ISLAND BREEZE TECHNICIAN CERTIFICATION PROGRAM

Load Calculation and Equipment Selection

MODULE M37 **MASTER TECH** **PREREQ A36**

In F3 you learned that heat flows from hot to cold at a rate set by the temperature difference and the resistance in its path, and every module since has been about moving, measuring, or fixing that flow. This module is about predicting it. Before a single piece of equipment gets ordered, somebody has to answer one question with a number: how much heat enters this house on the hottest design hour, and how much leaves it on the coldest? That number is the load, the calculation that produces it is ACCA Manual J, and the discipline that matches a real machine to it is Manual S. Everything downstream, equipment, refrigerant charge behavior, airflow, duct design, comfort, equipment life, inherits the quality of that first number. A tech at master level does not guess tonnage from square footage any more than they guess charge from line frost. This module teaches you what drives the load, how to walk a house and capture the inputs honestly, how to read manufacturer expanded performance data at the design conditions this market actually serves, and the sanity checks that catch a garbage calculation before it becomes a fifteen year mistake bolted to a pad.

Short Version

A cooling load is a physics answer, not a folklore answer: the sum of heat entering through glazing (driven by orientation and SHGC), conduction through walls, ceilings, and floors (driven by U-values, the inverse of R-values), infiltration (outside air leaking through the envelope), internal gains (people, appliances, lights), and the penalty for ducts running outside the thermal envelope. ACCA Manual J is the ANSI-recognized standard for that calculation, run at design conditions, the 1 percent cooling and 99 percent heating values for the location, never record extremes. This course's running design point is 112 F outdoor, 75 F indoor cooling. The output splits into sensible load (temperature) and latent load (moisture), and both carry forward. Square footage rules of thumb fail because the envelope, not the floor area, sets the load: three 2,000 square foot homes of different vintages can load at roughly 52,000, 38,000, and 26,000 BTU per hour, and one 500-square-feet-per-ton rule hands all three the same 4 ton answer, undersized for the first, sloppy for the second, and badly oversized for the third. Oversizing is not a safety margin, it is a failure chain: short cycles, comfort swings, weak moisture removal, start-stress wear on capacitors, contactors, and compressors, and airflow demands the existing ducts may not carry. Manual S turns the load into an equipment selection using OEM expanded performance data, not nameplate tons: a nominal ton is 12,000 BTU per hour at the AHRI rating point of 95 F outdoor with 80 F dry bulb and 67 F wet bulb entering air, and no house at 112 F design lives at that point. Read the expanded table at the actual ambient and the actual entering wet bulb, interpolate between rows, and apply the sizing window: single-stage cooling equipment lands between 90 and 115 percent of the total load, with wider windows for staged and variable equipment, and latent capacity must cover the latent load. Block load picks the box; room-by-room loads size the ducts in M38. Before trusting any calc, run the sanity checks: design conditions verified, BTU per square foot inside the envelope's expanded band, latent fraction believable for the climate, component breakdown matching the house you actually walked, and the existing equipment's runtime behavior on a near-design day used as a measuring instrument. A load calc you did not sanity-check is a number you copied, not a number you know.

Key Values

VALUE	NUMBER	WHAT IT MEANS
Course design conditions (cooling)	112 F outdoor, 75 F indoor	The Manual J design point used throughout this course. Verify, never assume software defaults.
Design condition philosophy	1 percent cooling, 99 percent heating	Loads are calculated at conditions exceeded only about 1 percent of hours, never at record extremes.
Indoor heating design	70 F	Standard Manual J indoor heating condition.
Nominal ton	12,000 BTU per hour	A label, not a delivery. Rated at AHRI conditions only.
AHRI cooling rating point	95 F outdoor, 80 F db / 67 F wb entering air	The lab point behind every nameplate. Not the design point of a hot dry climate.
Capacity loss at high ambient	Roughly 1 percent of total capacity per degree above 95 F (representative; read the OEM table)	A nominal 3.5 ton delivering about 42,000 on the label delivers in the neighborhood of 36,000 to 37,000 at 112 to 115 F.
Manual S window, single-stage cooling	90 to 115 percent of total load	Multi-stage up to about 120 percent, variable capacity up to about 130 percent. Smallest selection that covers wins.
Latent rule	Latent capacity at design conditions must meet or exceed latent load	A unit that meets sensible but misses latent does not control humidity.
Airflow recall (C12)	400 CFM per ton nominal, 350 to 450 window	Dry climates run the top of the window to favor sensible capacity.
Duct-outside-envelope penalty	Commonly 10 to 25 percent added load	Leakage plus conduction. Duct location is a load input, not an afterthought.
Duct leakage recall (A36)	20 to 30 percent typical	The same number that wrecked delivered capacity in A36 inflates the load here.
Dry-climate latent fraction	Roughly 5 to 10 percent of total cooling load	A desert calc showing 25 percent latent has a wrong input somewhere.
Sanity bands (cooling, at hot-dry design)	Pre-1990 unimproved envelope about 25 to 35 BTU/sq ft; 1990s partial upgrades about 18 to 25; 2015+ tight construction about 12 to 18	Sanity check bands only. Never sizing rules.

VALUE	NUMBER	WHAT IT MEANS
Short cycling wear recall	Capacitors 21 percent of service calls	Start-stress components die first on oversized, short-cycling equipment.

Field Checklist

Load input survey on any install or changeout candidate:

- ☐ Verify design conditions in the software: correct weather station, 112 F outdoor and 75 F indoor for this course's market, 70 F indoor heating. Never accept defaults sight unseen.
- ☐ Measure conditioned floor area yourself (laser or plans). Listing square footage lies, and it includes garages.
- ☐ Compass the house. Note which walls and glass face south and west.
- ☐ Windows: count, measure, pane count, frame material, low-E coating, interior shading, exterior screens, overhang depth. Glazing is usually the biggest single lever in a hot climate.
- ☐ Walls: construction type (masonry block, frame, foam), insulation evidence (outlet box peek, knock test, vintage).
- ☐ Attic: insulation type and depth (convert depth to R), radiant barrier, sealed or vented, and where the ducts run.
- ☐ Ducts: inside or outside the envelope, insulation R on the duct, visible leakage signatures from the A36 walk.
- ☐ Infiltration: age, weatherstripping condition, can lights, fireplace dampers, blower door number if one exists.
- ☐ Internal gains: occupant count (bedrooms plus one as the default convention), kitchen and appliance load, anything unusual (server racks, aquariums, hot tubs indoors).
- ☐ Interview the homeowner on the existing equipment: does it hold setpoint on the hottest afternoons, and does it run continuously or cycle? Runtime behavior brackets the real load.
- ☐ Run the calc, then sanity-check: BTU per square foot band, latent fraction, component breakdown versus the house you walked.
- ☐ Manual S: pull expanded performance data at design ambient (interpolate to 112 F), verify total capacity lands in the 90 to 115 percent window for single-stage, sensible covers sensible, latent covers latent.
- ☐ Confirm the selected airflow against the duct system's measured static history (C12, A36). Equipment selection writes a check the ducts have to cash, and M38 is where that check clears.
- ☐ Save the full calc output to the job record before any equipment is ordered.

IB STANDARD

A Manual J runs on every install and every changeout, no exceptions, and the complete output is saved to the job record in ServiceTitan before equipment is ordered. A like-for-like swap without a load calc repeats the original sizing decision, right or wrong, and signs the company's name to it for the life of the equipment.

Full Breakdown

The load is a physics problem, and you already know the physics

Recall F3: heat moves by conduction, convection, and radiation, always hot to cold, at a rate proportional to the temperature difference and inversely proportional to the resistance in the path. A house in summer is a box at 75 F sitting in air at 112 F with the sun pouring radiation onto its glass and roof. Heat is coming in through every surface, every crack, and every kilowatt-hour the occupants burn inside. The cooling load is simply the rate of that inbound flow at design conditions, in BTU per hour. The heating load is the same problem run in reverse on a winter design morning.

Break the inbound flow into its components and you have the anatomy of every Manual J:

1. **Glazing gains.** Windows admit heat two ways: conduction through the assembly (the U-value), and solar radiation straight through the glass, which is governed by SHGC, the Solar Heat Gain Coefficient, a 0-to-1 number for the fraction of solar energy a window passes. Orientation multiplies this: west glass takes the full afternoon sun at the same hour the outdoor air peaks, which is why a west-facing sliding door can out-load an entire insulated wall.
2. **Envelope conduction.** Walls, ceilings, floors. Each assembly has an R-value (resistance) and its inverse U-value (transmission). Area times U times temperature difference is the flow. The attic matters double because the roof deck above it can run far hotter than the outdoor air, so the ceiling sees a bigger effective temperature difference than any wall.
3. **Infiltration.** Outside air leaking in through the envelope carries both sensible heat (it is hotter than the inside air) and latent heat (whatever moisture it holds). Infiltration is the input with the widest honest uncertainty, estimated from construction age and tightness, or measured with a blower door when one is available.
4. **Internal gains.** People (a sitting adult gives off a couple hundred BTU per hour sensible plus moisture), cooking, appliances, lights, electronics. These do not care about the weather, which is why they dominate the load fraction in very tight homes.
5. **Duct gains.** Ducts running outside the thermal envelope, through a hot attic being the classic case, both leak conditioned air (the A36 number: 20 to 30 percent typical) and absorb heat through their walls. Manual J treats duct location and leakage as load inputs, because they are. Moving ducts inside the envelope, or sealing them, changes the load itself.

Sum the components and split the answer two ways: **sensible load**, the heat that raises air temperature, and **latent load**, the energy needed to condense moisture out of the air. Both numbers travel forward to equipment selection, because a machine has separate sensible and latent capacities and must cover each.

One more foundation stone: **design conditions**. Manual J calculates loads at the 1 percent cooling and 99 percent heating design values for the location, the conditions exceeded only about 1 percent of the hours in a year, never at record extremes. This course's anchor is 112 F outdoor and 75 F indoor for cooling, with 70 F indoor for heating. Designing for the all-time record would mean oversizing for 99-plus percent of the equipment's life, and the rest of this module is about why that trade is a bad one. Outdoor and indoor design values come from Manual J's own tables or ASHRAE data, and the standard is explicit that you do not mix sources within one calculation.

PHOENIX FIELD NOTE

Phoenix's Manual J design point is 112 F outdoor, 75 F indoor, and the winter heating design sits around 40 F, which is why cooling governs nearly every sizing decision here. Yes, the thermometer will touch 115 to 118 F a few afternoons a year. That is not a reason to inflate the load; it is a reason to read the equipment's expanded performance table at those ambients, which is exactly what Manual S does later in this module. Heat waves are handled by honest capacity data, not by padded loads.

The design chain: J, then S, then D, in that order

Manual J produces the loads. Manual S selects equipment whose real capacity at the real design conditions matches those loads. Manual D sizes the duct system to deliver the selected airflow to each room. The chain runs in that order because each step's output is the next step's input, and changing any link reopens the links after it: pick different equipment and the airflow and available static change, which changes the duct design. Manual J and Manual D are ANSI-recognized standards referenced by building codes, so skipping them is not just sloppy, in many jurisdictions it is a code violation. This module covers J and S. Manual D, friction rates, and duct sizing belong to M38, and every room-by-room number you generate here is raw material for that work.

The standard is also blunt about something this trade keeps relearning: **no rules of thumb**. Square feet per ton, CFM as floor area times a multiplier, "go up a half ton to be safe," all of it is explicitly rejected. Every house gets a calculation, including, especially, changeouts, because a like-for-like swap without a calc just re-signs the original installer's guess.

Walking the envelope: why two houses with one floor area load nothing alike

Here is the comparison that kills square-footage thinking permanently. Take three single-story homes, 2,000 square feet each, in the same hot-dry market at the same 112 F design point. The numbers below are representative round figures for teaching; every real house gets its own calc.

House one: mid-1980s masonry block, original everything. Uninsulated or barely insulated block walls (real-world R in the low single digits), single-pane metal-frame windows with no low-E coating, an attic insulated to an old code level, vented and hot, with the ductwork running through it, and an envelope that leaks air at every penetration. Representative component picture: glazing 14,000 BTU per hour, walls 8,500, ceiling 9,000, infiltration 7,500, attic ducts 8,000, internal gains 5,000. Total near 52,000 BTU per hour, about 26 BTU per square foot.

House two: mid-1990s frame and stucco, partially upgraded. Insulated 2x4 or 2x6 frame walls, dual-pane low-E windows retrofitted somewhere along the way, blown attic insulation topped up to a modern R, but the ducts still live in the hot attic and the envelope tightness is original. Representative picture: glazing 9,200, walls 5,500, ceiling 4,500, infiltration 5,500, attic ducts 7,500, internal 6,000. Total near 38,200 BTU per hour, about 19 BTU per square foot.

House three: 2020 spray-foam build. Foam-sealed attic that brings the roof deck inside the thermal envelope, so the attic runs within a few degrees of the house and the ducts in it are effectively indoors. High-R walls, low-SHGC dual-pane glazing, and infiltration a small fraction of the 1980s number. Representative picture: glazing 7,000, walls 4,000, ceiling 3,000, infiltration 2,500, ducts 1,500, internal 8,000. Total near 26,000 BTU per hour, about 13 BTU per square foot.

Same floor area. Loads of roughly 52,000, 38,200, and 26,000. A 500-square-feet-per-ton rule hands all three houses 4 tons: it cannot hold setpoint in house one, it is a sloppy oversize in house two, and it is nearly double the real load of house three. Notice also what the breakdown teaches: in the old house the envelope dominates and internal gains are a footnote; in the tight house the envelope has been beaten down so far that people, appliances, and glazing carry most of the load. That shift is why tight-house calcs feel "too small" to techs raised on old housing stock, and why the sanity bands later in this module are vintage-specific.

The attic duct penalty deserves its own sentence. Ducts outside the envelope are a double tax: they leak 20 to 30 percent of the air you conditioned (the A36 number), and the air that survives the trip absorbs heat through the duct walls the whole way. In the representative numbers above, attic ducts are 7,500 to 8,000 BTU per hour of load in the older homes, one of the largest single components, and effectively erased in the sealed-attic build. A duct sealing or duct relocation job changes the load calculation itself, which is the design-side echo of what A36 taught you from the service side.

PHOENIX FIELD NOTE

The vintage ladder is the fastest envelope read in this market. Pre-1990 Phoenix block homes with original glass and attic ducts routinely land in the 25 to 35 BTU per square foot band at 112 F design. Nineties and two-thousands frame-and-stucco with partial upgrades lands around 18 to 25. Tight 2015-plus construction with sealed attics lands around 12 to 18. Those attics, recall, run 140 to 160 F on a design afternoon, which is why duct location swings the calc here harder than in almost any other US market, and why a sealed-attic foam build loads like a different species. Use the bands to sanity-check a calc, never to replace one.

What oversizing actually costs

The intuition that bigger equipment is a safety margin dies hard, so walk the failure chain deliberately.

An oversized unit slams the thermostat to setpoint in a few minutes and shuts off. That is a **short cycle**, and on a mild day an oversized single-stage machine might run cycles of five to ten minutes, dozens of starts a day. Now collect the bills:

- **Comfort swings.** Big capacity in short bursts means the temperature near the thermostat oscillates while rooms far from it never get a long enough runtime to pull even. Long, steady runtimes are what mix a house; short blasts stratify it.
- **Weak moisture removal.** A cooling coil only dehumidifies after it has run long enough to get cold and wet, several minutes into each cycle, and when the unit slams off, moisture still on the coil re-evaporates into the

airstream. Short cycles spend most of their runtime in the worst part of the curve. In humid markets this is the headline failure: an oversized unit holds 75 F while the house feels clammy at high indoor humidity. In a dry market the latent stakes are lower, but the rest of the chain still collects in full.

- **Start-stress wear.** Every start is the hardest moment of a compressor's life: locked-rotor inrush, capacitor stress, contactor arcing, oil washout. Recall the failure ladder: capacitors alone are 21 percent of service calls, and the cycling machine multiplies starts per day by several fold. Oversizing converts envelope laziness into electrical and mechanical wear.
- **Efficiency loss.** Equipment is rated at steady state. A machine that spends its life starting and stopping never settles into the efficiency on its label, and the losses compound with the wear above.
- **The airflow check it writes.** A bigger unit demands more CFM, recall 400 per ton nominal, through whatever duct system exists. Drop a 5 ton onto ducts built for 3.5 and you have bought the C12 chain: high static, starved airflow per ton, cold coil, and a brand new machine running like a strangled old one. Oversizing is frequently an airflow failure wearing a sizing costume.

Undersizing has exactly one honest cost, and you should be able to state it fairly: a unit below the load cannot hold setpoint at design conditions, meaning the house drifts a few degrees above setpoint for some hours on the few hottest days. That is a real cost, which is why Manual S sets a floor of 90 percent, but notice the asymmetry: mild undersize hurts a handful of design hours per year, while oversize hurts every cycle of every mild day, which is most of the equipment's life. The window in the next section exists to balance those two costs deliberately, instead of by fear.

The field symptoms of a botched or skipped load calc are ones you already diagnose: short cycling, temperature swings, humidity complaints, comfort failures at design conditions, and repeat capacitor and compressor failures from cycling stress. At master level, when you see that cluster on a service call, "who sized this and how" belongs on your suspect list next to charge and airflow.

Manual S: matching a real machine to the calculated load

Manual J hands you three numbers: total cooling load, its sensible and latent split, and the heating load. Manual S turns them into an equipment selection, and its core rule is the one this trade violates most: **select from OEM expanded performance data at the actual design conditions, never from nominal tonnage.**

A nominal ton is 12,000 BTU per hour, and the nameplate earns that number at the AHRI rating point: 95 F outdoor ambient, 80 F dry bulb and 67 F wet bulb air entering the indoor coil, at rated airflow. That is a laboratory handshake, not a promise about your design day. At 112 F the condenser is rejecting heat into air 17 degrees hotter than the rating point, head pressure climbs, compressor mass flow drops, and total capacity falls with it, commonly in the neighborhood of 1 percent per degree above 95 F, with the exact curve belonging to the specific machine. And the indoor side of the rating point is just as wrong for a dry climate: 67 F entering wet bulb describes a humid house. A dry-climate home at 75 F and typical desert indoor humidity enters the coil around 62 to 63 F wet bulb, which shifts the coil's work away from latent and toward sensible. Both corrections come from the same place: the manufacturer's expanded performance tables, which list total and sensible capacity across outdoor ambient, entering wet bulb, and airflow.

With real capacities in hand, apply the **sizing windows**:

- **Single-stage cooling: total capacity between 90 and 115 percent of total load.** This is the workhorse rule.
- **Multi-stage equipment: up to about 120 percent**, because the low stage restores long runtimes on mild days.
- **Variable-capacity (inverter) equipment: up to about 130 percent**, because it can turn down far enough to stay matched, the same modulation logic that solved zoning in A36.
- **Sensible capacity must cover the sensible load, latent capacity must cover the latent load**, each checked at design conditions. A selection that covers total by drowning the sensible while missing latent fails the house in humid weather.
- **Heating gets checked too.** Furnaces carry generous allowed oversize because a furnace does not pay the short-cycle latent penalty a cooling coil does, but in a cooling-dominated market the cooling selection nearly always governs and heating rides along; verify it covers the heating load and move on. Heat pump heating capacity at the winter design temperature comes from the same expanded-data habit, and cold-climate selection is its own discipline this market rarely needs.

When two selections both land inside the window, **take the smallest machine that covers the load**. The smaller selection runs longer cycles, mixes the house better, asks less of the ducts, and starts fewer times per day. The window's upper half exists for availability and staging realities, not as an invitation.

One more Manual S habit that separates masters from catalog-shoppers: **airflow is part of the selection.** The expanded table is published at specific airflows, and the capacity you claim is only real at the CFM you can actually deliver. A dry climate wants the top of the 350 to 450 CFM per ton window to favor sensible capacity, and whether the duct system can deliver that airflow inside the blower's static budget is a question you already know how to ask from C12 and A36, and M38 answers in full.

Reading an expanded performance table at 112 F and beyond

Here is a representative expanded-data excerpt for a nominal 3.5 ton (42,000 BTU per hour) single-stage condenser matched to its rated coil, at 1,400 CFM and 75 F dry bulb entering air. The shape is what matters; the real numbers come from the real OEM document for the exact match you are quoting.

OUTDOOR AMBIENT	TOTAL / SENSIBLE AT 67 F ENTERING WB	TOTAL / SENSIBLE AT 63 F ENTERING WB
95 F	41,200 / 29,800	39,900 / 34,600
105 F	39,200 / 29,000	38,000 / 34,200
115 F	37,200 / 28,300	36,200 / 33,800

Read the lessons straight off the table:

1. **The nameplate is the best number the machine will ever post, and it posts it at 95 F in a humid house.** At 115 F and 67 wb this "3.5 ton" delivers 37,200, about 88 percent of nominal. Capacity derates as ambient climbs, exactly when the load peaks.
2. **Dry air moves capacity from latent to sensible.** Slide from the 67 wb column to the 63 wb column at 115 F: total drops slightly (36,200), but sensible jumps from 28,300 to 33,800, because the coil is condensing far

less moisture and spending its work on temperature. In a dry climate the 63 wb column is the honest one, and it flatters the sensible number you actually need.

3. **Interpolate to the design ambient.** Design is 112 F and the table prints 105 and 115. Interpolating the 63 wb column: total about 36,700, sensible about 33,900. Linear interpolation between adjacent rows is standard practice; reading the 95 F row and calling it a day is how nameplate-sized systems end up short on the only afternoons that matter.
4. **Airflow footnotes count.** The same table at 1,500 CFM would post more sensible and slightly less latent. The capacity you select at is the capacity at the airflow the duct system will really deliver.

Now finish the worked example. House two from the envelope section: total load 38,200, sensible 35,500, latent 2,700 (a 7 percent latent fraction, healthy for a dry climate). The Manual S window for single-stage: 90 percent of 38,200 is about 34,400, and 115 percent is about 43,900. The 3.5 ton above delivers about 36,700 total at 112 F, which is 96 percent of load: inside the window. Sensible 33,900 against 35,500 is 95 percent: covered within the window's tolerance. Latent delivered is the difference, about 2,800 against a 2,700 load: covered. A nominal 4 ton from the same family would deliver roughly 42,000 at design, 110 percent, also technically inside the window, and the smallest-that-covers rule says take the 3.5: longer cycles, easier ducts, fewer starts. The existing unit on that roof line of reasoning? A 5 ton delivering around 52,000 at design is 137 percent of load, far outside the window, and its short-cycling, capacitor-eating service history is the field evidence of exactly that arithmetic.

> **PHOENIX FIELD NOTE**
>
> Make 115 F-row reading a reflex here, not a refinement. Phoenix design is 112 F, the table rows bracketing it are usually 105 and 115, and several afternoons a year run past design. A unit selected off its 95 F AHRI numbers in this market is missing 10 to 15 percent of the capacity its installer thinks it has, at 5 PM in July, with the attic at 150 F feeding the duct penalty at the same time. When a tech asks why the "right-sized" unit cannot hold 75 F at 114 F outside, the answer is usually that nobody ever looked below the first row of the table.

Block load versus room-by-room

Manual J produces output at two resolutions, and a master tech knows which job needs which.

Block load treats the whole house as one box and produces the total sensible and latent loads. That is everything equipment selection needs: Manual S consumes block numbers. **Room-by-room** runs the same physics one room at a time and produces each room's load, which converts to each room's required airflow, which is the input Manual D needs to size trunks, branches, and registers in M38.

When is block enough? A straight changeout where the duct system is untouched, the envelope has not changed since the ducts were designed, and the comfort history is clean: block load picks the new box and the existing distribution keeps doing what it demonstrably already does. When is room-by-room mandatory? New construction, additions, any duct modification or redesign, zoning work, and any house with room-level comfort complaints, because a complaint room is a room whose airflow and load have never been formally introduced to each other. The honest default at master level: run room-by-room whenever the answer might change anything about the ducts, and remember that the marginal cost of room-by-room is a few more inputs in software that is already open.

The software walkthrough, tool-agnostic

Every credible load calc tool, whatever its name and whatever it looks like this year, walks the same flow, because the flow is the standard. Learn the flow and the brand never matters.

1. **Project and design conditions.** Enter the location; the tool loads design temperatures from its weather database. Verify them: right station, 112 F and 75 F cooling for this course's market, 70 F heating indoor. A wrong station or a default 95 F design here poisons everything downstream silently.
2. **Construction assemblies.** Define the wall, ceiling, floor, and window types: R-values, U-values, SHGC, from what you actually saw in the walk, not from optimistic menus. This is where the attic insulation depth you measured and the single-pane glass you flagged become numbers.
3. **Geometry, room by room.** Dimensions, ceiling heights, and each window and door with its size and orientation. Orientation entry is where careless calcs die: a west window entered as north quietly deletes the largest gain in the house.
4. **Infiltration.** Age-and-tightness category, or a blower door number when one exists. Resist the temptation of the "average" button on a house whose weatherstripping you watched daylight through.
5. **Internal gains.** Occupants (bedrooms plus one is the default convention), kitchen and appliance allowances, plus anything unusual you found.
6. **Duct system.** Location relative to the envelope, insulation R, and a leakage figure, assumed by condition or measured. You know from the envelope section how much this input moves.
7. **Outputs.** Block sensible, block latent, heating load, and room-by-room loads with airflow targets. Print or export everything; the output is a job-record document, not a screen you glance at.
8. **The sanity loop.** Before the calc leaves your hands, run the checks below, and when a number surprises you, change one suspect input at a time and watch the sensitivity. A load that swings 8,000 BTU on one toggle has just told you which input deserves a second site look.

Garbage in, garbage out is the whole risk. The software does arithmetic flawlessly on whatever fiction it is fed, and every input it asks for is something this module taught you to capture honestly with a tape, a compass, a flashlight, and an attic hatch.

Sanity checks before you trust any calc

A master tech treats a finished load calc like a finished diagnosis: it is not done until it survives cross-examination. The checklist:

1. **Design conditions verified.** The single most common silent error. 112 F outdoor, 75 F indoor cooling for this market, correct heating values, correct station.
2. **BTU per square foot lands in the vintage band.** Pre-1990 unimproved around 25 to 35, nineties partial upgrades around 18 to 25, tight 2015-plus around 12 to 18 at hot-dry design conditions. Outside the band is not automatically wrong, but it owes you an explanation you can point to in the inputs.
3. **Latent fraction is believable.** Dry climate: roughly 5 to 10 percent of total. A desert calc posting 20-plus percent latent has a humid-market default hiding in the infiltration or ventilation inputs.

4. **The component breakdown matches the house you walked.** Glass-heavy west-facing great room: glazing should top the list. Sealed-attic foam build: infiltration and ducts should be small change, and if infiltration posts as the biggest line, an input is wrong. The pie chart is the calc explaining itself; make it tell the same story your walk did.
5. **Floor area is yours.** Measured, conditioned space only, not the listing number.
6. **Duct location is set to reality.** The attic-versus-conditioned-space toggle moves more BTUs than almost any other single click.
7. **The existing equipment testifies.** This is the master move. The old unit is a measuring instrument with years of data: if the existing 5 ton held setpoint while cycling on a near-design afternoon, the true load is comfortably below 5 tons of delivered capacity, and a calc that says 60,000 BTU per hour is wrong no matter how the inputs look. If the old 3 ton ran flat-out from noon to midnight losing ground, the load exceeds about 30,000-some delivered. Runtime history brackets the answer before the software opens.
8. **Cooling and heating both checked.** In a cooling-dominated market the cooling number governs selection, but a heating load that looks absurd is still a flag that an input (usually infiltration or design temperature) is broken, and the same broken input is in your cooling number.

Then, and only then, the load goes to Manual S, the selection goes on the job record with the expanded-data excerpt that justified it, and the room-by-room sheet goes in the M38 folder. A number that survives all eight checks is a number you can defend to an inspector, an engineer, or the next master tech who pulls the file.

Common Mistakes

1. **Sizing the replacement to the nameplate being hauled away.** A like-for-like swap is a load calc done by a stranger decades ago, before the windows were replaced, the attic was topped up, and the ducts started leaking. The cost: re-signing someone else's guess for another 15 years. Every changeout gets its own Manual J.
2. **Using square feet per ton as a sizing method.** Three houses at the same square footage can load at 52,000, 38,000, and 26,000 BTU per hour. The rule of thumb hands them identical equipment, wrong three different ways. Bands are for sanity-checking a calc, never for replacing one.
3. **Selecting off nominal tons instead of expanded data.** A nominal 3.5 ton is not 42,000 BTU per hour at 112 F; it is about 36,000 to 37,000. Skipping the table builds a system that is short exactly and only on the afternoons it was bought for.
4. **Treating oversize as a safety margin.** The margin costs comfort swings, weak latent performance, multiplied starts, capacitor and compressor wear, and duct static the C12 way, on every mild day forever, to buy nothing on design day that the 90 to 115 percent window did not already provide.
5. **Accepting software defaults for design conditions, infiltration, or duct location.** The tool will happily calculate a humid 95 F market with sealed-attic assumptions for a leaky desert block home. Defaults are the most prolific author of garbage calcs in the industry.
6. **Entering window orientations carelessly.** A west sliding door logged as north deletes the biggest single gain in a hot-climate house. Compass in hand, every glazing entry, every time.

7. **Ignoring sensible versus latent in selection.** Covering total load while missing the latent column produces a house at setpoint and uncomfortable in humid weather. Both capacities get checked against both loads at design conditions, every selection.
8. **Skipping the sanity checks because the software printed something confident.** The calc that ships unexamined is the one with the wrong station, the default infiltration, and the 60,000 BTU answer on a house whose existing equipment already proved the load is half that. Eight checks, two minutes, every calc.

What Is Next

The load calc told you how much. Manual S told you which machine. M38 answers the question both of them handed forward: how does the right amount of air actually reach every room? Manual D, friction rates, trunk and branch sizing, and the static-pressure discipline you have been building since C12 finally get to design ducts instead of just diagnosing them. Bring the room-by-room sheet; it is the entry ticket.

Module Visuals

ENVELOPE HEAT GAIN PHOENIX

Same 2,000 sq ft, Three Very Different Loads

Representative cooling loads at 112 F design. The envelope sets the load, not the floor area.

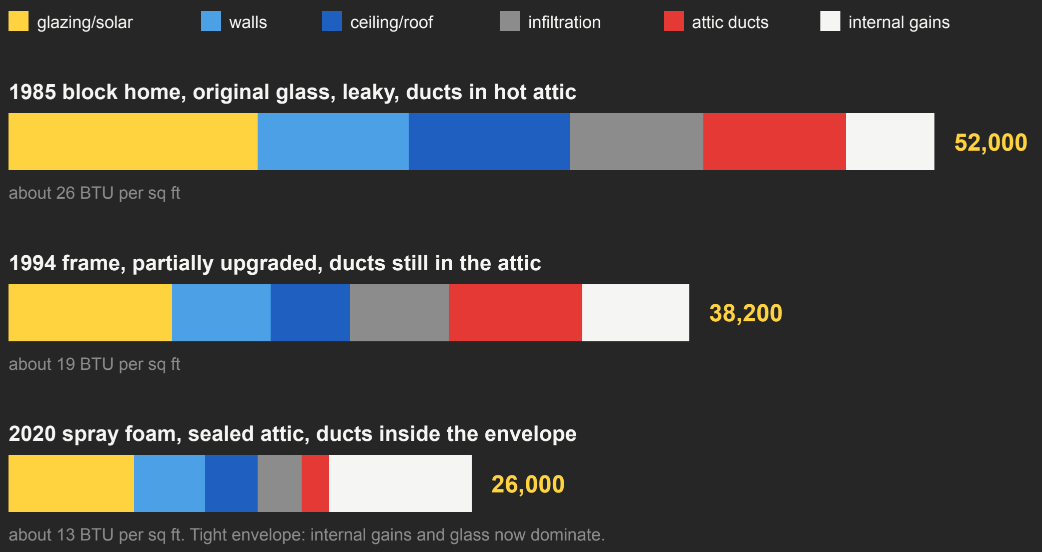

The 500 sq ft per ton rule hands all three houses 4 tons (48,000 nominal)

House 1: cannot hold setpoint. House 2: sloppy oversize. House 3: nearly double the load.

Bands are for sanity-checking a calc. Manual J replaces the guess with a number.

LOAD CALC SOFTWARE FLOW

The Load Calc Software Flow (Any Tool, Same Flow)

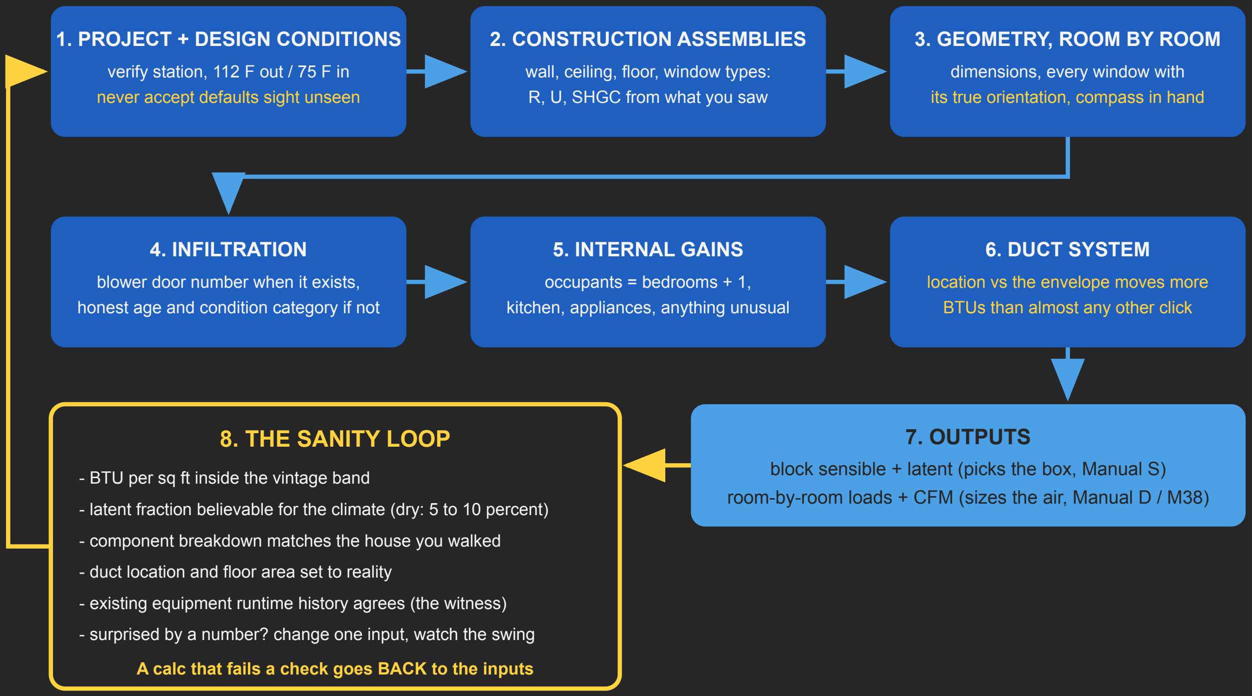

1. PROJECT + DESIGN CONDITIONS
verify station, 112 F out / 75 F in
never accept defaults sight unseen

2. CONSTRUCTION ASSEMBLIES
wall, ceiling, floor, window types:
R, U, SHGC from what you saw

3. GEOMETRY, ROOM BY ROOM
dimensions, every window with
its true orientation, compass in hand

4. INFILTRATION
blower door number when it exists,
honest age and condition category if not

5. INTERNAL GAINS
occupants = bedrooms + 1,
kitchen, appliances, anything unusual

6. DUCT SYSTEM
location vs the envelope moves more
BTUs than almost any other click

7. OUTPUTS
block sensible + latent (picks the box, Manual S)
room-by-room loads + CFM (sizes the air, Manual D / M38)

8. THE SANITY LOOP

- BTU per sq ft inside the vintage band
- latent fraction believable for the climate (dry: 5 to 10 percent)
- component breakdown matches the house you walked
- duct location and floor area set to reality
- existing equipment runtime history agrees (the witness)
- surprised by a number? change one input, watch the swing

A calc that fails a check goes BACK to the inputs

The software does flawless arithmetic on whatever fiction it is fed. Honest inputs are the whole job.

MANUAL J INPUTS MAP

Manual J: The Inputs That Set the Load

Design conditions: 112 F outdoor / 75 F indoor (1 percent design, never record extremes)

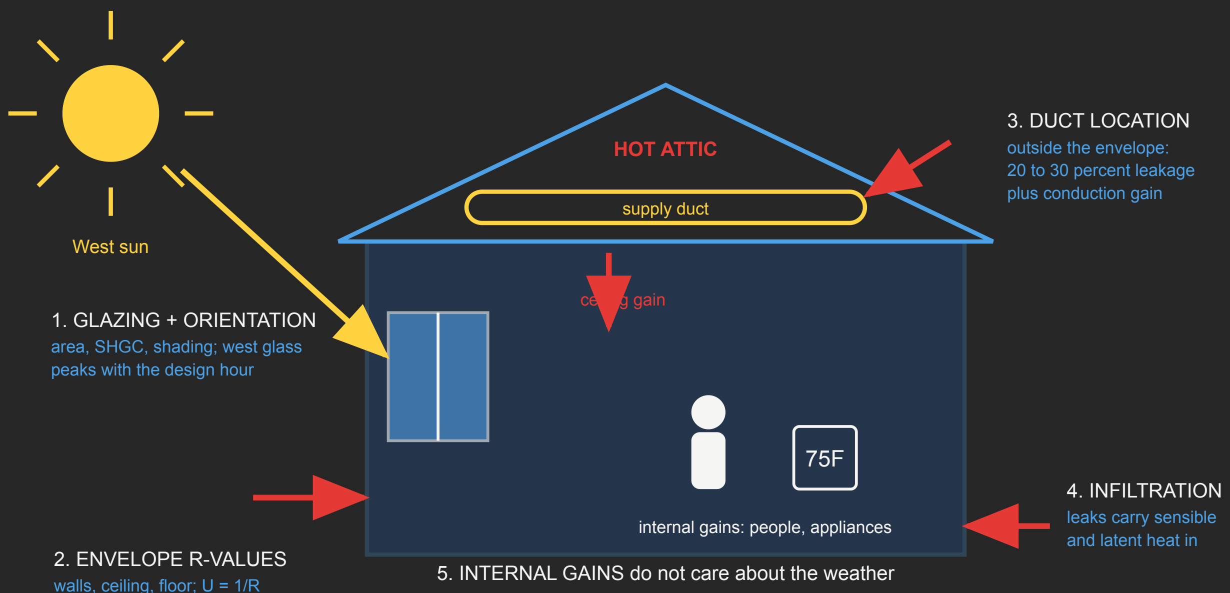

1. GLAZING + ORIENTATION
area, SHGC, shading; west glass peaks with the design hour

2. ENVELOPE R-VALUES
walls, ceiling, floor; U = 1/R

3. DUCT LOCATION
outside the envelope: 20 to 30 percent leakage plus conduction gain

4. INFILTRATION
leaks carry sensible and latent heat in

5. INTERNAL GAINS do not care about the weather

OUTPUT: total load in BTU per hour, split into SENSIBLE (temperature) + LATENT (moisture)

Both numbers carry forward to Manual S. No rules of thumb, every house gets a calculation.

MANUAL S SELECTION WINDOW

Manual S: The Sizing Window

Delivered total capacity at design conditions, as a percent of the calculated total load

Capacity comes from OEM expanded performance data at the real ambient and entering wet bulb.

Nameplate tons are rated at 95 F with 80/67 entering air. Nobody's design day lives there.

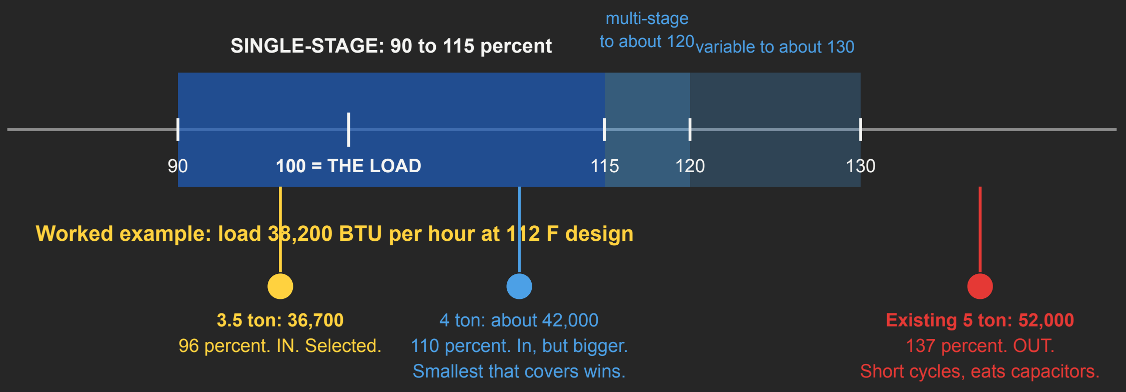

SINGLE-STAGE: 90 to 115 percent

multi-stage
to about 120

variable to about 130

Worked example: load 38,200 BTU per hour at 112 F design

3.5 ton: 36,700
96 percent. IN. Selected.

4 ton: about 42,000
110 percent. In, but bigger.
Smallest that covers wins.

Existing 5 ton: 52,000
137 percent. OUT.
Short cycles, eats capacitors.

Every selection passes three coverage checks at design conditions

1. TOTAL capacity inside the window for the equipment type
2. SENSIBLE capacity covers the sensible load 3. LATENT capacity covers the latent load

And the capacity claimed is only real at the airflow the ducts can deliver (C12, then M38).

OVERSIZING FAILURE CHAIN

Oversizing Is Not a Safety Margin, It Is a Failure Chain

OVERSIZED UNIT
capacity far above the load

→

SETPOINT IN MINUTES
thermostat satisfied, unit slams off

→

SHORT CYCLING
5 to 10 minute cycles, dozens of starts a day

The bills it sends, every mild day, for the life of the equipment:

COMFORT SWINGS
big bursts overshoot the thermostat; far rooms never get a runtime long enough to mix the house

WEAK LATENT REMOVAL
a coil dehumidifies only after minutes of wet runtime; short cycles re-evaporate the coil. Headline cost in humid markets

START-STRESS WEAR
inrush, capacitor stress, contactor arcing, oil washout. Recall: capacitors alone are 21 percent of service calls

DUCT STATIC (C12)
more tons demand more CFM through the same old ducts: high static, starved airflow per ton, cold coil chain

The asymmetry Manual S balances on purpose

Mild UNDERSIZE: loses a few degrees on a handful of design afternoons per year (the 90 percent floor exists for this).

OVERSIZE: pays the whole chain above on every cycle of every mild day, which is most of the equipment's life.

Field signature of a skipped or botched load calc

Short cycling + temperature swings + humidity complaints + repeat capacitor and compressor failures.
When you see the cluster, put sizing on the suspect list next to charge and airflow.