacitor strapped to that heat ages and drifts out of tolerance in weeks instead of years. Replace the cap without the map and you are subscribing the customer to a capacitor of the month club, with a side of iced coils.

The actual repair: replace the capacitor, yes, and re-hang and properly support the return flex. Post-repair map: return 0.12, filter 0.12, coil 0.26, supply 0.16, TESP 0.66. Measured airflow rises from roughly 1180 CFM (295 per ton, starving) to about 1500 CFM (375 per ton, above the 350 dry climate floor). Refrigerant recheck after 15 minutes: suction 130 psig (45 F), superheat 10, split 20. The motor now runs in moving air, and the new capacitor gets to die of old age instead of heat.

### Case Study 3: The Charge That Was Never Low

Failure pattern recalled: D22 pattern three, ignoring airflow while chasing refrigerant.

This one arrives as a callback you inherit. First visit, two weeks ago, a different tech: complaint was weak afternoon cooling, his invoice says "found low on charge, added refrigerant, cooling restored." Today the same house has an iced solid indoor coil and no cooling at all. The system: 3-ton, constant airflow ECM air handler programmed for 1200 CFM, ductwork original to the house.

Thaw the coil, run it, take the set. With the added refrigerant on board: suction 108 psig (35 F), suction line 38 F, superheat 3, head 340 psig (105 F), liquid line 91 F, subcooling 14, split 25 with almost nothing moving at the registers. Superheat 3 is brushing floodback. Subcooling 14 says the condenser is now holding more refrigerant than it wants. This system is overcharged and still starving for air.

Reconstruct visit one. Before the refrigerant went in, this system almost certainly read: suction 108 (35 F), superheat around 6, subcooling 10, split 25, weak flow. Superheat 6 was the entire diagnosis. Low superheat means the coil is flooded with refrigerant that cannot find warm air to boil, which is an air problem and never a charge problem. The first tech saw low suction, thought "low charge," and never let the superheat finish the sentence. Adding charge pushed superheat from 6 to 3, fed the ice, and converted a weak cooling complaint into a no cooling callback.

But the compound part is the history underneath. The static map explains why this duct system failed THIS month after limping for years:

PORT	READING (IN WC)
A, return path	0.22
B minus A, filter	0.38
C minus D, wet coil	0.28
D, supply path	0.17
TESP	1.05

The filter is a brand new 1 inch high MERV pleat the homeowner upgraded to last month, loading at 0.38 all by itself on top of a return path already double its budget. For years the constant airflow ECM masked the duct disease the D25 way: holding its programmed CFM by quietly eating watts, the silent compensation that keeps customers comfortable and ducts undiagnosed. The new filter pushed total static past the motor's compensation ceiling near 1.0 in WC, the ECM finally ran out of authority, CFM fell off the cliff, and the "sudden" low cooling appeared. The discriminating measurements were available on visit one: superheat already low, watt draw on the air handler nearly double what a healthy duct system asks of the same motor, and a static map screaming on two ports.

The actual repair: recover refrigerant and weigh the charge back to nameplate plus line set spec, replace the 1 inch pleat with a properly sized media cabinet so filtration stops costing a third of an inch of static, and open up the return path. Post-repair TESP lands at 0.70, the ECM comes back inside its comfortable range, and the verification set returns to baseline: suction 130 (45 F), head 390 (115 F), superheat 10, subcooling 10, split 20. Refrigerant removed from the repair: every ounce the first visit added. The charge was never low.

#### Case Study 4: The Compressor That Was Only Hot

Failure pattern recalled: D22 pattern four, condemning healthy compressors.

A 3-ton split on a 100 F afternoon. Complaint: quit cooling around 2 p.m. As found: thermostat calling, contactor pulled in, condenser fan spinning slowly, compressor silent, shell too hot to keep a hand on. Power off, leads off the compressor, ohmmeter on the terminals: C to R reads OL. C to S reads OL. S to R reads 3.1 ohms, smooth and steady.

The pattern matching tech writes "two open windings" and starts the replacement quote. He is about to condemn a healthy compressor, and the meter already told him so. Run the D26 sum check in your head: C to R plus C to S must equal S to R, because S to R is the two windings in series. If both windings were truly open, S to R could not read 3.1 ohms; the series path would be broken too. One device opening the common leg explains all three readings at once: the internal overload, sitting between the C terminal and both windings, tripped open on heat. OL to both from common, with the windings intact behind it, is the textbook D26 internal overload signature, and a shell too hot to touch is the confession.

So the compressor is not dead, it is protecting itself. The master question is from what. The slow condenser fan is the thread: pull and meter the dual capacitor. The fan section, rated 5 microfarads, reads 3.9, minus 22 percent, triple the minus 6 percent replacement threshold. The kill chain: weak fan capacitor, fan loses speed, condenser stops rejecting heat, condensing temperature climbs toward 130 F (475 psig on R-410A), discharge

gas temperature chases it upward until the thermal overload does its one job near 290 F discharge, and the compressor drops out mid afternoon, right when the customer noticed. Fault one is the tripped overload faking compressor death. Fault two is the capacitor that caused the trip. Replace only what is in your hand at the moment of discovery and you fix neither.

#### **PHOENIX FIELD NOTE**

A tripped internal overload on a Phoenix afternoon resets on its own schedule, not yours. The dome soaks at shell temperatures well past anything a meter lead should touch, and the overload will not re-close until the mass of the compressor cools, which can take hours in 110 F shade. Do not stand there re-ohming it every five minutes and do not let the wait talk you into condemning it. Stage the rest of the diagnosis, replace the root cause part, rinse the condenser coil, and give the shell its time. The overload resetting on a cooled shell IS a data point: it proves the protection works.

The actual repair: replace the dual capacitor, clean the condenser coil while the shell cools, then finish the D26 sequence before declaring victory. Windings cold: C to R 0.5, C to S 2.6, S to R 3.1, sum check passes. Megohm at 500 V DC with the system under pressure, never in vacuum: 120 megohms, strong. Restore power, run 15 minutes: suction 130 psig (45 F), head 390 (115 F, 15 over ambient), amps near 75 percent RLA, discharge line warm not screaming. The compressor that almost got replaced finishes the day moving heat, and the invoice carries a capacitor and a coil cleaning instead of a changeout the system never needed.

### **The Discrimination Habit**

Pull the four cases apart and the same skeleton is inside every one.

First, the full set, always. Every case turned on a reading the lazy version of the call would have skipped: the head conversion in case one, the static map in cases two and three, the sum check in case four. You cannot find a contradiction in numbers you did not take.

Second, name the contradiction out loud. "Subcooling says empty but head says full." "The cap is bad but caps do not die of natural causes in three weeks." "Superheat says flooded but the invoice says charge was added." "Two windings open but their series path reads perfect." A contradiction stated in one sentence points directly at its own discriminating test. If you cannot state it, you have not found it, and you are not ready to repair anything.

Third, change one variable, then re-read everything. Clean the coil, THEN re-judge charge. Replace the cap, THEN map the static. Cool the shell, THEN re-ohm. The D22 funnel called this the confirmation test; compound work just runs the funnel twice on the same system. The single most common compound fault failure is fixing fault one and walking away without the second full readings set that would have exposed fault two. Stabilization rules still apply: 10 to 15 minutes after any change before the numbers mean anything.

Fourth, the verification set is the receipt. A compound repair is finished when every reading agrees with every other reading and with the baseline for the conditions, not when the part is installed. Two faults means two fixes means the final set is the only proof both landed.

## Intermittent Faults: Diagnosing a Ghost

The compound fault hides in plain sight. The intermittent fault simply leaves before you arrive. The unit that "quits every afternoon" runs flawlessly for the entire two hours you stand in front of it. You cannot discriminate readings you cannot capture, so the entire game becomes capturing them.

Know the five usual ghosts. Thermal intermittents: a cracked solder joint on a control board that opens only at temperature, a flame sensor that reads clean microamps cold and sags below dropout after twenty minutes of heat soak. These faults are invisible at the moment of inspection by definition; the part works at the temperature you test it at. Load dependent electrical sags: a utility drop or undersized conductor that holds fine voltage at idle and dips below the 187 V start floor on a 208 to 230 V nameplate exactly when the compressor asks for inrush, killing the start and leaving no evidence. Pressure switch flutter: a switch sitting at the ragged edge of its setpoint, chattering open for half a second under a wind gust or a marginal vent, logging a lockout the system later runs away from. Defrost cycle faults: everything from D29 that only misbehaves inside or right after a defrost event, invisible in any reading taken more than an hour later. And humidity dependent faults: condensate float switches that trip only when peak season humidity finally loads the drain, board corrosion that leaks current only at high dew point.

### PHOENIX FIELD NOTE

Monsoon season manufactures a specific class of ghost. A system that ran perfectly from October through June starts nuisance tripping in July, and the trigger is water: the first real condensate load of the year finds the half plugged drain and the barely tilted float switch, and storm-hour electrical noise stacks fault histories with timestamps clustered at 2 a.m. Ask the customer one question early: did this start with the storms. A yes moves the drain pan, float switch, and board environment to the front of the line, and makes the timestamp pattern in any fault history your first witness.

The tools exist because the ghost will not perform on demand. A data logger or recording meter with min and max capture parked on the suspect circuit turns "it quits sometimes" into a timestamped voltage sag you can read off the screen tomorrow. A recording meter left across the pressure switch catches the half second flutter no human with a probe ever will. On communicating systems, the run capture is already done for you: the A32 fault history with timestamps is a flight recorder, and the A32 rule applies doubly here, photograph the history before anything gets cleared, because a hundred identical events clustered in one storm hour and the same hundred spread over three weeks are two different diagnoses.

The cheapest logger on the truck is the customer. Run the structured interview before any tool comes out, and aim every question at bounding WHEN: time of day or night, weather running at the time, how long the system had been running, which mode, what they heard or smelled, what resets it. "It dies on hot afternoons after running a while" is a thermal or load profile. "Only during storms" is moisture or noise. "Only on cold mornings" points at defrost. A ghost bounded in time is half caught: now you know when to be there, or what window the logger has to cover, or which fault family the history timestamps should cluster in. An intermittent diagnosed without bounding WHEN is a parts cannon firing blind.

## The Callback Autopsy

A callback is a diagnostic gift wrapped in an uncomfortable feeling. Some tech, maybe you, stood in front of this system with full access and reached a conclusion the system has since voted against. That failed conclusion is data. The unit just told you which assumption in the first diagnosis was false, and the autopsy method exists to find it on purpose instead of by luck.

Step one: reconstruct the first visit before touching the unit. Pull the invoice and the photo record and rebuild what the first tech saw, did, and concluded. What readings are written down. What parts went in. What the stated diagnosis was.

Step two: build the two column list, measured versus assumed. Everything with a number in the record goes in the measured column. Everything the diagnosis needed to be true that has no number behind it goes in the assumed column. This is the heart of the method. "Charge low" with a subcooling reading is measured; "airflow fine" with no static, no split, no CFM anywhere on the ticket is assumed. In case three above, the entire assumed column was one word, airflow, and that was the autopsy in thirty seconds.

Step three: the broken assumption is almost always in the assumed column, and the callback symptom points at which entry. Came back iced: the airflow assumption broke. Same part dead again: the root cause assumption broke. High head trip in a heat wave: the condenser condition assumption broke. Re-run the funnel with the broken assumption promoted to prime suspect, and take on visit two every measurement visit one skipped.

Step four: close the loop. The repair that survives is the one where the previously assumed item now has a number next to it.

### IB STANDARD

Island Breeze counts a callback as the customer calling back within 30 days reporting the same issue, and every callback gets reviewed against the original 8-photo ServiceTitan close-out record. That photo set is the autopsy's evidence locker: the readings screen, the static and split documentation, the as-left condition of coil, filter, and panel. A close-out photographed to standard makes the measured versus assumed list ten minutes of work. A thin photo record makes the first visit unreconstructable, which is itself the finding, and the review is run without blame: the autopsy hunts broken assumptions, not techs. The lesson from every autopsy gets written into the job record so the next tech inherits the conclusion, not the mystery.

## Escalate or Grind

Master techs do not get stuck less than other techs. They notice sooner that they are stuck, and they have a protocol for it instead of an ego about it.

Time-box the diagnosis. At 45 minutes without a falsifiable hypothesis, stop and restart the funnel from the top: re-verify the complaint, re-walk the survey, re-take the full set. Not because the numbers changed, but because the second pass with fresh eyes catches the reading you mentally rounded off the first time. At 90 minutes still stuck, phone a friend. The clock is not an insult; it is the same discipline as the stabilization timer. Past the box, the human tendency is to start guessing with parts, and a guess installed at hour three costs more than a phone call at hour two.

The phone call only works if you can hand your colleague the system in sixty seconds, which means the packet is built BEFORE the call: all seven readings with conversions, static map if airflow is in play, model and serial, what is ruled out and by which specific number, and the contradiction stated in one sentence. "Three ton TXV, suction 108 converts to 35, superheat 4, subcooling 10, split 26, TESP 0.92 with 0.38 on the return port, airflow convicted, but registers are still weak after the return repair" gets useful help. "It is acting weird" gets sympathy. And when the problem genuinely hands off, to a senior tech, to tomorrow, to a factory tech line, what a master writes down is the difference between a relay and a restart. The handoff note carries four things: the complete readings record with timestamps and conditions, the hypotheses already falsified and the measurement that killed each one, the staged state of the equipment (panels, jumpers, valve positions, anything not as found), and the single next test you would run if you were staying. A handoff with those four things means the next person starts at hour two of the diagnosis instead of minute one. Anything less and the customer pays for the same first hour twice.

#### IB STANDARD

At Island Breeze the 90 minute phone a friend goes to the senior tech on call, and the escalation itself gets documented in ServiceTitan: who was called, what was tried, what the handoff packet contained. An escalation logged to standard is a mark of judgment, not a mark against the tech; the review that matters is whether the packet let the second tech start where the first one stopped.

## Common Mistakes

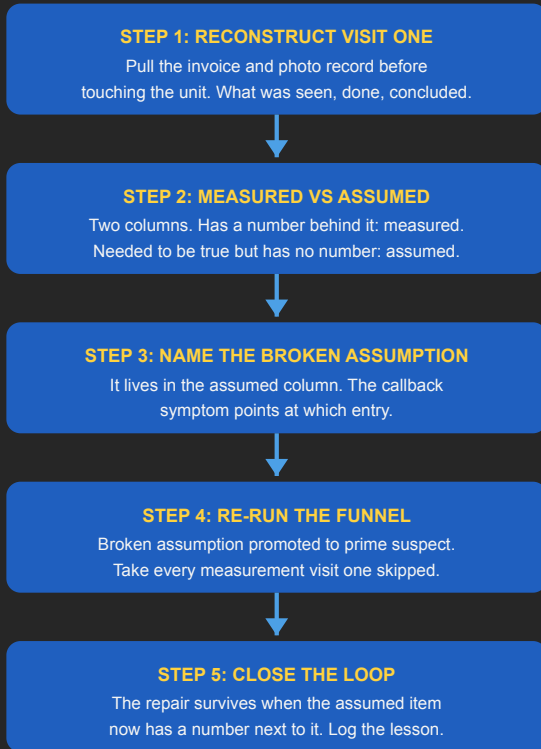
1. **Averaging away the contradiction.** Subcooling 4 with head dead normal gets mentally rounded to "close enough, low charge" and the second fault survives the visit. A contradiction between two readings is the diagnosis, not noise in it.
1. **Fixing fault one and skipping the second full readings set.** The condenser gets cleaned, the cap gets replaced, the tech feels the cold air and leaves. Ten of the fourteen readings that would have exposed fault two were never retaken. Every change gets 10 to 15 minutes of stabilization and a complete re-read, every time.
1. **Treating a bad part as a complete diagnosis.** The capacitor reading minus 10 percent is true and insufficient. Parts that die young were killed. No young dead part leaves the diagnosis until the killer is named with a number: static, voltage, heat, cycling.
1. **Adding refrigerant against low superheat.** Superheat 6 or below says the coil is flooded and cannot find heat, which is an air problem. Charge added there feeds floodback and ice and converts a weak cooling call into a no cooling callback. The funnel question order exists for exactly this: superheat first, and low superheat exits to the air side before charge is ever on the table.
1. **Condemning a hot compressor on OL readings.** OL from C to both terminals with S to R intact is a tripped internal overload, and the sum check proves the windings are whole. The shell gets time to cool, the root cause gets hunted in the meantime, and no compressor is condemned without the full D26 sequence on paper.

1. **Hunting a ghost without bounding WHEN.** Two hours of perfect readings on an intermittent system proves only that the fault is not present now. The customer interview, the logger, and the fault history timestamps bound the window first; tools deployed before the WHEN is bounded are pointed at nothing.
1. **Clearing fault histories on arrival.** On a communicating system the alert history with timestamps is the only witness to an intermittent. Cleared, it is gone forever. Photograph first, always, then diagnose, then clear.
1. **Treating the callback as shame instead of data.** The tech who avoids the autopsy, or runs it looking for someone to blame, leaves the broken assumption alive for visit three. The measured versus assumed list is run cold, the broken assumption gets named, and the lesson lands in the job record.
1. **Grinding past the time-box.** Hour three of a stuck diagnosis produces parts cannon behavior in good techs. The 45 and 90 minute checkpoints are discipline, not weakness, and the handoff packet is what makes escalation cheap.
1. **Verifying the repair with one reading.** A compound repair signed off on supply air temperature alone is half verified at best. Two faults, two fixes, one proof: the complete set in agreement with baseline, written down.

## CALLBACK AUTOPSY FLOW

### THE CALLBACK AUTOPSY

A returned call is data, not shame. Find the broken assumption on purpose.



**SYMPTOM POINTS AT THE ASSUMPTION**

**Came back ICED:**  
the airflow assumption broke

**Same part DEAD AGAIN:**  
the root cause assumption broke

**HIGH HEAD trip in a heat wave:**  
the condenser condition assumption broke

**Fault WILL NOT REPEAT on site:**  
the WHEN was never bounded; it is a ghost, deploy capture tools, not parts

**EXAMPLE: 30 SECOND AUTOPSY**

Invoice says: low on charge, refrigerant added.  
Measured column: suction pressure. That is all.  
Assumed column: airflow. No static, no split, no CFM anywhere on the ticket.

**Callback symptom: iced coil.**

Verdict: airflow was assumed, airflow broke. Done.

**THE AUTOPSY HUNTS BROKEN ASSUMPTIONS, NOT TECHS.**

## CASE STUDY GRID

# FOUR COMPOUND CASES, ONE SKELETON

Full set. Contradiction out loud. One variable changed. Full set again.

### THE PICTURE

#### 1. HEAD PRESSURE THAT LIED

SH 22, SC 4: says low charge  
Head 390, 15 over: says fine  
**Dirty condenser + undercharge**  
canceling on the head gauge

### THE WRONG CALL

**Charge to SC 10 behind a  
blanketed coil; head rides 445**  
High pressure trip waiting  
for the hottest week

### THE DISCRIMINATOR

Eyes on the coil. Clean it,  
restabilize, re-read: head  
falls to 340, SC 3, the true  
low charge picture appears

### BOTH REPAIRS

Clean coil, find the  
leak (A-coil first),  
repair, evacuate,  
weigh in, verify

#### 2. CAP THAT KEPT DYING

SH 4, SC 10: airflow fault  
Cap minus 10 pct, 3 weeks old  
**Weak cap + crushed return**  
cooking the blower motor

**Swap the cap again and leave.**  
The corpse got replaced,  
the killer kept the job.  
Capacitor of the month club

Four-port static map:  
return port 0.38 against a  
0.10 budget, TESP 0.92.  
The map convicts the duct

New cap AND re-hang  
the return. TESP 0.66,  
1500 CFM, motor  
cooled, verify

#### 3. CHARGE THAT WAS NEVER LOW

SH 6 said airflow; refrigerant  
added anyway; coil iced  
**Duct disease + new restrictive**  
filter past the ECM ceiling

**Low suction read as low charge.**  
Charge against SH 6 fed  
floodback and ice: weak cooling  
became NO cooling

Superheat verdict first: low  
SH exits to the air side.  
Filter port 0.38, TESP 1.05,  
ECM watts nearly doubled

Weigh charge back to  
spec, media cabinet,  
open the return,  
TESP 0.70, verify

#### 4. COMPRESSOR ONLY HOT

C to R OL, C to S OL,  
S to R 3.1 ohms, shell scorching  
**Tripped overload + weak**  
condenser fan capacitor

**Two open windings, quote a  
compressor.**  
A healthy compressor  
condemned by its own protection

The sum check: S to R intact  
means both windings whole.  
Fan cap reads 3.9 of 5,  
minus 22 pct: the root cause

New capacitor, clean  
coil, cool the shell,  
full D26 sequence,  
run and verify

**TWO FAULTS, TWO FIXES, ONE PROOF: A FINAL SET WITH ZERO CONTRADICTIONS.**

## COMPOUND FAULT WEB

# COMPOUND FAULTS: TWO VECTORS, ONE GAUGE

Each fault bends the readings the other way. The sum looks normal. The sum is lying.

### FAULT A: DIRTY CONDENSER

Coil cannot reject heat

**Pushes head pressure UP**

alone: head toward 445 psig (125 F)

### FAULT B: LOW CHARGE

Not enough refrigerant anywhere

**Pulls head pressure DOWN**

alone: head toward 340 psig (105 F)



**READS NORMAL. IS NOT.**

### THE FULL SET STILL TELLS ON THEM

**SUPERHEAT 22**

says: coil starving

**SUBCOOLING 4**

says: charge missing

**HEAD 390, NORMAL**

says: charge is fine

**SPLIT 14**

says: weak capacity

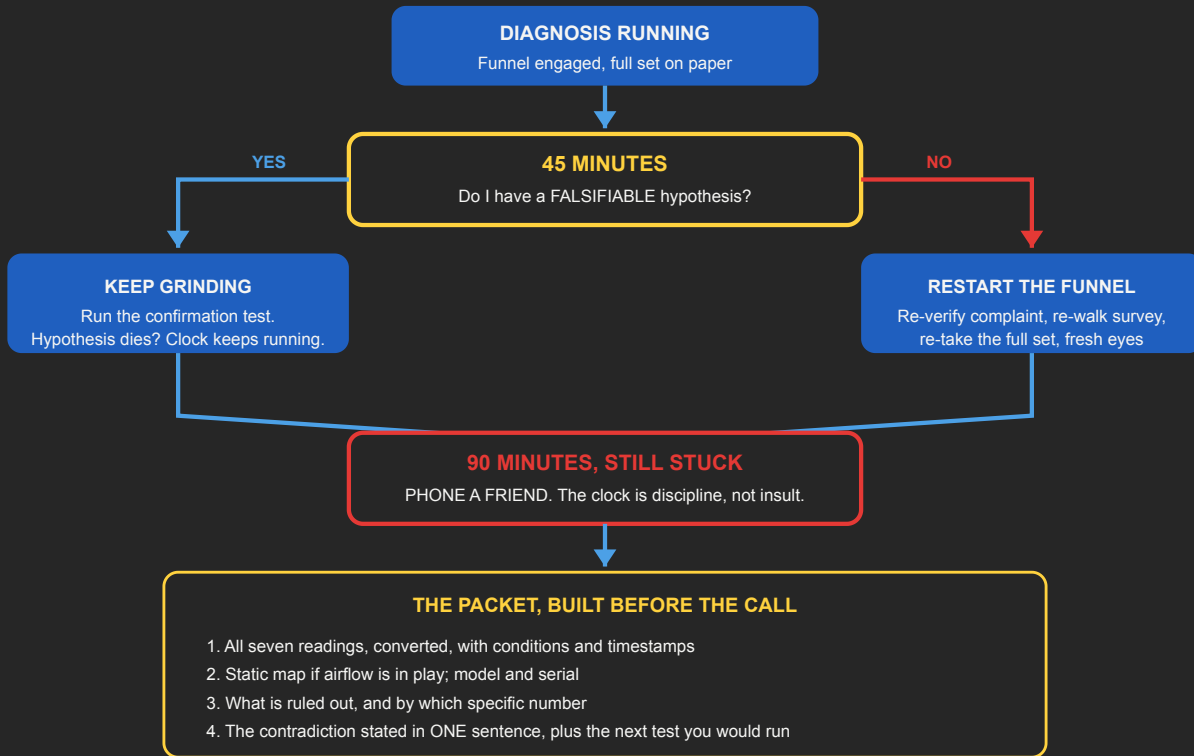
**ONE READING CONTRADICTING THE STORY IS THE SECOND FAULT WAVING AT YOU.**

Change one variable, restabilize 10 to 15 minutes, take the full set again.

## ESCALATION DECISION TREE

# ESCALATE OR GRIND: THE TIME-BOX

Masters do not get stuck less. They notice sooner, and they have a protocol.



**A GUESS AT HOUR THREE COSTS MORE THAN A PHONE CALL AT HOUR TWO.**

## INTERMITTENT FAULT TRAPS

# INTERMITTENT FAULT TRAPS

The fault leaves before you arrive. Bound WHEN, then capture it.

### THE GHOST

#### THERMAL

Cracked board joints, flame sensor that sags hot, fine when cold

#### LOAD DEPENDENT SAG

Voltage fine at idle, dips below 187 V start floor under inrush

#### PRESSURE SWITCH FLUTTER

Switch at the edge of setpoint, chatters open for half a second

#### DEFROST CYCLE ONLY

Misbehaves only inside or right after a defrost event (D29)

#### HUMIDITY ONLY

Float switch, plugged drain, board corrosion at high dew point

### WHEN IT SHOWS

After 20 plus minutes of runtime, hot afternoons, heat soak

Only at compressor start, worst at peak grid hours

Wind gusts, marginal venting, random lockouts with no witness

Cold mornings; invisible in any reading taken an hour later

First wet weeks of the season; storm-hour fault clusters

### HOW TO CATCH IT

Test AT temperature; heat the suspect, watch the reading live

Recording meter, min and max capture across the start event

Logger parked across the switch; no human probe catches it

Force a cycle and watch all four actions; 60 minute rule after

Fault history timestamps; photo the history BEFORE clearing

**THE CHEAPEST LOGGER IS THE CUSTOMER: BOUND WHEN BEFORE ANY TOOL COMES OUT.**