VENTIVA™

WHITE PAPER

Evaluation Framework Guide:

Air-Mover Cooling Technologies for Thin-and-Light Notebooks

A practical, technology-agnostic framework for OEM thermal and mechanical engineers to assess conventional and emerging air mover solutions before committing to a thin-and-light notebook system design

Copyright © 2026 Ventiva
Copyright © 2026 Ventiva

Evaluation Framework Guide:
Air-Mover Cooling Technologies for Thin-and-Light Notebooks

OVERVIEW
This Evaluation Framework Guide gives OEM thermal and mechanical engineers nine concrete, technology-agnostic evaluation steps for assessing the viability of any air mover solution - including conventional blowers and emerging technologies - before committing to a system design for a thin-and-light notebook computer.

Air mover evaluation should follow three stages: Design Feasibility, Test Conditions, and Measured System performance. Before beginning an evaluation, it is important to first agree on the product targets, including mechanical envelope, target sustained power, component and ergonomic temperature limits, the cooling power budget, and the acoustic noise ceiling. Steps 1–4 screen candidate designs; Steps 5–6 establish a test protocol that accurately represents in-system performance; Steps 7–9 assess the results.

WHY COOLING CLAIMS MISLEAD
Most disputed cooling claims trace back to one error: treating 'gross module heat dissipation', processor TDP, and total system power as interchangeable numbers. They are not.

- **Gross module rating** is the heat a module can dissipate under idealized test conditions. Any vendor claim for this rating assumes a particular temperature for the cooling air, as well as the surfaces of the target heat producing components. These assumed temperatures may or may not be representative of real in-system characteristics for a particular design. Gross module rating is often based on unrestricted inlet and outlet flow assumptions for the air mover and should be understood as a maximum possible cooling capacity that likely overestimates actual in-system cooling significantly.

- **Processor TDP (Thermal Design Power)** is a thermal design parameter defined by the xPU manufacturer which estimates the maximum sustained power dissipation of a processor when running a worst-case workload. It is not a measurement of sustained package power. Processor TDP describes only the power consumed by the processor package; it excludes power consumption/heat loads outside the package, such as SSD, memory, and display power.

- **Total system input power** includes the cooling solution, the processor, and all other components within the system. Typically, this is evaluated by measuring the total power delivered by the battery and/or the total power delivered from the AC wall adaptor (measurement typically does not include losses within an external AC adaptor itself, only power dissipated within the system chassis). This value may also include contributions from battery circuitry and charging inefficiencies within the battery itself.

AIR MOVER DESIGN FEASIBILITY

Step 1 — Account for required clearances when assessing mechanical fit

A device's stated thickness is rarely the full story — what matters for z-height budgeting is the total mechanical envelope it needs, including any tolerance clearances it requires beyond its own body as well as any open space maintained to serve as an inlet/outlet plenum to the device.

- **Top/bottom-breathing devices** (including many blowers and piezoelectric air movers) need a plenum above or below the device for air inlet, which adds to the effective height the system must budget, beyond the device's own thickness. The air mover vendor may or may not give guidance on recommended inlet clearances. Be aware that tight clearances can reduce performance dramatically for many devices.

- **Side-in/side-out devices** move air laterally and do not require the same extra vertical clearance, making them more amenable to constrained thickness systems, but they typically require lateral clearance to allow for an inlet plenum adjacent to the device.

Request the full mechanical keep-out zone — not just component thickness — for every candidate air mover, so the comparison accurately reflects the space required to accommodate each technology. Include lateral ducts, drive electronics, insulation, and assembly tolerances in that envelope. If the vendor does not provide inlet plenum clearances, ask for performance curves measured specifically at the clearances available within your design.

Step 2 — Evaluate thermal-module flexibility

If a cooling technology only achieves its rated performance when paired with one specific thermal module — for example, a vapor chamber of a single fixed size or geometry — that is a real constraint on the system, not a neutral implementation detail.

Ask every candidate vendor to demonstrate performance across more than one thermal module configuration, and treat a single-configuration-only result as a limitation on design freedom. A technology that performs well across a range of thermal solution geometries gives the OEM more latitude in battery placement, board layout, industrial design, and

future chassis variants; one that is locked to a specific module constrains the OEM into that module as well.

Step 3 — Require a published flow rate

An air-mover supplier should publish a volumetric flow rate (typically in CFM) and a pressure–flow curve at a stated drive power. Be aware that reported free-air maximum flow rates are very rarely achieved in-system. Pay close attention to how performance/flow drops off with increasing pressure.

Require inlet and outlet air temperatures and the density basis of the flow measurement. Some air mover's performance will vary significantly with temperature and pressure. For an aggressive thin and light notebook operating under a full workload, the inlet temperature to the air mover is typically 50-60 °C (yielding air densities roughly 10% lower than ambient air). Require the vendor to state whether the reported flow rate is Actual Cubic Feet per Minute (ACFM) or Standard Cubic Feet per Minute (SCFM), which is temperature and pressure corrected to some established reference condition. Due to altitude differences, ambient pressure can easily vary by 10% across labs in different global locations. To make an accurate assessment of device viability, it is critical to understand both the temperature and pressure assumptions, as well as the test conditions for any reported result.

Cooling outcome is a function of both the air mover's flow and the OEM's own thermal solution design (heat pipe or vapor chamber routing, fin stack, duct geometry, etc.). The air mover vendor may recommend a preferred layout to optimize performance against their device's flow curve, but it is the OEM's role — not the vendor's — to decide how that flow is used in a specific system.

A flow-rate figure that stands on its own — paired with its measurement temperature and pressure — lets an OEM independently model expected performance in their own design. Relying on a single, vendor-chosen "example notebook" result may not be representative of thermal performance in a different chassis or system architecture. In general, air flow testing should always be done in a test geometry that mimics the final system geometry as closely as possible (which often means rapid prototyping an entire chassis for testing, including motherboard, cables, thermal solution, etc.)

Step 4 — Check the thermodynamic limit

Use the airflow data from Step 3 to screen feasibility before building a functioning test system, then verify the result with measured installed flow and exhaust temperature. Devices that create high local velocities can improve local heat transfer rates but do not change the overall air-stream energy balance unless the heated air is removed from the system. Any internal air that recirculates, even at high velocity, does not actually remove heat from the system—it only mixes heat around within the chassis. For a given flow and allowable temperature rise:

$$Q_{total} = Q_{active} + Q_{passive}$$

$$Q_{total} = \dot{m}\, c_p\, \Delta T + Q_{passive}$$

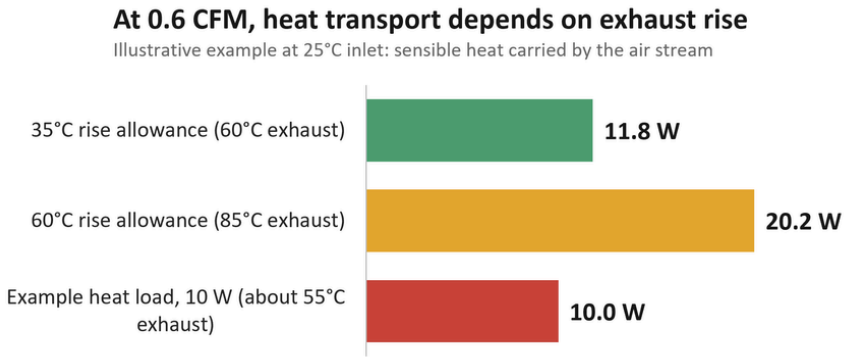
where Q_{total} is the total heat dissipated, $Q_{passive}$ is the passive cooling, and Q_{active} is the incremental active cooling provided by the air mover. Q_{active} is further bounded by the thermodynamic limit energy balance where $\dot{m}$ is its net mass flow of air through the system, c_p is the specific heat of the air, and ΔT is the mass-flow-weighted exhaust temperature minus inlet temperature. Note that these are mass flows of air entering and leaving the system as a whole. Any internal air that recirculates, even at high velocity, does not actually remove heat from the system—it only mixes heat around within the chassis.

Direct chassis-to-ambient heat loss via natural convection and radiation is a separate heat transfer path that supplements the total cooling in any system. This cooling is often referred to as "passive cooling" and is largely independent of air mover selection. Passive can be hard to accurately predict for a variety of reasons so we will measure it directly in Step 7 (typical expected value is 5-10 W for a thin-and-light notebook).

The ΔT in the thermodynamic limit essentially defines the maximum exhaust temperature, so it matters for two practical reasons:

- **Ergonomic limits.** Exhaust air that is too hot is uncomfortable or unsafe to sit near, which caps the usable ΔT regardless of what the air mover could theoretically achieve.
- **Localized hot spots.** Hot exhaust will create a hot spot at the outlet vent and may create additional hotspots at other locations within the system due to the passage of warm air.

Use installed net flow and the OEM's allowable exhaust temperature to sanity check a Gross Module Cooling capability claim; remembering that maximum free-air flow gives only an optimistic upper bound. Contact-surface limits and exhaust-air limits are different. Measure both rather than assuming a universal allowable temperature rise.



Illustrative example only. The 35°C and 60°C rises are assumed design allowances, not universal limits. At 25°C inlet, a 60°C rise means 85°C average exhaust, not surface temperature. In this example, if the maximum allowed exhaust temperature is 60°C and the net flow rate through the system is 0.6 CFM, then the maximum possible cooling that could be provided by the air mover is 11.8 W.

TEST CONDITIONS AND INSTRUMENTATION

Step 5 — Define temperature limits and monitor hot spots

Before testing, define the allowed skin (chassis surface) and maximum component temperatures and instrument the limiting locations. In ultra-thin systems, the binding constraint is frequently the skin temperature ("ergonomic limit"), not the internal junction temperature of the xPU—yet many evaluation setups only log the processor's own internal sensor(s), not chassis temperature.

- If the system has onboard thermistors placed at the relevant skin locations and used for closed-loop throttling, log their readings directly.

- If it does not, instrument the chassis independently. IR cameras are useful for identifying surface hotspot locations but are often difficult to calibrate. Surface-mount thermocouples are the preferred method. If using IR thermography exclusively, be sure to correct for surface emissivities and reflections.

A cooling solution that looks adequate by xPU junction temperature alone can and often does still fail a skin-temperature limit in a thin, high-density system. Hot-spot location(s) should be reported alongside peak value, since a highly localized hot spot is a different design problem than a broadly elevated surface.

Step 6 — Standardize the load

Cooling comparisons are only meaningful if every candidate is tested against the same electrical and thermal load. Mixed workloads, undisclosed durations, or vendor-chosen benchmarks make cross-vendor comparison less meaningful.

Recommended load generation:

- **Intel-based platforms:** use Intel Power Thermal Analysis Tool (PTAT), where supported, or Prime95 as a defined stress workload, with power/thermal telemetry alongside it. Both of these options use a simulated steady workload at a controllable level to exercise the xPU and other components in the system. Do not run two stress generators simultaneously unless that combination is intentional and documented.

- **Other xPU vendors**: use the vendor's equivalent power/thermal monitoring utility, paired with the same class of stress tool, to match workload intensity.

- **Fix and disclose:** ambient temperature, workload/tool version, test duration (minimum 20 minutes), thread count, instruction-set settings, BIOS power policy and duration. The system should be plugged and battery at full charge level to avoid variations in charging power during the test. Log the actual package power, xPU and chassis temperatures, any throttling or warning notifications from the monitoring utility, and

processor frequency. Define steady state and repeat key runs; report variation. Distinguish sustained from burst power.

Use the same workload type and settings across candidate designs. For the capacity test in Step 7, vary the power limit to find each configuration's maximum sustainable power under the agreed temperature and acoustic limits. For a fixed-load comparison, hold measured power constant by fixing the workload to an established level.

SYSTEM PERFORMANCE AND EFFICIENCY

Step 7 — Separate active and passive cooling

Every laptop chassis dissipates some heat on its own via radiation and convection from external surfaces including the chassis, keyboard, palm rests, and even the underside surface. A system with a very small, low-flow air mover can still report a relatively high total cooling number simply because it is riding on top of a large passive cooling budget, not the air mover. Total passive cooling for a thin and light notebook is typically in the range of 5-10 W. Passive cooling is greater for larger systems and those with better heat spreading within the chassis. It also varies significantly with surface temperature limits and surface material choices.

Test protocol:

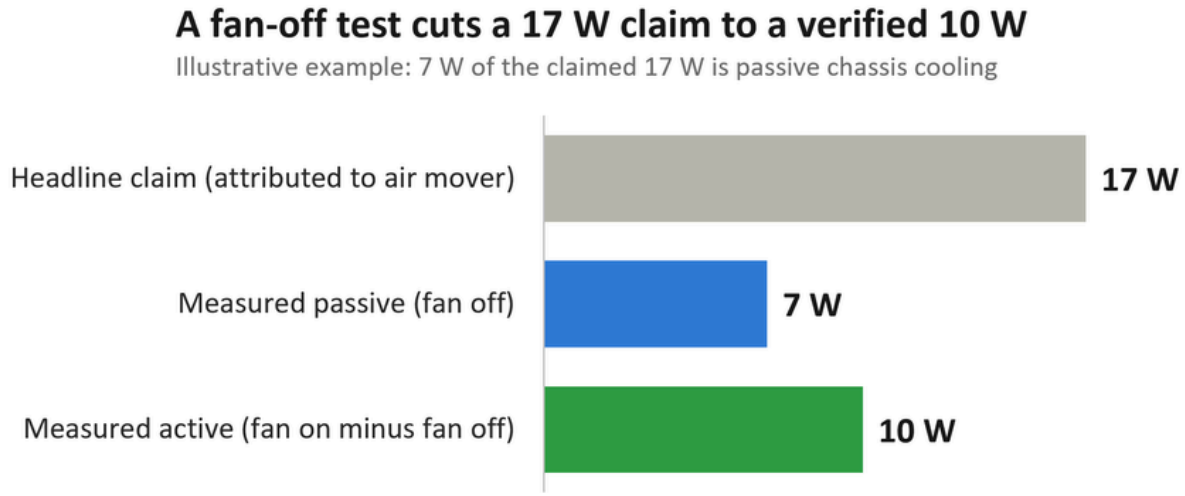
1. Conduct this test in a controlled air flow environment. A large, perforated enclosure is commonly used to minimize the impact of air currents created by HVAC and workers in the area of the test.

2. With the air mover disabled but still installed, increase the workload power limit to find the maximum sustainable processor power within agreed skin and junction temperature limits at a fixed ambient. Each step increment in power should be held for a minimum of 20 minutes. The maximum power achieved within these limits is assumed to be the system's passive cooling capability.

3. Repeat with the air mover enabled, using the same workload type and temperature limits. Adjust the power limit to find the new maximum sustainable processor power; record which temperature limit is reached in each state (i.e., which component temperature limit is hit or which skin temperature limit and location).

4. The difference in sustainable power between these two measurements is the incremental active cooling capacity. The fan-off result is the passive baseline of that assembly. The fan-on result is the total system cooling available.

Any vendor claim for a system power cooling capability that does not disclose the fan-off baseline (or other measure of passive cooling) cannot be used to isolate what the device contribution to the total system cooling capability actually is.



Illustrative example, not a specific vendor's published figures. Here, active means incremental sustainable processor power at the same temperature limits.

Step 8 — Measure Coefficient of Performance (COP)

Once the active cooling portion is isolated (Step 7), evaluate it on an efficiency basis, not just a raw-watts basis. A device that removes more heat but draws disproportionately more power is not necessarily the better system-level choice, especially in a battery-powered product.

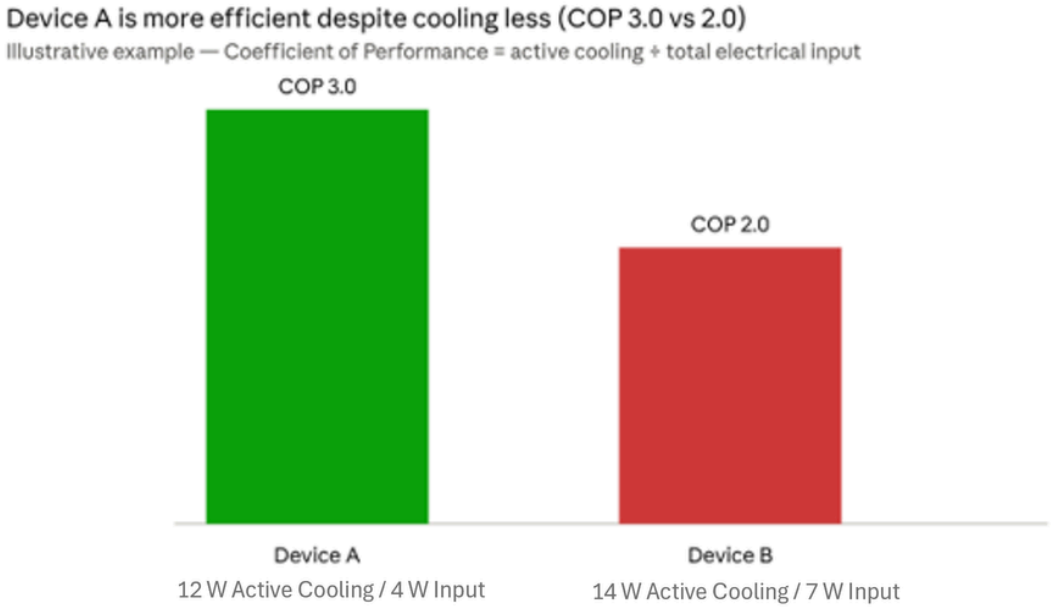
$$\text{COP} = \frac{\text{Incremental active cooling (W)}}{\text{Total air mover electrical input (W)}}$$

The denominator must include **all** power the air mover consumes to operate — not just the rated module input. This means:

- Driver and control electronics power, if separate from the module
- Any step-up/step-down conversion losses in the power supply feeding the air mover
- Idle or standby draw during duty-cycled operation

Report this system COP at defined drive settings using the Step 7 incremental active cooling capacity, not total processor power. Measure average cooling input power upstream of its dedicated converters; do not add converter losses twice or subtract cooling power again from the incremental active cooling capacity.

It is also worth evaluating the idle state power draw of the air mover (if any) when the system is powered on, but not running a workload that requires the thermal solution to be active. This state of representative of most battery-life assessments, so the idle power of the air mover will directly impact notebook battery life.



Illustrative example, not a specific vendor's published figures

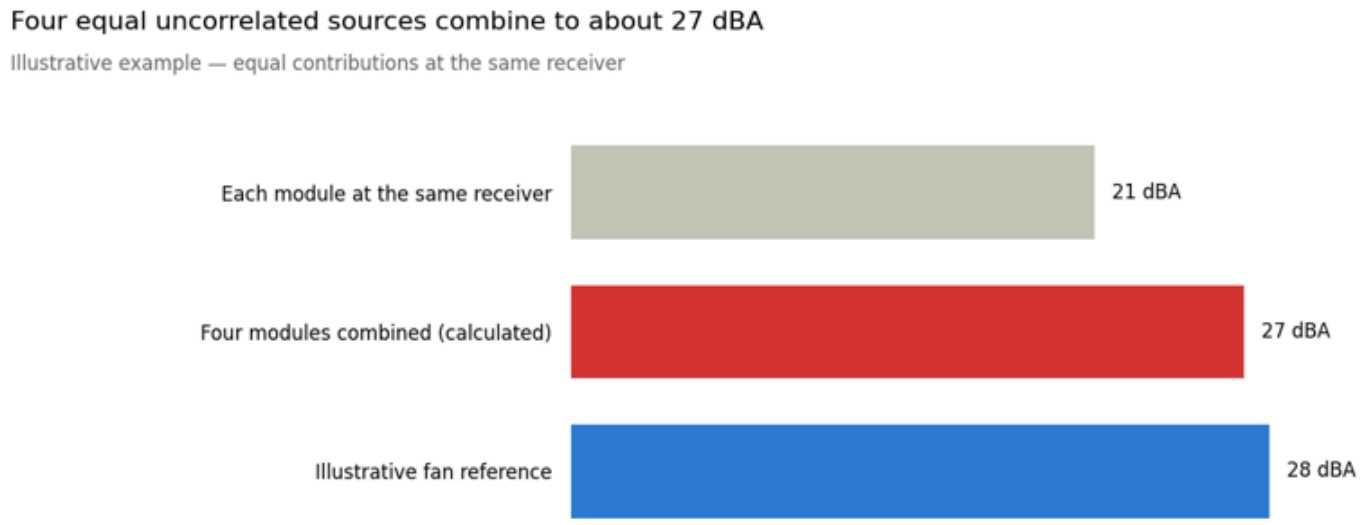
Step 9 — Check acoustics at the system level, not per module

The system chassis will act as an acoustic cavity for any air mover, amplifying some frequency ranges and diminishing others. To correctly assess the acoustic noise level created by an air mover in any application, it is imperative that acoustic testing be done at a system level, not at the device level. An acoustic rating is frequently quoted for a single module in isolation, then inappropriately used to describe the noise level of a system build from several of them operating together. Independent sound sources do not add arithmetically in dB — but they do add.

Also, be careful about acoustic measurements being carried over into marketing language such as "silent" or “quiet”. These descriptors are qualitative and subjective. A noise measurement reported in dBA is insufficient to determine human audibility – which depends not only on the overall noise level but also on the frequency distribution of that noise energy generated by a particular device.

$$L_{\text{total}} = L_{\text{single}} + 10\,log_{10}(N)$$

For N equal, uncorrelated sources, this gives roughly +6 dBA for four modules and +3 dBA for two. Four sources each contributing 21 dBA at the same receiver would combine to about 27 dBA. Actual laptop results depend on mounting, directionality, and enclosure effects; verify at equal sustained load and report distance (or use "Operator Position" as defined by the ECMA 74 standard), background noise level, and operating mode.



Beyond the combined sound level, request or independently verify:

- **Measurement, not assumption, of ultrasonic content.** Many solid-state actuators operate at ultrasonic carrier frequencies, which sit outside typical hearing range — but that should be confirmed by measurement, not simply asserted. The primary driving frequency may also not be the only acoustic band in which noise is created. Measuring ultrasonic frequencies can also confirm that any noise created is inaudible to pets and wildlife as well as humans.

- **Tonal or beat artifacts.** Multiple resonant actuators and their control modes can create audible sidebands, subharmonics, or beating under some operating conditions, even if the fundamental tone is inaudible.

- **Structure-borne noise.** Poor mounting techniques can couple actuator vibration into a large, light chassis panel, producing audible buzz, rattle, or resonant amplification that a module-level anechoic measurement will not capture. Even with a quality mounting strategy, vibro-acoustic coupling can lead to significant noise amplification. This effect is responsible for the high-pitched tone commonly audible from digital electronics (although mitigation strategies have greatly reduced the prevalence in more modern devices).

- **Mockup acoustic testing.** If a complete and operational system is not available a detailed 3D rapid prototype may be used to evaluate air mover acoustics. This is not perfect, but still preferable to open-air acoustic testing which is largely not representative of in-system performance. Any mockup should have detailed internal and vent geometries that match the planned system exactly, including cables, motherboard layout, etc. It is also critical that the air mover mounting method and materials are matched to that of the planned system.

Claims of "quiet" or "silent" should be backed by complete-system measurements under realistic mounting and workload conditions, not inferred from a single module's datasheet.

EVALUATION CHECKLIST

Step	What to Request or Measure	Why It Matters
1. Mechanical clearance	Full keep-out envelope, including any required inlet plenum(s)	Device thickness alone can understate the z-height a technology actually requires
2. Thermal module flexibility	Performance across more than one thermal module configuration	A single-module-only result constrains system design freedom
3. Flow rate	Reported in a standard unit (e.g., CFM), density basis, pressure– flow curve, drive power, and installed flow	Allows the OEM to model performance in its own design rather than relying on one vendor example
4. Thermodynamic limit	Spot coolers enhance local convection, but heat is only dissipated when the hot air is removed from the system. Installed flow and allowed exhaust rise should be checked against $Q = \dot{m}c_p\Delta T$	Defines minimum system flow requirements, bounds temperature in the exhaust air stream; chassis heat loss is separate
5. Temperature limits and hot spots	Agreed skin/component limits and sensors *at the actual hot spots*	Skin temperature, not junction temperature, is often the limiting constraint in thin systems

Step	What to Request or Measure	Why It Matters
6. Standardized load	Defined PTAT or Prime95 workload (or equivalent), with telemetry and steady-state criteria (typically 20 minutes minimum)	Enables cross-vendor comparison of results
7. Active vs. passive	Fan-OFF/ON maximum sustainable power at the same temperature limits	Measures incremental active cooling capacity enabled by the air mover in that assembly
8. COP	Incremental capacity (W) ÷ total cooling supply input (W), including conversion losses	Flags devices that trade battery life for marginal cooling gains
9. Acoustics	System-level dBA with all modules running in an installed condition, plus a check for tonal/beat artifacts and structure-borne noise	A single-module “silent” rating often understates combined noise; uncorrelated sources add as $+10 \cdot \log_{10}(N)$; installed air movers are typically significantly louder than they are in open air.

Disclaimer

This Evaluation Framework Guide is provided for informational purposes only as a general engineering framework for evaluating air-mover cooling technologies in thin-and-light notebook systems. It is intended to support consistent, technology-agnostic evaluation of alternative cooling approaches and is not an industry standard, certification program, design specification, or performance guarantee. Although prepared by Ventiva, the framework is intended to be applied equally to conventional blowers, solid-state air movers, and other current or future cooling technologies. Actual system performance will depend on factors including system architecture, thermal design, airflow path, component selection, operating conditions, and test methodology. Engineers should exercise independent judgment and conduct testing appropriate to their specific products and requirements. Any examples, calculations, assumptions, or illustrative data in this guide are provided solely to demonstrate evaluation methodology and should not be interpreted as performance claims or representations regarding any specific product, technology, or vendor. This guide reflects current engineering practices and may be revised as technologies, test methods, and industry practices evolve.