How to Design a TEC Cooling System

Sizing the TEC Module, Heat Sink, Sensor, and Mounting Stack for Precision OEM Thermal Control

White Paper AWP-TECC-03 · Revision 2.9 — June 2026

Analog Technologies, Inc. · San Jose, California

June 2026 · www.analogtechnologies.com

Summary — How to Design a TEC Cooling System

Designing a TEC cooling system means sizing four components together against the application's thermal load, ambient envelope, and target setpoint: the TEC module (heat-pumping capacity at the worst-case ΔT), the heat sink (Q_hot rejection at worst-case ambient), the temperature sensor (placement and mounting that determines what the loop can resolve), and the mechanical stack that holds them together (TIM, compression, orientation, sealing). The TEC controller — selected per companion paper AWP-TECC-02 (TEC Controller: What It Is and How to Choose One) — connects to this system and provides the closed-loop regulation, but the achievable stability at the load is set by the system around the controller, not the controller alone. This guide walks the engineer through the full system-design methodology: thermal-load characterization (active and passive contributions), TEC module sizing against datasheet performance curves (with the Q_max + ΔT_max trap to avoid), heat-sink sizing for Q_hot with environment-appropriate margin, sensor selection and mounting for low thermal resistance to the load, the TEC mounting and assembly procedure between load and heat sink, the additional design considerations for sub-10 mK-class stability targets at the sensor under defined laboratory conditions (achievable stability at the controlled object depends on the full thermal stack), and worst-case validation testing that proves the design before production. Five worked application examples (laser-diode wavelength stabilization, PCR / high-current instrumentation, outdoor kiosk cooling, miniaturized photonic sensors, LIDAR receivers) each produce a starting design path — TEC module direction, heat-sink direction, controller family, sensor placement, evaluation board, validation step.

How do I design a TEC cooling system?

Short answer: Designing a TEC cooling system is a six-step process — characterize the cold-side thermal load, define the worst-case ambient envelope, select a TEC module from its datasheet performance curves at the actual operating point, size the heat sink for the resulting hot-side heat, choose the sensor placement and mounting that determines what the loop can resolve, and validate the assembled stack against a worst-case test envelope before production. The TEC controller is selected separately using companion paper AWP-TECC-02 after the thermal-stack inputs are known.

Key equations (final). The relationships that anchor system-design:

$$Q_{hot} = Q_c + P_{electrical}$$

$$P_{electrical} = V_{TEC} \times I_{TEC}$$

ΔT_TEC = T_hot − T_cold,setpoint

R_θSA ≤ (T_hot,max − T_ambient,max) / Q_hot

T_hot = T_ambient + Q_hot × (R_θSA + R_TIM,hot + R_spreader, if significant)

Quick definitions: Q_load = heat that must be removed from the controlled object; Q_c = heat the TEC pumps at the selected operating point, read from the TEC datasheet performance curves; P_electrical = V_TEC × I_TEC at the same operating point; Q_hot = heat rejected at the hot side = Q_c + P_electrical. Full definitions follow.

Variable definitions: Q_c is the heat the TEC actually pumps at the selected operating point — the actual cold-side heat being moved at steady state, normally equal to Q_load at steady state (for ramping or thermal-cycling loads such as PCR blocks, the TEC must additionally pump transient ramp energy, sized from Q_transient = m · c_p · dT/dt; the total cooling capacity required during a ramp is therefore Q_c,required = Q_load,steady + Q_transient). It is read from the TEC datasheet performance curves at the operating ΔT_TEC and T_hot. Q_c is not the TEC module's unused maximum cooling capacity — using catalog Q_max here will over-estimate Q_hot and over-size the heat sink. P_electrical is the TEC's electrical input at the same operating point. ΔT_TEC is the temperature difference the TEC must develop between its hot and cold faces to hold the cold-side setpoint. R_θSA is the heat-sink-to-ambient thermal resistance under actual airflow and enclosure conditions (sometimes written R_θSA). If TIM and base-spreading resistance on the hot side are non-negligible, subtract their contribution from the allowable R_θSA before selecting the heat sink — the full thermal path from TEC hot face to ambient is R_TIM,hot + R_spreader + R_θSA. T_hot,max is the maximum hot-side temperature that still allows the TEC to pump Q_c at the required ΔT_TEC — read from the TEC datasheet performance curves, not the TEC's absolute maximum rating. For first-pass estimates with Q_c selected close to Q_load, Q_hot ≈ Q_load + P_electrical is acceptable; final heat-sink sizing uses the datasheet operating point.

Notation: R_θSA is the heat-sink-to-ambient thermal resistance used throughout this paper; some TEC controller and heat-sink datasheets use the equivalent symbol R_hs.

Read this paper if...

You need to design the thermal stack for an OEM precision-cooling product: TEC module sizing from datasheet performance curves, Q_hot calculation, heat-sink sizing for worst-case ambient, sensor placement and mounting (step-hole geometry), TIM and compression, condensation and sealing, and worst-case validation. If instead you need to choose the TEC controller — current and voltage headroom, output-stage topology, precision grade, suffix variant, and evaluation board — the companion paper AWP-TECC-02 (TEC Controller: What It Is and How to Choose One) is the right starting point.

Critical caveats. (a) Size the TEC from datasheet *performance curves* at the actual operating point — never from Q_max and ΔT_max read separately, which are non-simultaneous end-points. (b) Mount the temperature sensor as close to the controlled object as the geometry allows, in a step-hole geometry rather than mid-hole in potting compound — stability quoted in controller datasheets is measured at the sensor under defined laboratory conditions, and the temperature at the controlled load may differ. (c) Use thin uniform thermal-interface material on both TEC ceramic faces; sealing, condensation, and uniform compression matter. (d) Validate the assembled stack against a worst-case test envelope (hot ambient × deepest setpoint × peak load × low V_PS) before locking the design.

Pass/fail metrics — what to measure under the worst-case envelope

Defining test conditions is necessary but not sufficient — the engineer must also define what to measure and what pass/fail criteria apply. For a TEC cooling system, record at minimum:

- **Sensor stability** (peak-to-peak and standard deviation over a defined window, e.g., 60 min after thermal soak).
- **Controlled-object / load stability** (if instrumentable) at the location that matters for the application.
- **Maximum hot-side temperature T_hot** at peak load and peak ambient, against the TEC datasheet limit for the operating point.
- **Controller saturation indicators** — output duty cycle, current-limit flag, voltage-headroom margin to V_PS rail.
- **Time-to-setpoint and overshoot** on a step from ambient (and on a step from the opposite-direction extreme, if the application bi-directional regulates).
- **Recovery from a load step** (step the active load up and down, measure settling time and excursion).
- **Condensation observation** on the cold zone at deepest setpoint and worst-case humidity.
- **Fan / airflow condition** (speed, RPM, filter loading at end-of-clean cycle).
- **Supply-current margin** from V_PS at worst-case ambient against the worst-case current the controller demands.

Specific pass/fail thresholds for each metric are application-specific and must be derived from the OEM's system specification, not from this guide.

Controller note. Once the TEC operating-point current and voltage, the heat-sink design, the sensor placement, and the stability target are defined by this guide, select the controller using companion paper *AWP-TECC-02: TEC Controller: What It Is and How to Choose One* — which covers controller architecture, output-stage topologies, precision grades, ATI controller families, evaluation boards, and the six-step controller-selection methodology.

Why ATI for TEC-controller-centered cooling-system design

ATI is a supplier of matched TEC controllers, TEC modules, precision thermistors, and evaluation boards, supporting OEM customers across North America, Europe, and Asia. Sourcing TEC controller, TEC module, and precision thermistor from a single vendor can reduce the engineering effort that would otherwise be spent reconciling datasheets across separate manufacturers, although the magnitude of the saving depends on the specific design and qualification scope. ATI supports OEM programs with lifecycle communication, replacement guidance, and applications-engineering support where available; for any specific part-level migration questions, contact ATI sales.

In a hurry? Jump to *§4 System Design Methodology* for the six steps, *§5 sub-10 mK system-design considerations* for high-precision designs, *§5d Mounting and Assembly* for the TEC stack procedure, or *§8 Application Examples* to see worked solutions in your application area. The selection-support form at the end is short and helps ATI applications engineering review qualified OEM projects and suggest a starting thermal-stack direction, a controller-selection path, and a suggested evaluation path.

Companion Paper Map

Topic	Canonical paper
Beginner / visual explanation of a TEC controller	AWP-TECC-01 — What Is a TEC Controller? (Visual Engineering Guide)
TEC controller — what it is and how to choose one	AWP-TECC-02 — TEC Controller: What It Is and How to Choose One
TEC cooling system design — TEC module sizing, heat-sink sizing, sensor placement, mounting, validation (this paper)	**AWP-TECC-03 (this paper)**
Thermoelectric cooler fundamentals — what a TEC module is	AWP-TECM-01 — What Is a Thermoelectric Cooler?
Multi-TEC arrays and cascade systems for high cooling loads	AWP-TECM-03 — Multi-TEC Thermal System Design

Part I — Foundations

§1. Why TEC Cooling System Design Matters

Designing a TEC cooling system that meets its temperature-stability, efficiency, and lifetime targets in production depends on getting the system around the controller right — the heat sink sized for worst-case ambient, the TEC module operated at a defensible point on its datasheet performance curves (not at the Q_max / ΔT_max trap), the sensor mounted close to the controlled object with low thermal resistance, the mounting stack TIM and compression specified correctly, and the assembled system validated against a worst-case envelope before production release. The controller closes the loop, but it cannot regulate what the sensor cannot see, the heat sink cannot remove, or the TEC cannot pump.

Consequences of skipping system-design discipline. Heat-sink saturation at hot ambient; TEC current and voltage demand exceeding the controller's headroom under field conditions; the regulated sensor temperature differing from the controlled object temperature by tens or hundreds of millikelvin because of sensor-to-load thermal resistance; cold-side condensation when cooling below dew point without sealing or moisture management; warranty events that the design team did not see during typical-condition testing. Each is addressed in the methodology below.

Controller-side failure modes (insufficient current or voltage headroom, wrong topology for the noise environment, inappropriate precision grade) are covered in companion paper AWP-TECC-02, §1. This guide focuses on the system-side failure modes listed below.

This guide is a structured framework for system-side design. It walks the engineer from the cold-side thermal load through TEC module sizing against datasheet performance curves, heat-sink sizing for Q_hot at worst-case ambient, sensor selection and placement, the TEC mounting and assembly procedure, and the worst-case validation envelope that proves the design before production. Controller selection — architecture, output-stage topology, precision grade, ATI product families — is the subject of companion paper AWP-TECC-02, which this paper feeds into at Step 4 (controller match) once the system inputs are defined.

What is a TEC cooling system?

A TEC cooling system is the complete physical assembly that holds a sensitive component at a regulated temperature using a thermoelectric (Peltier) cooler. It comprises five elements engineered together: the TEC module (heat-pumping element between cold plate and hot side), the heat sink (rejects Q_hot to ambient), the temperature sensor (closes the loop), the mechanical mounting stack (TIM, compression, sealing — sets the achievable thermal resistance between load and sensor), and the TEC controller (the closed-loop electronics that drive bidirectional current into the TEC and is selected per companion paper *AWP-TECC-02*). System-level stability at the controlled object depends on all five elements; the controller alone is necessary but not sufficient.

Common TEC cooling system failure modes after correct component selection

Short answer: Common TEC cooling system failure modes include undersized heat sinks, incorrect TEC datasheet interpretation (Q_max / ΔT_max read separately), insufficient controller voltage headroom, poor sensor placement, excessive thermal-interface material, uneven mounting compression, wrong TEC orientation, condensation below dew point, and incomplete worst-case validation.

Even after the TEC module, controller, sensor, and heat sink are correctly specified, the assembled system can fall short of its stability and lifetime target for any of the following causes. Each is addressed in the methodology below.

- **Undersized heat sink at worst-case ambient.** Heat sink sized at typical ambient saturates at the field-worst-case ambient, T_hot climbs, and the TEC can no longer pump Q_load at the required ΔT.
- **TEC module sized from Q_max and ΔT_max read separately.** Q_max (at ΔT = 0) and ΔT_max (at Q_c = 0) cannot both be achieved at the same operating point — see §5.
- **Insufficient controller voltage headroom at hot ambient.** Current headroom alone does not solve voltage headroom; the loop saturates at deep ΔT — see AWP-TECC-02 §4 Steps 4 and 4b.
- **Poor sensor placement.** Sensor mounted far from the controlled object, in a straight through-hole filled with potting compound, or thermally short-circuited to ambient — measured temperature does not represent the load temperature.
- **Excessive or non-uniform thermal-interface material.** Thick TIM, voids, or uneven compression adds thermal resistance between TEC and cold plate (or TEC and heat sink) — directly degrades cooling.
- **Wrong TEC orientation.** For standard ATI TEC modules, the labeled side is normally the cold side. Verify orientation on the specific module drawing before assembly, especially for custom, sealed, high-temperature, or non-standard modules. Reversed orientation drives heat the wrong way.
- **Condensation below dew point.** Operating the cold surface below the local dew point without sealing or moisture management causes water condensation, corrosion, electrical leakage, and optical contamination.
- **Incomplete worst-case validation.** System tested only at typical room conditions; field failure at worst-case ambient × deepest setpoint × peak load × low V_PS is a common field-failure mode in OEM thermal-control products.

Scope of this guide

This paper covers TEC cooling system design for OEM thermal-control products with cooling-load requirements from 0.5 W to approximately 60 W per TEC channel. It addresses thermal-load characterization, TEC module sizing, heat-sink sizing, sensor placement, mounting and assembly, and worst-case validation. Loads above the

per-channel range are typically served by paralleled or multi-TEC arrays, covered in the companion paper *AWP-TECM-03: Multi-TEC Thermal System Design*. Controller selection — architecture, output-stage topology, precision grade, ATI product families, evaluation boards — is the subject of the companion paper *AWP-TECC-02: TEC Controller: What It Is and How to Choose One*, referenced throughout. The guide does not cover bench-top laboratory instruments, large industrial cooling assemblies, or two-phase refrigeration.

Who this guide is for

The primary reader is the OEM design engineer integrating a TEC controller into a product: photonics and laser-diode designers, diagnostic and laboratory instrumentation engineers, imaging-system designers, and the technical purchasing managers who specify the bill of materials. The guide assumes working knowledge of analog electronics, basic control theory, and the Peltier effect. Readers new to thermoelectrics should consult the companion paper *AWP-TECM-01: What Is a Thermoelectric Cooler?* before proceeding here.

Quick-reference starting point

If the application's thermal load, target temperature, and worst-case ambient are already known, the quickest path to a starting controller family is the comparison table and worked examples in companion paper *AWP-TECC-02*, §1 and §8. This guide concentrates on the system-design inputs that feed that controller selection — TEC module, heat sink, sensor placement, mounting, and validation envelope.

Recommended starting design-path by application — system inputs

The table below summarizes the system-design inputs for five common applications: thermal load range, ΔT direction, TEC module size direction, heat-sink direction, sensor-placement risk, and dominant mounting / condensation issue. The matching ATI TEC controller family for each application is in the companion controller-selection paper *AWP-TECC-02, §1 and §8*.

Application	Cold-side load (≈)	ΔT direction	TEC module direction	Heat-sink direction	Sensor / mounting risk
Miniaturized photonic sensor (Raman, gas analyzer, biosensor)	≤ 1 W	Small (≤ 20 °C)	Compact single-stage	Passive or small forced-air	Sensor proximity to detector; thin sensor leads
Laser-diode wavelength stabilization (DFB, VCSEL, F-P)	0.5–5 W	Moderate (15–35 °C)	Single-stage matched to package	Forced-air, controlled airflow	Sensor on cold plate close to diode; step-hole geometry
Mid-power precision (OCXO, photodetector, ADC reference)	1–10 W	Moderate (15–40 °C)	Single-stage; oversize for COP if precision-critical	Forced-air with environment margin	Sensor mounting; airflow stability for sub-10 mK
PCR / instrumentation thermal block (illustrative, not regulated)	10–30 W (ramp)	Wide (40–90 °C; bi-direction)	Single-stage; cycle-life-rated	Forced-air with thermal-cycle margin	Mounting under thermal cycling; condensation control
Outdoor kiosk /	20–60 W	Wide; window-	Single-stage	Forced-air with	Cold-zone

Application	Cold-side load (≈)	ΔT direction	TEC module direction	Heat-sink direction	Sensor / mounting risk
enclosure cooling		mode setpoint	enclosure-class	dust / hot-ambient margin	sealing; dew-point management

System-design inputs only. Matching ATI controller families, precision grades, and evaluation boards are in companion paper AWP-TECC-02.

Component-level recommendations only. ATI TEC controllers, TEC modules, and thermistors are commercial-grade components; system-level qualification for medical / IVD (FDA, IEC 60601, IVDR), automotive / LIDAR (AEC-Q, ISO 26262), aerospace, and other regulated end products remains the OEM's responsibility.

What are the common TEC cooling system failure modes that occur after correct component selection?

In OEM TEC cooling systems where the catalog components are reasonable for the application, the common failure modes that occur after correct component selection are not component-choice failures — they are system-design and assembly failures. The recurring ones are: undersized heat sinks (using the TEC absolute-maximum hot-side rating in place of the temperature that still allows pumping Q_load at the required ΔT); insufficient controller voltage headroom at deep ΔT and worst-case hot ambient (V_TEC rises with T_hot, and V_PS − V_dropout must still cover it); thermistor placed on the TEC ceramic instead of the cold plate, or far enough from the controlled object that a thermal gradient appears between sensor and load; excessive TIM thickness on the hot- and cold-side interfaces (more grease is not better — fill the voids and squeeze the rest out); uneven mechanical compression that bows a thin cold plate and creates a center air gap; condensation when the cold-side setpoint drops below the local dew point without sealing or dry-gas purge; and inadequate worst-case validation — passing at design ambient and failing at the upper ambient bound the product will actually see in service. §4 through §5d cover each of these in the design-time sequence that prevents them.

§2. How a TEC Controller Works — Essentials

A closed-loop TEC controller samples the temperature of the load, compares the reading to a programmable setpoint, and adjusts the current through a Peltier element until the error converges. The loop runs continuously inside the controller. The amount of host involvement depends on the family: ATI precision-analog families regulate without host firmware once analog setup is complete; the high-current series requires digital configuration of PID, Auto-PID, limits, and protection thresholds, after which regulation runs autonomously; the ATFC106D uses a digital window-mode setup with its own configuration model. In all cases, after initial setup the closed loop runs continuously inside the controller without per-cycle host involvement. The amount of configuration required differs by family. The precision-analog families (TEC14M, TECA1, TEC5V4A, TEC5V6A) are largely standalone analog controllers — the setpoint, compensation network, and current/voltage limits are set with external components or pots, after which the loop runs without host involvement. The high-current series (TEC18V15A, TEC24V10A, TEC24V15A) is firmware-managed and exposes digital configuration of PID and Auto-PID parameters, current and voltage limits, and protection thresholds — once configured, the loop also runs without host involvement, but the initial setup and any field updates use the digital interface. The ATFC106D is a digital window-mode enclosure-cooling controller with its own setup model. Across all families, host-side monitoring is recommended in production designs but is not required to maintain regulation. Five functional blocks are common to every closed-loop TEC controller:

- **Temperature sensor** (NTC thermistor, platinum RTD, or IC sensor) bonded to the load.
- **Error amplifier** that subtracts sensor voltage from setpoint voltage.
- **Compensation network** (PID, with R-C components setting proportional gain, integral action, and derivative damping). Covered in §5b.
- **Power output stage** that converts the compensated signal to regulated bidirectional TEC current. Topology choices are covered in §3.
- **Protection and monitoring** including current limit, voltage clamp, thermal shutdown, and diagnostic outputs.

What is the difference between a TEC controller and a TEC driver?

The closed-loop architecture distinguishes a TEC controller from a TEC driver. A driver delivers commanded current without temperature feedback, leaving the host system to close the loop in firmware. A controller closes the loop in its own analog or digital domain. For OEM applications where temperature stability is the deliverable, the controller is generally the appropriate choice because it integrates sensor feedback, compensation, protection, and TEC power-stage drive, reducing host firmware and loop-design burden.

Heating and cooling direction control

A Peltier element pumps heat in the direction of current flow. Reversing the current reverses the direction of heat pumping — the same face that was absorbing heat now rejects it. Bidirectional TEC controllers exploit this through an H-bridge output stage, in which four switching elements arrange to connect either polarity of the supply rail across the TEC terminals. Unidirectional controllers — useful in cooling-only or heating-only applications — use a simpler single-polarity output stage; this guide focuses on the bidirectional case that is common in precision OEM products that must heat and cool around a setpoint.

When the loop calls for cooling (load above setpoint), the H-bridge routes current in the polarity that makes the load-side TEC face absorb heat. When the loop calls for heating (load below setpoint), the H-bridge reverses polarity and the same face now rejects heat. The transition is automatic, continuous, and managed entirely inside the controller — no external mode switch, no host firmware command, no relay required for TEC polarity reversal inside the control loop (the final product may still require fuses, disconnects, interlocks, or safety shutdowns dictated by its end-application and applicable safety standards).

Bidirectional drive matters when the application alternates between modes during normal operation: a laser-diode wavelength-stabilization loop that must cool in summer and heat in winter to hold the same wavelength; a PCR thermal cycler slewing between 4 °C and 95 °C; a refrigerated kiosk that cools by day and gently heats overnight to prevent condensation. Unidirectional controllers exist for applications that only ever cool or only ever heat, and they cost less, but bidirectional is the common default for precision OEM products that must both heat and cool around a setpoint.

How does a TEC controller switch between cooling and heating?

In ATI's bidirectional continuous-current TEC controllers — the precision-analog families (TEC14M, TECA1, TEC5V4A, TEC5V6A) and the high-current series (TEC18V15A, TEC24V10A, TEC24V15A); controller output-stage topology is covered in AWP-TECC-02. Output-stage topology and family-specific implementation are covered in AWP-TECC-02. The transition is continuous and managed entirely inside the controller — no external mode switch, no host firmware command, no relay is required. (Note: the ATFC106D operates differently — it is a

digital window-mode controller for enclosure-cooling applications that switches the TEC on and off within a temperature window rather than running closed-loop bidirectional drive; see §7 for its application envelope.)

Key specifications that drive selection

Four parameters dominate datasheets and selection decisions. The remainder of this guide returns to each in detail.

- **Maximum output current (I_max).** Highest TEC current the controller can deliver continuously. Must exceed the TEC's operating-point current — §4.
- **Maximum output voltage (V_max).** Peak voltage across the TEC terminals. Sets achievable ΔT — §4.
- **Temperature stability.** Peak-to-peak variation of the regulated temperature at steady state over a stated window. "Regulated temperature" means the temperature the controller actively regulates — i.e., the temperature at the sensor connected to the controller's feedback input. The temperature at the load (e.g., laser-diode junction) will differ from the sensor temperature due to thermal gradients across the package, sensor placement, and mounting variability. ATI datasheet stability figures are measured at the sensor under defined test conditions. Distinct from accuracy and resolution — §5.
- **Efficiency.** Power delivered to the TEC divided by power drawn from the supply. Topology-dependent — §3.

Part II — Controller Topology (brief)

§3. Output-Stage Topology — Summary

The TEC controller's output stage delivers regulated bidirectional current to the TEC module and comes in three principal topologies: **linear** (low ripple and low EMI but high self-heating, suited to small loads and ultra-low-noise instruments), **PWM / switch-mode** (high efficiency but switching ripple that must be filtered, suited to higher-power applications where EMI is managed at the system level), and **hybrid** (a switching stage delivers the bulk current, a linear stage post-regulates the residual — combining high efficiency with low output ripple in a single module). ATI's continuous-current TEC controller families covered in companion paper *AWP-TECC-02* use the hybrid topology, which is the architectural basis for both the precision-analog families (TEC14M, TECA1, TEC5V4A, TEC5V6A) and the high-current series (TEC18V15A, TEC24V10A, TEC24V15A). The ATFC106D is a separate window-mode enclosure-cooling controller, not a continuous-current precision controller.

Topology comparison (linear vs PWM vs hybrid), quantified hybrid-topology benefits, and topology-selection guidance by application class are detailed in AWP-TECC-02, §3.

Part III — System Design Methodology

§4. Six-Step TEC Cooling System Design Methodology

Designing a TEC cooling system is a six-step calculation. Each step constrains the next: the cold-side load sizes the TEC module; the TEC operating point sizes the controller and the heat sink; the design priority chooses between compactness and efficiency; and the heat sink anchors the worst-case ambient envelope. By the end the engineer has a specific TEC, heat sink, sensor placement, mounting stack, and a controller family to pair from companion paper *AWP-TECC-02: TEC Controller: What It Is and How to Choose One*. This guide is the system-side

companion to TECC-02; Steps 4 and 4b below summarize the controller match and supply-rail choice, with full detail in TECC-02. Steps 1, 2, 3, 5, and 6 are detailed in this paper.

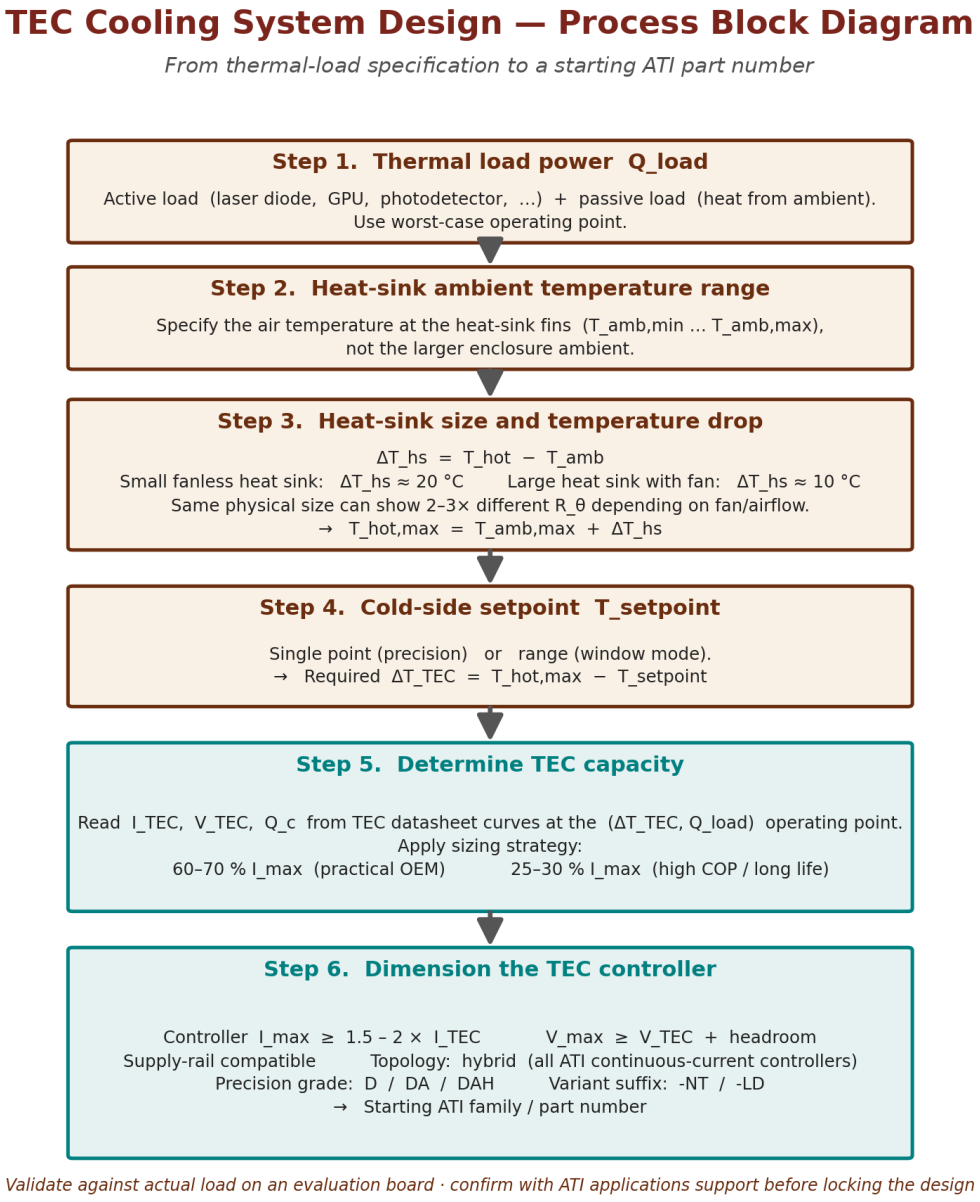

Figure 5 — TEC Cooling System Design Flow: process block diagram from thermal-load specification through heat-sink sizing, TEC capacity, and controller dimensioning to a starting ATI part number.

Step 1 — Determine the thermal load on the cold side

The thermal load Q_load is the heat power the TEC must move from the cold side to the hot side to hold the load at setpoint. The total Q_load is the sum of contributions from two physically distinct mechanisms — internal heat generation by the load object itself (active load) and heat flowing into the load object from the surroundings (passive load) — that must be added together.

Active loads — heat generated inside the load object. Two categories cover the great majority of OEM applications:

- **Heater-type sources (laser diodes, GPUs, CPUs, ASICs, FPGAs, power semiconductors).** These dissipate electrical input power as heat. For digital electronics (GPU, CPU, ASIC, FPGA), substantially all electrical input becomes heat, so Q_load ≈ P_electrical at the worst-case operating point. For laser diodes, only part of the electrical input becomes optical output and leaves the package; the rest becomes heat. The thermal load is therefore Q_load = P_electrical − P_optical_output. A 1 W optical-output laser diode with 25% wall-plug efficiency draws 4 W electrical and dissipates 3 W as heat into the package; the TEC must remove that 3 W (plus any passive contribution). Use the worst-case operating current and worst-case wall-plug efficiency from the laser-diode datasheet, not the typical values, for sizing.
- **Optical absorbers / receivers (photodetectors, sensors exposed to laser or solar illumination, beam dumps, calorimeters).** These absorb incident optical or solar power, which becomes heat in the device. For a photodetector exposed to a laser beam, Q_load ≈ P_optical_incident × (1 − reflectivity) plus the small electrical dissipation from photocurrent flowing through the detector's series resistance and bias network. For a surface exposed to direct sunlight, peak solar flux is approximately 1000 W/m² × surface area × absorptivity at the operating wavelength — for a 25 mm × 25 mm absorbing surface, that's about 0.6 W at full sun, before any internal generation. Light-receiver applications often require both worst-case (full beam) and minimum (dark) operating-point calculations to bound the controller's required current range in both directions.

Passive loads — heat flowing into the load object from the surroundings. If the load object is held below the ambient temperature — the typical TEC cooling case — heat flows continuously from the warm ambient into the cold load object via convection (still or forced air across exposed surfaces), conduction (through wires, sensor leads, mounting hardware, and any structural connection to the warm housing), and radiation (significant when ΔT is large or the cold object's emitting surface area is large). The passive load can be calculated when each thermal path is known — Q_passive ≈ (T_ambient,max − T_setpoint) ÷ R_thermal_path summed over all parallel paths — or measured directly with a thermal-isolation experiment in which the TEC is run open-loop and the steady-state cold-side temperature is recorded as a function of TEC current — within strict safety boundaries: enforce a current limit at or below the TEC datasheet maximum, monitor hot-side temperature against the TEC operating-point limit, monitor cold-side temperature against the dew point and overcooling, stop the experiment on sensor fault or any out-of-range reading, and never exceed the TEC current or temperature ratings during the open-loop sweep. Even a load object with no internal generation has a passive Q_load when held below ambient; engineers designing photodetector cooling sometimes overlook this and undersize the TEC.

Adding active and passive contributions: the total cold-side thermal load is Q_load = Q_active + Q_passive at the worst-case operating point (worst-case ambient, worst-case device dissipation, all sources active simultaneously if relevant). Typical values from real applications: fiber-coupled telecom laser diode 0.5–2 W; free-space DFB at higher power 1–5 W; single PCR well during ramp 10–30 W; OCXO atomic-clock oven 1–3 W; cooled CMOS image sensor at deep ΔT 3–15 W; refrigerated kiosk display compartment 20–60 W. Carry this total Q_load value into the TEC-sizing step (Step 3) and the controller-sizing step (Step 4). The 1.5–2× controller current headroom in Step 4 applies to the TEC operating-point current I_TEC read from the TEC datasheet, not to Q_load directly.

Step 2 — Define temperature requirements

Three temperatures matter, not one.

- **Target object temperature (setpoint).** Fixed value (a laser at 25.000 °C) or programmable range (a thermal cycler running 4 °C to 95 °C). The range determines whether a fixed-range variant suffices or whether the engineer needs the wider configurability of a -NT variant — §5b.
- **Worst-case ambient temperature at the heat sink.** Not room temperature — the worst-case temperature the heat sink will actually see in the field. Account for sunlit enclosure walls, dust accumulation, adjacent equipment, and seasonal extremes. A thermal-management design that closes at 25 °C lab ambient may fail at 50 °C in a hot environment.
- **Required ΔT across the TEC.** Defined as T_hot_side − T_cold_side, where T_cold_side is the object temperature and T_hot_side is the heat-sink-side TEC face at worst-case ambient. Larger ΔT collapses coefficient of performance and demands more controller voltage headroom.

Step 3 — Select the TEC module

The TEC module sits between the load and the heat sink. Its job is to pump Q_load (the cold-side heat load) against the temperature gradient ΔT_TEC = T_hot − T_cold,setpoint, where T_cold,setpoint is the required cold-side temperature and T_hot is the resulting hot-side temperature determined by Q_hot and the heat-sink path. The TEC's required cooling capacity at the operating point, Q_c, must equal or exceed Q_load (Q_c ≥ Q_load). The TEC datasheet's published Q_max is the heat-pumping capacity at ΔT = 0; at non-zero ΔT, the available cooling falls — to a first-order approximation, linearly with ΔT/ΔT_max. The actual relationship is non-linear once Joule self-heating, hot-side temperature dependence, and Fourier back-conduction through the pellets are considered; for precision sizing, work from the TEC datasheet performance curves at the operating point rather than from the linear approximation.

> **Common TEC sizing mistake: combining Q_max and ΔT_max**
>
> *Q_max (the maximum cooling capacity, published at ΔT = 0) and ΔT_max (the maximum achievable temperature difference, published at Q_c = 0) are end-points of the TEC's performance curve and cannot both be achieved simultaneously. A TEC pumping at Q_max delivers ΔT = 0; a TEC operating at ΔT_max pumps Q_c = 0. Selecting a module by reading these two end-points off the datasheet and assuming the TEC can do both at once is one of the most common TEC sizing errors. The correct approach is to read Q_c at the actual operating point (I_TEC, ΔT_TEC, T_hot) from the datasheet performance curves and verify Q_c ≥ Q_load there. The TEC's absolute maximum hot-side temperature is a damage / survival limit, not a design operating point — long-term reliability is significantly better below the absolute maximum.*

Two practical TEC-sizing strategies, both valid for different design priorities:

Strategy	TEC Operating Point	Use When
Practical sizing	60–70% of TEC I_max at worst-case ΔT	General OEM cooling. Balances cooling capacity, efficiency, and controller headroom. The standard recommendation for most applications.
High-COP / long-life sizing	25–30% of TEC I_max at worst-case ΔT	Battery-powered instruments, low-power field deployments, or applications targeting extended TEC service life. Requires an oversized TEC and larger heat sink, but minimizes electrical input and Joule self-

Strategy	TEC Operating Point	Use When
		heating.

From the TEC datasheet at the chosen operating point, read off the operating current I_tec and operating voltage V_tec. These are the controller's input requirements.

Companion paper What Is a Thermoelectric Cooler? covers TEC module selection in depth. ATI's ATE-series TEC modules — including the ATE1-127 single-stage and ATE2/ATE3 multi-stage families — provide a known starting point for pairing with ATI TEC controllers; final pairing must be verified by I_TEC, V_TEC, Q_c, heat-sink design, and the controller datasheet limits. See analogtechnologies.com/tec-module.html.

Step 4 — Match controller to TEC (brief)

Match a TEC controller whose I_max provides 1.5–2× headroom over the TEC operating-point current I_tec from Step 3, and whose V_max plus supply rail cover V_tec + V_dropout at worst-case hot ambient. Current headroom and voltage headroom are separate budgets — the 1.5–2× factor applies to current, not voltage. Full controller-selection methodology, output-stage topology, precision grades, and ATI TEC controller family comparison are in companion paper AWP-TECC-02, §4 Steps 4 and 4b.

Step 4b — Select input (supply) voltage (brief)

Select the lowest available supply rail that satisfies **V_PS ≥ V_TEC_operating + V_dropout(hot)** at worst-case ambient, where V_dropout(hot) accounts for the controller's output-stage dropout (which rises with junction temperature in the linear leg) plus harness drop and supply tolerance. ATI families pair to common rails as follows: 3.3 V → TECA1 (-3V-3V); 5 V → TEC14M, TECA1 (-5V-5V), TEC5V4A, TEC5V6A; 6–18 V → TEC18V15A; 5.5–24 V → TEC24V15A; 24 V → TEC24V10A; 12 V (enclosure cooling) → ATFC106D.

Step 4 + 4b complete details — current/voltage headroom rationale, V_dropout calculation from R_dson with hot-junction derating, the full controller-family ⟷ supply-rail compatibility table with output voltage and use-when guidance, and the AEO answer block on input-voltage choice — are in companion paper AWP-TECC-02, §4 Steps 4 and 4b.

Step 5 — Choose design priority

Every TEC system involves trade-offs. Decide which constraint dominates before picking a family.

Priority	What It Looks Like	System trade-off
Minimize input power	Battery-powered instrument, energy-harvesting node, long-duty-cycle field deployment. Use the high-COP sizing strategy (25–30% I_max). Accept a slightly larger TEC and heat sink in exchange for lower electrical draw and reduced Joule self-heating.	Oversize the TEC, operate at a lower I/I_max fraction, accept a larger heat sink. Controller-family routing is in AWP-TECC-02.
Minimize size / cost	Wearable, drone, dense OEM PCB, high-volume consumer or medical device. Use the practical sizing strategy (60–70% I_max).	Smaller TEC, higher current fraction, more Q_hot stress on the heat sink. Controller-family routing is in AWP-TECC-02.
Maximize ΔT (deep cooling)	Cooled CCD/CMOS far below ambient, mid-IR detector, low-noise photodetector. Need both current and voltage headroom.	Higher voltage demand, lower COP, larger Q_hot and heat sink. Controller voltage / current

Priority	What It Looks Like	System trade-off
		headroom selection is in AWP-TECC-02.
Autonomous enclosure cooling	Kiosk, refrigerated display, outdoor electronics enclosure, lab incubator. Window-mode setpoint, integrated fan control.	Hysteretic duty-cycle window-mode operation with integrated fan control. Controller-family routing is in AWP-TECC-02.

Step 6 — Size the heat sink

The heat sink is sized to dissipate the hot-side heat Q_hot, where Q_hot is computed from the TEC operating point:

Q_hot = Q_c + P_electrical *(exact, from datasheet operating point)*

Q_hot ≈ Q_load + P_electrical *(first-pass approximation when Q_c is selected close to Q_load)*

Use Q_hot = Q_c + P_electrical as the final equation for heat-sink sizing, where Q_c is the heat the TEC actually pumps at the selected operating point — read from the TEC datasheet performance curves at the actual operating point (I_TEC, ΔT_TEC, T_hot) — and P_electrical = V_TEC × I_TEC at the same operating point. This expression captures the heat pumped plus the electrical input that becomes heat at the hot side, and it already incorporates Fourier back-conduction through the pellets within the measured Q_c. The simpler **Q_hot ≈ Q_load + P_electrical** is a first-pass approximation, valid only when Q_c is selected close to Q_load (typical first-cut sizing). Final heat-sink sizing must use the datasheet-based Q_c calculation, particularly for deep-ΔT, vacuum, or multi-stage applications where the linear approximation under- or over-states the actual hot-side heat by tens of percent.

Terminology reminder. Use Q_load for the cold-side thermal load that must be removed from the controlled object, and Q_c for the heat the TEC actually pumps at the selected operating point as read from the datasheet performance curves. At steady state with a correctly sized TEC, Q_c equals Q_load, but they are not interchangeable concepts — in particular, Q_max (the datasheet end-point at ΔT = 0) is not Q_c at the operating point, and substituting Q_max into the Q_hot equation is the sizing error called out in the Q_max / ΔT_max trap callout above.

Thermal interface resistance. The R_θSA ≤ (T_hot,max − T_ambient,max) / Q_hot relation gives the required heat-sink thermal resistance to ambient. It does not include the thermal interface material (TIM) layer between the TEC hot face and the heat sink, which adds R_TIM to the total path. For typical thermal greases R_TIM is small relative to R_θSA; for thicker pads or imperfect mating surfaces R_TIM can be a significant additional drop. The full path is R_total = R_TIM + R_θSA; subtract the expected R_TIM from your computed R_θSA allowance, or specify a lower-R_TIM interface, to avoid undersizing.

Required heat sink thermal resistance:

R_θSA ≤ (T_hot,max − T_ambient,max) / Q_hot

Specify a heat sink with R_θSA comfortably below the calculated maximum. T_hot,max in the equation above is the maximum hot-side temperature that still allows the TEC to pump Q_load at the required cold-side setpoint and ΔT — read from the TEC datasheet performance curves at your actual operating point — and is generally lower than the TEC's absolute maximum hot-side rating. Engineers sometimes substitute the absolute

maximum, which yields a heat sink that is too small to support the required cooling capacity. The appropriate margin depends on operating environment. In the bullets below, an "Nx margin" means selecting a heat sink whose thermal resistance is approximately 1/N of the calculated R_θSA maximum — i.e., lower °C/W by a factor of N — so the temperature rise from heat-sink loading, dust, fan aging, and ambient excursions still keeps T_hot below its limit:

- **Controlled lab / production environment:** 1.5× to 2× margin (R_θSA,selected ≈ R_θSA,max ÷ 1.5 to ÷ 2) is usually sufficient.
- **Industrial / general OEM:** 2× to 3× margin to absorb dust accumulation and fan aging.
- **Outdoor, dusty, or harsh environments:** 3× to 5× margin to absorb fan-flow degradation, filter loading, and ambient excursions above design.

Important: the margin recommendations above are based on the operating environment of the end product, not the environment where the engineering design is performed. A lab-developed instrument that ships to customers and operates in uncontrolled or field conditions should be treated as "industrial" (2×–3×) or "outdoor" (3×–5×) for heat-sink margin, not as "controlled lab." The margin should reflect the conditions the product will encounter in service.

Passive sinks are often practical for low-power designs with sufficient volume and allowable temperature rise; forced air becomes common as Q_hot rises or space is constrained; liquid cold plates become attractive above 50 W when compactness, acoustic limits, or ambient constraints make forced air insufficient. Detailed heat-sink design — fin geometry, fan selection, thermal interface materials, R_θSA measurement — is covered in the companion paper *AWP-HS-01: How to Design a Heat Sink for a TEC Controller System*.

The complete selection process is summarized in Figure 5 at the start of Part III (the "§4 Six-Step TEC Cooling System Design Methodology"), which shows the full flow from thermal load through heat-sink sizing, TEC capacity, and the six-step controller dimensioning above to a starting ATI part number.

Key Insight — the six-step methodology in one paragraph

Calculate the cold-side thermal load (Q_load), then read the TEC's operating-point current and voltage at worst-case ΔT from its datasheet (not its I_max/V_max). Match the controller's I_max to that operating current with a 1.5×–2× safety margin and its V_max with the appropriate voltage headroom (either the 10–20% heuristic or the explicit Step 4b calculation — not both). Confirm the controller's supply rail matches the available system rail. Choose the precision-grade variant (D, DA, DAH) that matches the required setpoint precision and the suffix (-NT, -LD) that matches the load type. Size the heat sink for Q_hot at worst-case ambient. Each step is detailed above; final selection is confirmed with bench evaluation against the actual load.

Send these inputs to ATI applications engineering for a thermal-stack review

Engineers who want ATI to review their calculated starting point can send the inputs below to sales@analogtechnologies.com; ATI applications engineering can review qualified OEM projects and suggest a starting TEC module, heat-sink direction, sensor placement approach, and matched controller-family pointer to AWP-TECC-02. The full selection-support form is in §10.

1. Cold-side thermal load Q_load (W) and target setpoint or range
2. Worst-case ambient envelope at the heat sink and required T_hot,max
3. Heat-sink configuration (passive / forced-air / liquid) and any size or acoustic constraints
4. Sensor placement plan (distance to load, mounting method) and stability target at the controlled object
5. Condensation risk (cold-side temperature vs. expected dew point) and any sealing requirement

6. Validation target (worst-case test envelope, duration, instrument) and application class

§5. System-Level Stability (brief on the controller side)

The controller defines an upper bound on regulated stability *at the sensor* under defined laboratory test conditions — precision grade D (≤5 mV setpoint accuracy, ≈±0.05 °C), DA (≤2 mV, ≈±0.01 °C), or DAH (≤0.5 mV, ≈±0.001 °C achievable at-sensor). In a production TEC cooling system, three different stabilities must be distinguished and they degrade in this order: (1) controller regulation stability at the sensor (what the closed loop holds the sensor reading to — the figure quoted in controller datasheets); (2) cold-plate stability near the sensor (the actual metal temperature in the immediate vicinity of the sensor bead, which differs from (1) by the sensor mounting thermal resistance and any sensor self-heating); (3) device or load stability (the temperature of the actual controlled object — laser-diode junction, photodetector chip, biosensor surface — which differs from (2) by the thermal path between cold plate and device including TIM, lead-frame, and package thermal resistances). The controller controls (1); the system design controls (2) and (3).

> **Important: at-sensor vs. at-load stability**
>
> *The millikelvin stability figures cited in controller datasheets are measured at the sensor under standard test conditions. In a production system, thermal gradients between sensor and load, sensor mounting variability, ambient drift, ADC noise, and PID tuning quality on the actual thermal mass will degrade the stability at the load relative to the stability at the sensor. The controller alone does not guarantee 1 mK at the load — validate stability on the actual assembly, at the location that matters for the application, before locking the design.*

Controller-side stability vs. accuracy vs. resolution interpretation, the controller datasheet test conditions (sensor type, mounting, ambient, bandwidth, duration), and the controller specification framework are in companion paper AWP-TECC-02, §5.

Design considerations for sub-10 mK (<0.01 °C)-class stability targets

For applications targeting regulated temperature stability tighter than approximately 10 mK at the load, the controller's contribution is only a fraction of the total stability budget. The following design rules address the other contributors — controller body-temperature stability, heat-sink ventilation, ambient air, sensor mounting, and supply / setpoint noise — that typically dominate at this level.

Controller placement and topology — keep the controller body temperature stable

Two design choices about the controller itself become important at sub-10 mK, before the heat sink, sensor, and supply rules below.

- **Place the TEC controller in a well-ventilated location.** The controller dissipates some power as heat — primarily in the hybrid linear leg at significant output current, plus modest losses in the switching leg and the control electronics. If the controller sits in a stagnant or thermally-coupled space (enclosed compartment, adjacent to other heat sources, mounted on a hot surface), its body temperature drifts with load, ambient, and time. The voltage references, sensing circuits, comparators, and — in firmware-managed families — ADCs and digital-control circuits inside the controller all have small but non-zero temperature coefficients; controller body-temperature drift therefore translates to setpoint drift that the closed loop cannot correct — because the setpoint moved, not the load. For sub-10 mK applications, mount the controller in a position with free airflow to ambient, away from other heat sources, and not

inside a sealed compartment. A small dedicated airflow over the controller package, or a chassis-mount thermal interface to a large stable thermal mass, removes body-temperature drift as a stability contributor.

- **Prefer hybrid or well-filtered switching topologies for sub-10 mK applications.** A pure-linear TEC controller dissipates V_drop × I_out continuously in its linear pass element across the full output range. For sub-10 mK applications this introduces two practical problems: (a) the controller body temperature follows load current and ambient temperature directly, producing the setpoint drift described above; and (b) the total input power is much higher than an equivalent hybrid or well-filtered switching design, often with no system-level precision benefit once sensor placement, heat-sink stability, airflow, and setpoint noise dominate the result at the load. Carefully engineered pure-linear designs can achieve very low ripple and may suit specific niche cases (very small loads, very tight EMI budgets where switching is prohibited entirely), but for general-purpose sub-10 mK OEM designs, hybrid or well-filtered switching topologies are usually preferable. ATI's continuous-current TEC controller families covered in this guide use the hybrid topology — the switching leg delivers the bulk current, the linear leg only post-regulates the residual — which keeps controller self-heating modest and roughly constant across the operating range.

Heat sink: constant ventilation, no variable airflow

Convection on the hot-side heat sink is the largest single source of slow ambient-driven instability in most TEC systems. Two engineering choices matter at sub-10 mK:

- **Constant fan speed (or slow scheduled control).** A heat-sink fan running at a fixed speed produces a constant heat-sink temperature offset above ambient (ΔT_hs). Any variation in fan speed — including temperature-feedback fan control that ramps with measured load temperature — appears as a slow ΔT_hs variation that the TEC controller must compensate for; the compensation itself introduces a small setpoint disturbance. For sub-10 mK applications, run the heat-sink fan at a single constant speed and size the heat sink so that this fixed-speed operating point handles worst-case Q_hot. If variable fan speed is unavoidable for system-level acoustic or power constraints, drive it as a slow open-loop schedule with a time constant much longer (>10×) than the thermal time constant of the load — or with feed-forward compensation derived from the known heat-load profile — rather than as a closed loop tied directly to the regulated load temperature. Verify the effect of any fan-speed change on regulated temperature using the actual assembly before locking the design.
- **No ambient wind on the load or the heat sink.** The distinction is between intentional, constant forced convection on the heat sink (which is controlled and stable) and unintentional, variable external air currents (HVAC, doors, traffic, fume-hood draft) that impinge on the load or the heat sink. The heat-sink fan provides the former; the enclosure excludes the latter. Even a small variable airflow across the load object or across the heat-sink fins modulates the system thermally and shows up directly in the regulated temperature. For sub-10 mK applications, the entire load + sensor + heat-sink assembly should sit inside an enclosure that excludes external air currents; the intentional internal fan provides the necessary forced convection on the heat sink, and everything else should be still.

Temperature sensor: small, thermally well-coupled, electrically clean

The sensor sets the upper limit on what the closed loop can resolve and how quickly it can react. For sub-10 mK applications:

- **Use a small NTC thermistor.** Choose the smallest bead diameter consistent with the required self-heating budget. Smaller thermal mass means faster response, and faster response means the loop can correct disturbances before they propagate to the load.

- **Use thin sensor leads.** Thick leads conduct parasitic heat from the warmer ambient end down to the sensor bead and bias the reading; thin leads minimize this heat-leak.
- **Mount the sensor as close to the target object as the geometry allows.** Every millimetre of physical separation between the sensor bead and the temperature you actually care about adds thermal-resistance error and slows the measured response to load changes.
- **Use a step (counter-bored) hole when potting, not a straight through-hole.** See Figure 6 below. The step geometry seats the sensor bead at the bottom of a narrow well that is in close thermal contact with the target metal; the leads route through the wider upper channel, which can be filled with potting compound for mechanical retention. A straight through-hole leaves the bead suspended in potting compound mid-hole, adding significant sensor-to-target thermal resistance and slowing the measured response. The ATI thermistor datasheet at www.analogtechnologies.com shows the dimensioned drawing of the recommended step-hole geometry.

Figure 6 — Thermistor potting: step hole vs straight hole

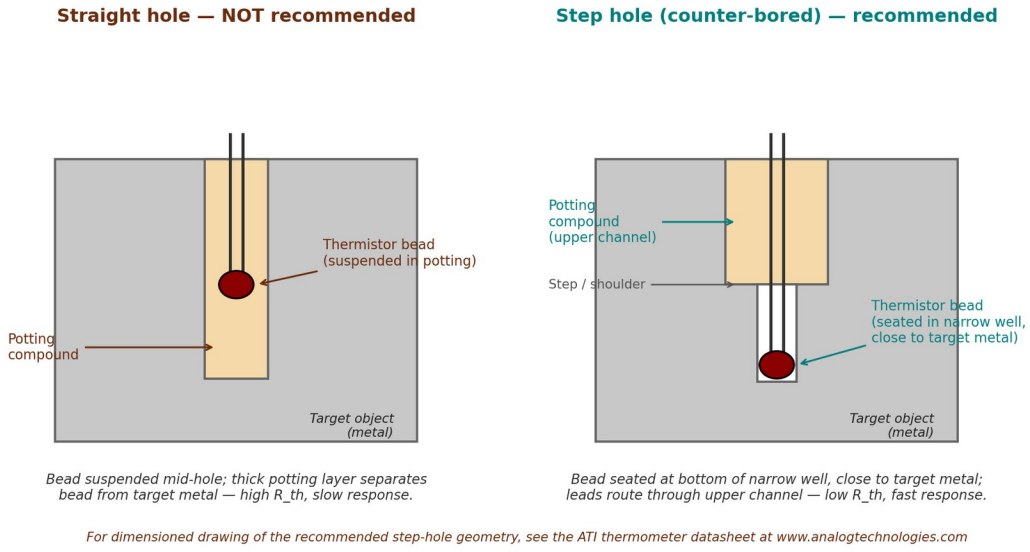

Figure 6 — Thermistor mounting and potting: step (counter-bored) hole vs straight through-hole. Seat the sensor bead at the bottom of the narrow well, close to the target metal, for low thermal resistance and fast response. Refer to the ATI thermistor datasheet at www.analogtechnologies.com for the dimensioned drawing.

Power supply and setpoint pin: no low-frequency noise below 1 kHz

The closed loop attenuates output disturbances roughly in proportion to loop gain at the disturbance frequency. Above the loop bandwidth, the loop cannot react. Below the loop bandwidth, low-frequency noise on either the TEC supply rail or the analog setpoint (TMS) control pin appears directly at the regulated temperature.

- **TEC power supply (V_PS) — manage low-frequency drift and audio-band noise.** V_PS noise becomes important when it (a) reduces output-voltage headroom enough to push the controller into saturation, (b) couples into the controller's sensing or reference circuitry, (c) modulates the output stage in a way that appears at the TEC, or (d) appears at frequencies within the controller's loop bandwidth in a manner the loop cannot reject — particularly under saturation or near-saturation. Line-frequency hum (50/60 Hz and harmonics), audio-band switching noise, and slow load-dependent rail drift are the typical offenders. Most modern switching supplies put their dominant noise above 50 kHz, well outside the temperature-loop

bandwidth, and is easily filtered with standard LC components — switching noise in this range is generally acceptable. Verify with an oscilloscope and low-frequency spectrum analysis on the V_PS rail under representative load.

- **TMS (setpoint) control pin — directly converts noise into temperature command error.** Unlike supply-rail noise, TMS noise is interpreted by the loop as a real setpoint change, so any noise on the TMS pin appears directly at the regulated temperature scaled by the controller's V/°C constant from the family datasheet. There is no rejection of TMS noise by the loop — the loop is following it. For sub-10 mK applications, route the TMS signal away from switching nodes and digital lines, use shielded or twisted-pair wiring if the signal travels off-board, and bypass the pin with a small ceramic capacitor close to the controller. Derive an allowable noise budget from the controller's V/°C scaling and the required temperature stability target, and confirm that measured noise on TMS is below that calculated threshold at all frequencies — TMS noise at any frequency that produces a temperature command above the stability target will appear at the regulated load.

Summary: at sub-10 mK, the system is no longer dominated by what the controller can do in isolation. The dominant contributors become controller body-temperature stability (placement and topology), ambient air movement (heat-sink fan constancy, enclosure isolation), sensor mounting quality (small bead, thin leads, step hole), and low-frequency electrical noise on the V_PS rail and the TMS pin. Address each of these design choices before attributing any residual instability to the controller; the controller can only regulate as well as its own body temperature, the sensor sees, and the supply and setpoint allow.

§5b. Compensation Network and Variant Suffixes (brief)

The compensation network is the cluster of resistors and capacitors that sets the PID loop's frequency response. ATI precision controllers ship with an internal compensation network factory-tuned for typical thermal loads, with three precision grades (-D, -DA, -DAH) selectable by part number suffix. For non-standard thermal loads, the -NT variant exposes the network for external configuration. The -LD variant is pre-configured for laser-diode thermal masses. From a system-design perspective, the relevant choices are (a) precision-grade matched to the application's stability target, and (b) -NT only if the thermal load is unusually large, small, or fast-slewing relative to typical OEM cooling applications.

Compensation theory, the variant-suffix decoder table (-D, -DA, -DAH, -NT, -LD with internal scope and applicability), and the AEO question/answer block on compensation networks are detailed in companion paper AWP-TECC-02, §5b.

§5c. Validation Testing and Long-Term Lifetime Margin

High-stability, high-efficiency TEC systems are not adequately validated by typical-condition testing alone. Two practical considerations drive thorough validation: (a) worst-case operating conditions can hide failures that day-to-day testing does not surface; and (b) the TEC module's own electrical and thermal characteristics drift over service life, so a system that passes specification at the start of its life can fall out of specification before the program's required service window ends. recommended for systems targeting high stability or long service life

Worst-case validation test conditions

Test the system under every combination of worst-case inputs the application can credibly encounter, not just at typical room conditions. The most stressful combination for a TEC controller is usually all of the following simultaneously:

- **Hottest possible heat-sink temperature.** Set the ambient (or environmental chamber) to the upper-end specification temperature, with the heat sink fully loaded and all internal heat sources operating at maximum.
- **Lowest setpoint temperature the application demands.** Drives the deepest required ΔT across the TEC, which sets the maximum V_TEC and the worst COP.
- **Highest thermal load the cold side will see in service.** Includes the worst-case combination of active load (laser at maximum drive current, GPU at sustained TDP, photodetector at peak illumination) and passive load (worst-case ambient conduction and convection into the cold object).
- **Lowest power supply voltage within the allowed V_PS tolerance band.** Includes line sag, brown-out conditions, low-end of the V_PS specification, and end-of-life battery voltage in battery-powered products. Low V_PS at high I_out is where headroom limitations show up first.

This combination drives the controller's output current and voltage to their maximums simultaneously. A controller that operates comfortably at room temperature and nominal supply may saturate (current-limited or voltage-limited) under this combined worst case, with the load drifting off setpoint and no closed-loop authority to correct it. If the controller saturates anywhere in this envelope, the design margin is insufficient — step up to the next-larger family, improve the heat sink, or relax the system requirements until the worst-case envelope is comfortably inside the controller's capability.

Reserve maximum output-power capacity for long-term lifetime margin

A new TEC module is at its best electrical and thermal state on day one. Over service life, the module's AC resistance (R_AC, also called ACR — measured at approximately 1 kHz) gradually rises as the bismuth-telluride pellet solder joints and contact interfaces age, particularly under thermal cycling. This is normal wear for any thermoelectric module. As R_AC rises:

- **The TEC requires more voltage to deliver the same current** → V_TEC at the operating point rises.
- **The TEC's coefficient of performance (COP) falls** → more electrical input power is required to pump the same Q_load.
- **The controller must deliver more current and voltage to hold the setpoint** → controller dissipation and Q_hot at the heat sink both rise.

A controller selected with exactly the day-one I_TEC and V_TEC operating point will fall out of specification before the TEC's service life ends. Reserve appropriate current and voltage margin based on service-life target, thermal cycling, ambient envelope, and controller capability — the 1.5× to 2× current safety margin discussed in §4 Step 4 is intended to absorb this normal end-of-life degradation in addition to ambient and cool-down transients. Conservative margin at design time keeps the system operating within specification across the full TEC service life, and avoids the alternative of finding out during a fielded-product warranty period that the controller has saturated as the TEC aged.

ATI long-life TEC modules — slower R_AC drift

The rate of R_AC rise depends strongly on the TEC's construction — pellet metallurgy, solder alloy selection, contact interface quality, and packaging. ATI offers long-life TEC modules with construction optimized for slow R_AC drift; under ATI's representative thermal-cycling conditions, these modules are designed for slower R_AC drift under thermal cycling, with internal ATI characterization data available for qualified OEM evaluations (comparator class, cycle profile, temperature extremes, and failure criterion are documented in the test report). Specific accelerated-life test data — cycle profile, temperature extremes, comparator class, and failure criterion

— is available on request from ATI engineering for customers evaluating long-life modules against application-specific reliability targets. For OEM products with multi-year service life requirements — laser-diode wavelength stabilization in deployed instruments, diagnostic and laboratory instrumentation with refurbishment cycles measured in years, telecom and industrial equipment in remote installations, and any application where field replacement is expensive or service-disruptive — ATI's long-life TECs are recommended over generic commercial modules. Contact ATI sales for the recommended long-life part number matched to your application's thermal load, ΔT, and service-life target.

§5d. TEC Mounting and Assembly — Between the Load and the Heat Sink

A correctly selected TEC, controller, sensor, and heat sink will still underperform — or fail — if the TEC is mounted incorrectly between the load object and the heat sink. Correct mounting is equally critical to achieving the designed performance: a poorly assembled stack can lose 20–40% of its theoretical cooling capacity to parasitic thermal resistance at the two TIM interfaces, and reversed orientation produces the opposite of the intended effect (heating the load instead of cooling it). This section covers the practical mounting details, sourced from ATI's TEC module application white paper AWP-TECM-01, that separate a reliable production assembly from a frustrating prototype.

How do I mount a TEC module correctly between the load and the heat sink?

Short answer: (1) Verify orientation against the module's mechanical drawing — for the standard ATI TEC modules covered by this guide, the label side is the cold side and must face the load; sealed, miniature, or custom variants may differ. (2) Apply a thin uniform layer of thermal-interface material (target 25–100 µm) on both ceramic faces — fill voids, do not build up bulk. (3) Lap both mating surfaces flat to ±0.025 mm (TIR ≤ 0.076 mm). (4) Use four screws with spring washers and apply uniform compression to a target contact pressure recommended by the specific TEC manufacturer's datasheet — a 1.0–2.0 MPa range over the ceramic-face area is a typical starting target, but verify against the specific module's datasheet because allowable pressure varies with module construction, ceramic thickness, and substrate stiffness. (5) Apply torque in small increments alternating between opposite screws, never apply shear force to the TEC, retorque after one hour, and verify pressure distribution with pressure-indicating film for thin cold plates or large-footprint modules. (6) Place the temperature sensor on the cold plate within 5 mm of the load — not on the TEC ceramic. Full procedure and rationale follow below.

Compression pressure is not screw torque

Specified TEC compression is a target contact pressure (e.g., 1.0–2.0 MPa over the TEC's ceramic-face area) — not a direct screw-torque value. Translating contact pressure to torque depends on the number and size of mounting screws, washer stack, thread lubrication, friction coefficient, TEC footprint, cold-plate thickness and stiffness, and TIM compliance. Thin or poorly-supported cold plates can bow under non-uniform compression. Over-compression can crack the TEC's ceramic faces; under-compression raises thermal-interface resistance and degrades cooling. For a specific TEC module, use the manufacturer's recommended compression pressure (or torque) and verify the load distribution and cold-plate flatness experimentally before locking the assembly procedure.

Critical orientation rule

For the standard ATI TEC modules covered by this guide, the label and printing are on the COLD side

of the module. When mounting, the cold side (label side) must face UP toward the load object; the hot side (wire side) faces DOWN toward the heat sink. Reversing this orientation drives heat the wrong way and heats the load instead of cooling it. Always verify orientation against the specific module's mechanical drawing before applying TIM or torquing screws — sealed, miniature, round, custom, or high-temperature TEC variants may have orientation conventions specific to that variant.

The mounting stack

The TEC sits between the cold-plate load above and the heat sink below, sandwiched between two thin layers of thermal interface material (TIM). Figure 7 shows the cross-section. From top to bottom the stack is:

1. **Load object / cold plate** — the device whose temperature you regulate (laser-diode mount, photodetector cold plate, sensor housing, well block, localized GPU / ASIC hotspot cold plate, etc.)
2. **TIM layer** — thin uniform thermal grease, target 50–100 µm thickness, between cold plate and the TEC's cold-side ceramic
3. **TEC module** — cold side (label) facing UP toward the load; hot side (wires) facing DOWN toward the heat sink. Bi_2Te_3 pellets sandwiched between the two ceramic plates, with the wires exiting from the hot ceramic side
4. **TIM layer** — thin uniform thermal grease between the TEC's hot-side ceramic and the heat-sink mating surface
5. **Heat sink** — sized for Q_hot at worst-case ambient (§4 Step 6), with the mating surface lapped flat to the same tolerance as the cold plate

The whole stack is held in uniform compression by four mounting screws with spring washers, two on each side of the module (clear of the ceramic). The mounting forces transmit through the cold plate and the heat-sink base, not through the brittle ceramic substrates of the TEC itself.

Figure 7 — TEC mounting stack cross-section

Label / printing on COLD side (faces UP toward load); wires exit HOT side (faces DOWN toward heat sink)

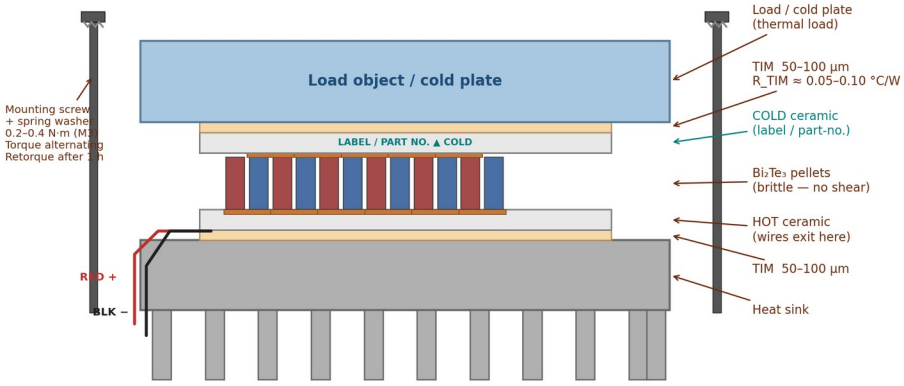

Target compression 1.0–2.0 MPa (150–300 PSI) · surfaces lapped flat to ±0.025 mm · never apply shear force · sensor on cold plate ≤ 5 mm from load (not on TEC ceramic)

Figure 7 — TEC mounting stack cross-section: cold-plate load on top, thin TIM, TEC with label / cold side facing the load, thin TIM, heat sink at bottom; mounting screws and spring washers on each side, wires exiting the hot side. Critical: label / cold side faces the load; reversing this heats the load instead of cooling it.

Surface preparation

Both the cold-plate mating surface and the heat-sink mating surface must be ground or lapped flat to within ±0.025 mm across the TEC module footprint, with a maximum total indicated reading (TIR) of 0.076 mm. Remove all burrs, machining chips, and foreign matter from the mounting area before applying TIM. For multi-module arrays sharing a common base plate, the TEC thickness variation between modules must not exceed ±0.025 mm — specify tolerance-lapped modules from ATI when ordering.

Thermal interface material (TIM) — thin and uniform

Apply a thin, continuous film of thermal grease to BOTH ceramic faces of the TEC and to the corresponding cold-plate and heat-sink mating areas. The target thickness is 25 ± 13 µm (fine-screen roller, production) or 50–100 µm (hand-applied with a flat-edge spatula). The goal is to fill microscopic surface voids — not to create a thick insulating layer. More TIM is not better. The following table from AWP-TECM-01 shows how thickness affects per-interface thermal resistance and the resulting temperature penalty at a representative 12 W heat-flux operating point. The R_TIM values are for a representative ATI test geometry — approximately 30 mm × 30 mm TEC footprint and a thermal grease with k ≈ 1 W/(m·K) (representative of standard zinc-oxide / silicone TEC mounting greases) — and scale inversely with contact area and proportionally with bond-line thickness; for a different module area or TIM material, scale by R_TIM = δ / (k × A) where δ is bond-line thickness, k is TIM thermal conductivity, and A is contact area. High-performance greases with k in the 3–8 W/(m·K) range reduce R_TIM proportionally and are worth specifying for tight thermal budgets.

TIM thickness	R_TIM per interface	Temperature penalty at 12 W
50 µm (ideal)	0.05 °C/W	0.6 °C
100 µm (acceptable)	0.10 °C/W	1.2 °C
200 µm (too much)	0.20 °C/W	2.4 °C
500 µm (excessive)	0.50 °C/W	6.0 °C

If thermal grease squeezes out the sides of the module when the assembly is clamped, too much TIM was applied. Wipe off the excess and reapply more thinly. The two TIM interfaces (cold side and hot side) together can easily contribute several degrees C of additional ΔT that the TEC must pump against — this directly subtracts from the available cooling capacity at the load.

Compression and torque

Use four screws with stainless-steel spring washers (Belleville or split-lock type) to apply uniform compression across the module footprint. Stainless-steel screws with fiber-insulating shoulder washers prevent electrical contact between the screw and any conductive trace. The 1.0–2.0 MPa (150–300 PSI) compression range cited here is a typical starting target for many TEC modules in standard mounting fixtures; the allowable compression varies with module construction, ceramic plate thickness, pellet height, and substrate stiffness, and final compression must be verified against the specific TEC manufacturer's datasheet for the module in service. For M3 screws on a typical 30 mm × 30 mm class module, this typical range corresponds to roughly 0.2–0.4 N·m of torque per screw. Always verify the recommended contact pressure (or torque) against the specific TEC manufacturer's mechanical recommendations for your module size, ceramic-plate thickness, and cold-plate / heat-sink-base stiffness before locking the assembly procedure.

Over-compression and thin-substrate caution

Excessive compression can fracture the TEC's ceramic plates, crush solder joints between pellets and pads, or — equally damaging — bow a thin cold plate or heat-sink base inward, lifting the corners off the TEC and creating a central air gap that defeats the entire purpose of the assembly. For thin cold plates, thin heat-sink bases, or large-footprint TECs, increase plate stiffness (thicker plate, ribbing, edge fastening) before increasing torque. When in doubt, start at the low end of the target range, verify the pressure distribution with pressure-indicating film, and adjust only as needed to close the contact.

The exact torque per screw for a target compression pressure is calculated from:

$T = (K \times D \times p\_target \times A) / N$

where T is the torque per screw in N·m; K is the nut factor (also called the torque coefficient), typically ≈ 0.20 for dry stainless-steel threads and ≈ 0.15 for lightly lubricated threads; D is the nominal screw diameter in metres; p_target is the target compression pressure in pascals; A is the TEC ceramic surface area in square metres; and N is the number of mounting screws. The factor (p_target × A) is the total clamping force required across the module; dividing by N gives the per-screw force, and the K × D term converts force to torque for a standard threaded fastener. Apply torque in small increments, alternating between opposite screws, using a torque-limiting screwdriver. For multi-module arrays, use three screws along the module centerline and torque the middle screw first. Verify the resulting compression in practice using pressure-indicating film (Fujifilm Prescale or equivalent) across the module footprint, especially when transitioning to production assembly or when the cold plate or heat-sink base is on the thin side; torque is only an indirect proxy for compression, and the relationship depends on screw friction, washer stack, lubrication, and stack geometry.

Critical compression cautions

Excessive torque on thin substrates causes the cold plate or heat-sink base to bow, which opens an air gap at the module center and dramatically increases thermal resistance — the opposite of the intended effect. Reduce torque if either mating surface is thinner than 3.2 mm copper or 6.4 mm aluminum. When in doubt, use pressure-indicating film (Fujifilm Prescale or equivalent) to verify uniform contact pressure across the module footprint before locking the design.

Never apply shear (lateral) force to the TEC during assembly. The Bi_2Te_3 pellets are brittle ceramic and crack under lateral load. Use alignment pins or a fixture to position the module, then apply only compressive (downward) force.

Check torque again one hour after initial assembly and retighten if necessary — thermal grease compresses and settles under initial load, and the assembly relaxes slightly in the first minutes after torquing.

Step-by-step assembly procedure

A repeatable production assembly follows the same sequence every time. The procedure below, summarized from AWP-TECM-01 §10.3, has been refined over decades of OEM field experience:

1. **Verify surface preparation** — both mating surfaces lapped flat to ±0.025 mm TIR ≤ 0.076 mm; all burrs, chips, and foreign matter removed from the mounting area

2. **Apply TIM to the hot-side interface** — thin continuous film on the heat-sink mating area and on the TEC's hot ceramic (wire side), target 25 µm (fine-screen roller) to 100 µm (hand-applied)
3. **Place the TEC on the heat sink** — hot side DOWN (wires exit toward the heat sink). Gently oscillate the module back and forth while applying uniform downward pressure until thermal compound efflux is visible around the edges and mechanical resistance is felt — this confirms intimate surface contact with minimal trapped air
4. **Apply TIM to the cold-side interface** — thin continuous film on the TEC's cold ceramic (label side, facing up) and on the matching cold-plate area
5. **Place the cold plate on the TEC** — repeat the oscillation technique to remove trapped air; keep the cold plate centered between the four mounting screws to prevent uneven compression that could crack the ceramic substrate
6. **Preload the assembly** — use a light clamp or calibrated weight aligned with the module center before final bolting
7. **Torque the screws in alternating sequence** — apply torque in small increments, alternating between opposite screws, using a torque-limiting screwdriver set to 0.2–0.4 N·m (M3) for a 1.0–2.0 MPa compression target
8. **Verify and retorque after 1 hour** — thermal grease settles and the stack relaxes; check final torque and adjust if necessary
9. **Verify orientation electrically before applying setpoint** — at low TEC current with a known cooling demand, confirm the cold side actually cools (and the hot side warms). If the system heats the load instead, the most likely cause is reversed wire polarity (red to TEC+, black to TEC−) or — far worse — an upside-down module that requires complete disassembly

Temperature sensor mounting (cross-reference)

The NTC thermistor or other temperature sensor must be mounted on the cold-plate / load object — NOT on the TEC's cold ceramic — and as close to the actual thermal load as the geometry allows (typically within 5 mm). Mounting the sensor on the TEC ceramic puts the TIM thermal resistance between the measurement and the regulated object, producing an uncontrolled offset that varies with heat load and TIM thickness. For sub-10 mK applications use a step-hole potting geometry rather than a straight through-hole (see §5 sub-10 mK design rules and Figure 6 for the cross-section and dimensioned-drawing reference).

Sealing and condensation management

When the cold-side temperature operates below the local dew point, water condenses on the cold ceramic and the cold plate — corroding the Bi_2Te_3 elements, shorting traces, contaminating optical surfaces, and eventually causing open-circuit failure. At 25 °C ambient and 60% relative humidity, the dew point is approximately 16.7 °C (standard psychrometric value); any cold-side temperature below this threshold requires condensation management. Two practical options: (a) use sealed (-S suffix) TEC variants designed for sub-dew-point operation, and (b) seal the assembly inside an enclosure with a dry-gas purge (nitrogen or dry air) or a desiccant pack. The mounting cavity around the TEC pellets should also be insulated with closed-cell polyurethane foam and perimeter-sealed with RTV silicone (or epoxy for higher mechanical integrity) to block parasitic air-gap conduction and exclude moisture. For representative ATI test geometries at high ΔT, this insulation has been observed to recover on the order of 5–15% of net cooling capacity, where parasitic conduction across the air gap between hot and cold ceramic plates otherwise becomes a significant performance limiter — the actual

recovery depends on module size, ΔT, and the geometry of the surrounding air volume, and must be confirmed in the application.

§6. Temperature Sensor Selection

The temperature sensor is the eyes of the controller. A noisy or miscalibrated sensor makes a precision controller useless. Three sensor families cover the great majority of OEM TEC applications. The numerical comparisons below are typical and sensor-dependent — actual figures vary with grade, readout circuit, and operating conditions; consult the specific sensor datasheet for committed numbers.

What is the best temperature sensor for laser-diode cooling with an ATI TEC controller?

Short answer: a precision 10 kΩ NTC thermistor with B-value matched to the ATI controller's setpoint calculation. NTCs deliver the highest resistance-change sensitivity near typical laser-diode operating points (15–35 °C), the small bead size (≤ 1 mm) places the sensing element close to the diode submount with minimal thermal gradient, and the simple voltage-divider readout is directly compatible with ATI's precision-analog DAH-grade controllers. For applications above 85 °C, where NTC sensitivity falls and self-heating becomes problematic, a thin-film Pt100 RTD is the alternative. Avoid integrated IC sensors (TMP-class) for sub-50 mK laser-wavelength stabilization — their digital quantization and self-heating set a noise floor too high for tight wavelength control. ATI applications support can confirm the recommended sensor part number for the specific diode package and target stability.

NTC thermistors

A negative-temperature-coefficient thermistor is a sintered metal-oxide bead whose resistance falls steeply with temperature — approximately 4% per °C near 25 °C for the standard 10 kΩ value (≈400 Ω/°C at 25 °C). The high resistance-change sensitivity makes NTCs the default sensor for precision TEC control near room temperature; however, the achievable system temperature resolution also depends on the ADC resolution, reference stability, and wiring. With a 16-bit ADC and stable reference (the precision-analog ATI controllers use a precision ADC and internal reference), the NTC voltage-divider front end can resolve approximately 1–5 mK near 25 °C; with a 12-bit ADC, the same circuit resolves approximately 6–10 mK. The sensitivity claim by itself does not guarantee millikelvin resolution at the system level.

Temperature dependence of NTC sensitivity. The 4%/°C and ≈400 Ω/°C figures above are for a standard 10 kΩ NTC at 25 °C. Near room-temperature TEC applications, a typical 10 kΩ NTC is on the order of 3.5–4.5%/°C depending on β value; the absolute Ω/°C falls with operating temperature because the resistance itself falls — at approximately 50 °C the resistance is around 3.6 kΩ and the absolute sensitivity is on the order of 140 Ω/°C; at 85 °C the resistance is closer to 1.5 kΩ and the absolute sensitivity is on the order of 60 Ω/°C (exact figures depend on the NTC's β and tolerance class). For applications operating significantly above room temperature, verify that the NTC's absolute sensitivity together with the ADC resolution still meets the required stability and accuracy at the actual operating point.

- **Strengths:** highest sensitivity near 25 °C, small physical size, inexpensive in volume, simple voltage-divider readout, low self-heating with proper drive-current limiting.
- **Limitations:** non-linear (handled by Steinhart-Hart or lookup table), limited range, modest long-term drift.
- **Best fit:** the majority of precision TEC applications — laser-diode cooling, photodetector regulation, OCXO ovens, imaging-sensor cooling.

Platinum RTDs

A platinum resistance temperature detector (Pt100 or Pt1000) is a precision platinum wire or thin film whose resistance increases nearly linearly with temperature.

- **Strengths:** near-linear response, lowest long-term drift of any precision sensor, widest temperature range, standardized tolerance classes (IEC 60751).
- **Limitations:** lower sensitivity than NTC near 25 °C, requires precision current source for excitation, lead-resistance compensation needed for 2-wire connections, higher per-sensor cost.
- **Best fit:** regulated medical, metrology, and reference applications; wide-temperature-range applications; long sensor cable runs where 4-wire Kelvin connection cancels lead resistance.

Semiconductor IC sensors

Silicon-based temperature sensors integrating sensing and readout on one die, producing analog voltage, analog current, or digital interface output.

- **Strengths:** simple interface, factory-calibrated, low cost, negligible self-heating.
- **Limitations:** lower absolute precision than NTC or RTD, larger thermal mass (slower response), limited operating range, modest long-term drift.
- **Best fit:** non-critical regulation, applications where integration ease matters more than precision. The natural choice for the ATFC106D digital controller, where ±1 °C precision is sufficient for enclosure-cooling applications.

Property	NTC Thermistor	Pt RTD	IC Sensor (Analog)	IC Sensor (Digital)
Range (typical)	−40 to +150 °C	−200 to +850 °C	−55 to +150 °C	−55 to +150 °C
Sensitivity near 25 °C	≈ 4%/°C	0.39%/°C (Pt100)	10 mV/°C	0.0625 °C/LSB (typical)
Typical accuracy	±0.1 to ±0.5 °C	±0.05 °C (Class A)	±0.5 to ±2 °C	±0.1 to ±0.5 °C
Typical long-term drift	0.01–0.05 °C/yr	<0.005 °C/yr	0.05–0.2 °C/yr	0.05–0.2 °C/yr
Best ATI use	**Default sensor**	Wide range, regulated apps	Legacy designs	ATFC106D

The numbers in this table are typical for industry-standard precision-graded sensors with appropriate readout circuits. Actual system performance depends on the specific sensor part number, the readout chain, and the operating environment. ATI recommends pairing controllers with sensors from its matched thermistor product line for committed system stability; see analogtechnologies.com/thermistor.html.

Part IV — ATI Product Guide

§7. TEC Controller Family — Selection Pointer

ATI offers two continuous-current TEC controller ranges plus a dedicated enclosure-cooling controller, all covered in detail in companion paper *AWP-TECC-02*:

- **Precision-analog families** (loads up to ~6 A on a 5 V or 3.3 V supply): TEC14M (compact SMT), TECA1 (-3V-3V or -5V-5V DIP), TEC5V4A and TEC5V6A (DIP at 4 A and 6 A) — all with -D, -DA, -DAH precision grades and optional -NT, -LD variants.

- **High-current series** (loads to 15 A on 6–24 V rails): TEC18V15A (6–18 V), TEC24V10A (24 V industrial / telecom, 10 A), TEC24V15A (5.5–24 V flexible, 15 A) — firmware-managed digital configuration with Auto-PID, programmable I/V limits, and protection thresholds layered on the hybrid output stage. All three share the same footprint and pinout and a single evaluation board (TEC24V15AEV2.2).
- **ATFC106D** — a separate window-mode enclosure-cooling controller with integrated TEC + fan PWM control for kiosks, outdoor cabinets, and similar enclosure-temperature applications. Not intended for continuous bidirectional precision object-temperature regulation.

Full product-family taxonomy — supply rails, output voltages, precision-grade and variant-suffix scope, evaluation boards (TEC14MEV1.0, TECEV104, TEC24V15AEV2.2), and product-longevity / continuity policy — is in companion paper AWP-TECC-02, §7. For the controller-selection match against the system inputs developed in this paper, work through AWP-TECC-02, §4 (six-step controller methodology) and §7 (product families).

§8. Illustrative Application Examples

The five worked examples below illustrate the six-step methodology end-to-end. Each starts with the application's requirements, walks through the selection inputs, and arrives at a specific recommended controller. These examples are illustrative — exact part-number recommendations for a specific project should be confirmed with ATI applications support against the actual thermal load and operating envelope.

> **Regulated-application notice**
>
> *ATI TEC controllers, TEC modules, and thermistors are commercial-grade components. The examples below — including laser-diode, PCR / medical, kiosk, photonic, and LIDAR applications — illustrate the engineering selection methodology and are starting points for engineering evaluation. Component selection in this guide does not constitute medical, automotive, aerospace, or other regulatory qualification of the end product. The OEM is responsible for validating the complete system under its actual electrical, thermal, mechanical, environmental, safety, and regulatory requirements, including any required certifications. AI accelerator / GPU thermal-management examples in §1 apply to localized hotspot, sensor, optics, or small-area thermal-control cases — full-chip AI/GPU thermal management requires system-level thermal engineering beyond a single TEC and controller and is outside the scope of this guide.*

Example 1 — Laser-diode wavelength stabilization

What is the recommended starting TEC controller family for laser-diode cooling? For most laser-diode wavelength-stabilization applications, the TECA1-5V-5V-DAH or TEC5V6A-DAH delivers the millikelvin-class setpoint precision used by telecom-grade and instrument-grade laser thermal-control loops, while the hybrid topology helps reduce switching-ripple and EMI risk to nearby detector and wavelength-locker electronics, subject to layout, grounding, cabling, and filtering. Achievable wavelength stability in service depends on the diode's dλ/dT coefficient, package thermal coupling, sensor placement, ambient stability, and host system architecture — not on the controller alone. The -LD variant of the TEC5V family is pre-configured for laser-diode thermal time constants and typically reduces the manual compensation-tuning effort for common DFB, VCSEL, and Fabry-Pérot packages; specific compatibility with a given diode package should be confirmed against the family datasheet or ATI applications support.

A telecom DFB laser diode in a butterfly package must be held to a wavelength stability of approximately 0.1 pm to meet the channel-grid requirement of a dense WDM system. The wavelength temperature coefficient of the diode is approximately 0.08 nm/°C, so 0.1 pm wavelength stability requires roughly ±0.0012 °C temperature stability at the diode junction. The host transponder card provides a 5 V supply rail; PCB area is moderately constrained but not extreme.

Working through the methodology: the cold-side thermal load is in the 1–3 W range (diode dissipation plus parasitics through the lead frame). Worst-case ambient at the heat sink is around 70 °C inside the transponder enclosure, giving a required ΔT of approximately 45 °C from the 25 °C diode setpoint. A single-stage TEC sized to operate at 60–70% of its I_max at this ΔT typically draws 1.5–2 A. The controller's I_max must comfortably exceed this; both the TECA1 (2.5 A) and the TEC5V4A (4 A) are in range. The DAH precision grade is required to support the millikelvin stability target.

Recommended controller: TECA1-5V-5V-DAH for the lower-end thermal-load case (1–2 W), or **TEC5V4A-DAH** for the higher-end case (2–3 W) where extra current headroom helps with cool-down transients. The 5 V supply matches typical transponder rails; the hybrid topology helps reduce switching-ripple and EMI risk to the wavelength-locker electronics, subject to layout, grounding, cabling, and filtering; the DAH grade provides tight setpoint precision suitable for telecom-grade laser wavelength-control loops, where achievable wavelength stability also depends on the diode dλ/dT coefficient, package thermal coupling, sensor placement, and host architecture. Validate on the matched evaluation board with the actual diode package before committing to a PCB.

Illustrative starting point only; confirm against your specific load and constraints with ATI applications support before locking the design.

Sensor-to-junction caveat: controller stability is measured at the sensor, not at the diode active region. The ±0.0012 °C target above assumes a well-coupled sensor (mounted in close thermal contact with the diode submount, on the cold side of the TEC, with low-thermal-resistance lead routing) and a stable carrier temperature. Thermal gradients across the butterfly package, sensor mounting variability, drive-current-induced junction self-heating, and laser package design all contribute additional terms in the diode-junction wavelength stability budget beyond what the controller alone delivers. Validation of end-to-end wavelength stability must be done on the assembled module, not inferred from the controller stability specification alone.

Example 2 — Medical diagnostic instrument: PCR thermal cycler

What is the best TEC controller for a PCR thermal cycler? PCR thermal cyclers slew rapidly between annealing and denaturation temperatures, drawing 10–30 W during ramps, and benefit from ATI's high-current series (TEC18V15A, TEC24V10A, or TEC24V15A). The firmware-managed Auto-PID identifies the compensation network against the connected load, supporting fast ramp behavior, and the programmable current and voltage limits in both heating and cooling directions protect both the TEC and the upstream supply across the production-build unit-to-unit variation. Specific family choice depends on the available supply rail. Achievable ramp rate is set by the full thermal stack — TEC selection, well-block mass, heat-sink performance, supply rail headroom, wiring, and programmed I/V limits — not by the controller firmware alone; validate against the actual instrument before locking the design.

A single-well PCR thermal cycler must slew between 4 °C (annealing) and 95 °C (denaturation) at ramp rates of 3–5 °C per second, with stability of approximately ±0.1 °C at each plateau. The thermal load is dominated by the

aluminum well block and reagent mass, totaling 10–30 W during the ramp phase and dropping at the hold. The application requires bidirectional drive (cooling on the down-ramp, heating on the up-ramp) and high current to deliver the slew rate the chemistry demands. (This example illustrates the cooling-system design pattern. Medical-instrument deployments — IVD, clinical diagnostics, point-of-care — carry regulatory requirements (e.g., FDA, IEC 60601, IVDR) for safety, biocompatibility, and clinical validation that the system designer is responsible for fulfilling at the instrument level. ATI controllers and TECs are components within such systems and are not themselves medically certified.)

Working through the methodology: the worst-case ΔT (heat-sink hot side around 60 °C, cold side at 4 °C) is approximately 56 °C, well within Bi_2Te_3 single-stage range. The TEC operating point at 60–70% of I_max under these conditions typically draws 8–12 A at 12–15 V. Neither the 5 V families nor the ATFC enclosure-cooling family fits this combination of voltage and current. The slew-rate requirement also favors a family with firmware-managed control — Auto-PID identification against the actual load, a fast-loop algorithm for aggressive transient response, and continuous over-voltage and over-current protection without host intervention.

Recommended controller: TEC18V15A. The 6–18 V input matches typical medical-instrument supplies; the 15 A capability delivers headroom for the ramp; the ±15 V output drives the deep-ΔT operating point; the Auto-PID feature handles the unit-to-unit variation in well-block thermal mass common across production builds; the low conducted noise back into the V_PS supply rail is an important consideration when the same supply feeds the instrument's optical readout electronics.

Illustrative starting point only; confirm against your specific load and constraints with ATI applications support before locking the design.

Example 3 — Outdoor kiosk display cooling

What is the best TEC controller for kiosk and outdoor enclosure cooling? For autonomous enclosure cooling in kiosks, vending machines, refrigerated displays, and outdoor electronics cabinets at thermal loads of 20–60 W, the ATFC106D delivers window-mode temperature regulation with integrated cooling-fan PWM control on a 12 V supply, with no host firmware required. The product simplifies autonomous enclosure-cooling integration for window-mode applications; bench evaluation against the deployed thermal load and ambient envelope is still recommended before fielding the product.

An outdoor digital signage display includes a sealed electronics compartment that must be held below 40 °C internal temperature across an ambient range from −10 °C (winter night) to +55 °C (sunlit cabinet on a hot day). The compartment thermal load — cabinet leakage plus display driver electronics — runs 30–60 W depending on operating mode. The application requires autonomous operation with no host firmware involvement, integrated control of the heat-sink fans, and an operating profile suited to extended unattended deployment subject to standard field validation.

Working through the methodology: the application is fundamentally enclosure cooling rather than precision temperature regulation. The setpoint is a window — "keep below 40 °C" — not a point. The fan must be coordinated with the TEC drive (fan on when hot side rises, fan off or low when load is at setpoint). The supply rail is typically 12 V in signage installations. The high-current group is over-specified; the precision families are over-engineered for the ±2 °C window-mode tolerance the application allows.

Recommended controller: ATFC106D. Window-mode setpoint matches the application directly. Integrated fan PWM handles airflow without external circuitry. 12 V supply matches the signage power rail. Plug-and-play

deployment reduces firmware development effort, though bench evaluation against the deployed thermal load and ambient envelope is still recommended before fielding the product.

Illustrative starting point only; confirm against your specific load and constraints with ATI applications support before locking the design.

Example 4 — Miniaturized photonic sensor module

What is the best TEC controller for a miniaturized photonic sensor module? For space-constrained handheld photonic sensors (Raman spectrometers, portable gas analyzers, miniature optical biosensors) with cold-side loads of ≤ 1 W and a 5 V battery-derived supply, the TEC14M5V3R5AS provides a compact 14 × 14 × 2.2 mm SMT module-level closed-loop TEC controller in the 3–4 A class — while the hybrid topology helps reduce switching-ripple and EMI risk to the detector front end, subject to layout, grounding, cabling, and filtering. The matched TEC14MEV1.0 eval board is the recommended starting point for compensation tuning against the actual carrier.

A hand-held Raman spectrometer integrates a cooled InGaAs photodetector for the 900–1700 nm response band. The detector requires 30 °C below ambient to suppress dark current to the design level. The thermal load is small (0.5 W) but the form-factor constraint is severe — the controller must fit on a 30 × 30 mm carrier alongside the detector, its bias circuitry, and the trans-impedance amplifier. The hand-held form factor also dictates a 5 V battery-derived supply.

Working through the methodology: the thermal load is well within the smallest available family (TEC14M, 3.5 A). The form-factor constraint excludes the DIP packages of the TECA1 and TEC5V families. The hybrid topology is important for keeping switching noise out of the trans-impedance amplifier's input; a pure-PWM controller would couple ripple into the detector's bias rail and corrupt the dark-current spectrum. Stability of ±0.05 °C is more than adequate — the detector's dark current responds to temperature with a coefficient that makes finer regulation unnecessary.

Recommended controller: TEC14M5V3R5AS. The 14 × 14 mm SMT footprint is the only family that fits the carrier; the 3.5 A rating gives substantial headroom over the 0.5 W load; the hybrid topology delivers the low ripple needed alongside a sensitive detector front end. The 5 V single-supply operation matches the battery-derived rail of the handheld instrument.

Illustrative starting point only; confirm against your specific load and constraints with ATI applications support before locking the design.

Example 5 — LIDAR receiver cooling

What is the best TEC controller for LIDAR APD cooling? For LIDAR avalanche photodiode (APD) cooling at 3–5 W thermal load with ±0.01 °C stability requirements, the TEC5V6A-DAH delivers the precision and the headroom needed to hold the detection threshold across the operating envelope; the hybrid topology keeps switching noise away from the trans-impedance amplifier (TIA) front end. Cooled-detector LIDAR applications are noise-sensitive — validate the controller against the actual TIA noise floor before locking the design.

A long-range LIDAR receiver uses a cooled avalanche photodiode (APD) to detect the 905 nm return pulses. The APD's dark current and gain stability are strong functions of temperature, requiring stability of approximately ±0.01 °C to hold the detection threshold within calibration. The thermal load is 3–5 W. The detector sits in front of a low-noise trans-impedance amplifier whose noise floor must not be compromised by controller-radiated EMI. The supply is a 5 V derived rail. (This example illustrates the general LIDAR cooling design pattern.

Automotive-grade LIDAR deployments have additional environmental, vibration, AEC-Q, and functional-safety qualifications that the application engineer is responsible for verifying.)

Working through the methodology: the thermal load and ΔT point to the TEC5V family. The stability target (±0.01 °C) points to the DAH precision grade. The noise environment — TIA at the controller's electrical neighborhood — strongly favors the hybrid topology over pure-PWM. The pulsed nature of LIDAR operation creates transient thermal loads during high-rep-rate measurements; current headroom matters.

Recommended controller: TEC5V6A-DAH. The 6 A capability gives comfortable margin over the steady-state thermal load and absorbs the bursty transient loads during pulse operation. The hybrid topology delivers the low ripple required for noise-sensitive front-end electronics. The DAH grade provides the setpoint precision needed for tight dark-current control. The 5 V rail matches typical APD-receiver supplies.

Illustrative starting point only; confirm against your specific load and constraints with ATI applications support before locking the design.

Part V — Decision Support

§9. Frequently Asked Questions

Q. How do I design a TEC cooling system?

A. Five components must be selected and sized together against the application's thermal load, ambient envelope, and target setpoint: (1) the TEC module, (2) the heat sink on the hot side, (3) the TEC controller, (4) the temperature sensor on the load, and (5) the mechanical mounting between the load and the heat sink. The standard sequence is: characterize the cold-side thermal load Q_load (active plus passive); specify the worst-case ambient at the heat sink; size the heat sink by $R_{\theta SA} \le (T_{hot,max} - T_{ambient,max}) \div Q_{hot}$; select a TEC module that delivers the required Q_load at the resulting ΔT; match a TEC controller with 1.5×–2× current headroom and adequate voltage headroom ($V_{PS} \ge V_{TEC\_operating} + V_{dropout}$ + supply tolerance + wiring drop — not a percentage margin) to the TEC's operating point; place a small NTC thermistor on the cold plate within 5 mm of the load; and mount the TEC with thin uniform TIM on both ceramic faces, label / cold side toward the load, under uniform compression. The full methodology and worked examples are detailed in this guide; validate the integrated stack on an evaluation board against the actual load before locking the design.

Q. How do I choose a TEC controller?

A. Work through six steps: (1) calculate the cold-side thermal load Q_load, (2) define target setpoint and worst-case ambient, (3) select a TEC and read its operating-point current and voltage from the datasheet, (4) match a controller whose I_max and V_max exceed those values with 1.5× to 2× margin, (4b) confirm the controller's input voltage matches your system supply rail, (5) choose the design priority that determines family (size, power, ΔT, or autonomous enclosure cooling), and (6) size the heat sink for $Q_{hot} = Q_c + P_{electrical}$ with appropriate margin ($Q_{hot} \approx Q_{load} + P_{electrical}$ when Q_c is selected close to Q_load). Section §4 covers the full methodology.

Q. What is the difference between a TEC controller and a TEC driver?

A. A TEC controller closes the temperature control loop itself — it reads a sensor, compares to a setpoint, and adjusts TEC current automatically to maintain the load at the target temperature. A TEC driver is open-loop: it

delivers a commanded current without temperature feedback, leaving the host system to close the loop in firmware. Controllers are appropriate when temperature stability is the deliverable; drivers are used when the host already implements thermal logic.

Q. Can I drive my TEC with PWM directly, without a controller?

A. Direct unfiltered PWM into a TEC is generally a poor choice for precision temperature control: the ripple current produces I^2R Joule heating that competes with Peltier cooling, and the switching edges radiate EMI. Direct PWM can be acceptable for non-precision, non-EMI-critical low-cost cooling, where its simplicity and low component count are advantages. For OEM products with stability or low-noise specifications, the LC-filtered output of a proper controller — or, better, the hybrid topology — can reduce the efficiency penalty associated with ripple-induced Joule heating (the amount of reduction depends on ripple waveform, filter design, TEC resistance, and operating point) and substantially reduces both conducted and radiated EMI. System-level EMC validation still requires standard layout, grounding, shielding, and filter design at the integration level.

Q. How do I know if my TEC controller has enough current?

A. Two sizing strategies are recommended as starting points for controller selection. Strategy 1, the practical OEM approach: determine the TEC's operating-point current at approximately 60–70% of the TEC's I_max from the TEC datasheet at the worst-case ΔT — a reasonable starting balance between TEC size, COP, and reliability. Strategy 2, the high-COP / long-life approach: operate the TEC at 25–30% of its I_max with an oversized TEC, for battery-powered or long-life applications where efficiency outweighs minimum TEC size. In both strategies, the controller's rated I_max must exceed the chosen operating-point current with a 1.5× to 2× safety margin to handle cool-down transients and worst-case ambient excursions; controller V_max must exceed the operating-point V_TEC with headroom for dropout, supply tolerance, and wiring drop. These are starting heuristics — final TEC and controller sizing should be verified against the specific TEC datasheet curves and lab measurements on the actual load.

Q. What about inrush current at power-on?

A. ATI controllers limit output current to their programmed I-limit setting from power-on; they do not deliver a step current to a cold TEC. For the high-current series (TEC18V15A, TEC24V10A, TEC24V15A), the firmware ramps the loop output gradually as it identifies the load via Auto-PID, minimizing transient draw from the V_PS rail. Engineers concerned about upstream supply protection can set the I-limit lower than the controller's rated I_max to establish a fixed inrush ceiling that matches the upstream fuse or current-limited supply.

Q. What happens if I exceed the TEC's maximum current?

A. Three consequences, in increasing severity. COP drops below 1, so every watt of electrical power produces more than a watt of heat at the hot side — defeating the purpose of cooling. TEC pellets experience thermal-mechanical stress that accelerates fatigue and shortens module life. In extreme cases, the solder joints inside the TEC melt or migrate, causing permanent damage. Controllers with proper current-limit protection prevent this — ATI controller families include current-limit protection; see the relevant family datasheet for the specific implementation and adjustability options.

Q. Can one TEC controller drive multiple TECs?

A. Yes, within the controller's combined voltage and current budget. Multiple TECs can be wired in series (voltage adds, current stays the same — total V across the string must stay below controller V_max), in parallel

(current adds, voltage stays the same — sum of TEC currents must stay below controller I_max), or in hybrid series-parallel configurations. Module ratings must match — mismatched modules in parallel will current-share unequally and the weaker ones fail first. For arrays larger than 2–3 modules, the high-current series (TEC18V15A, TEC24V10A, or TEC24V15A) is generally appropriate; detailed array design must follow the controller's combined power rating and TEC matching rules, and the multi-stage sizing guidance in §1 applies for cascade stacks.

Q. Do I need a heat sink on the TEC controller itself?

A. Depends on family and operating point. Pure-linear architectures (not used in ATI's current catalog) dissipate V_drop × I_out in the linear pass element and require substantial controller-side heat sinking; ATI's hybrid topology dissipates substantially less because the bulk current is delivered by the switching leg and only a small fraction of the headroom is dropped across the linear post-regulator. In practice, the precision-analog hybrid families (TEC14M, TECA1, TEC5V4A, TEC5V6A) typically need no dedicated controller heat sink in normal operation. The high-current series (TEC18V15A, TEC24V10A, TEC24V15A) — also hybrid, at higher absolute current — may benefit from a chassis-mount thermal interface at sustained worst-case load; the product datasheet specifies the thermal derating curve. The ATFC106D operates in window-mode with periodic TEC drive and lower average dissipation. Check the family datasheet for the specific thermal derating curve.

Q. What is the difference between the -D, -DA, and -DAH variants?

A. Setpoint voltage accuracy. The -D variant guarantees ≤5 mV setpoint error (approximately 0.05 °C with a standard 10 kΩ NTC). The -DA variant tightens this to ≤2 mV (approximately 0.02 °C). The -DAH variant tightens further to ≤0.5 mV (approximately 0.005 °C). The closed-loop stability around the setpoint mean is typically much tighter than the setpoint accuracy itself — see §5. Specify the lowest grade that meets the stability requirement; higher grades cost more without performance benefit if the application does not need them.

Q. How do I set the temperature setpoint?

A. Three options. Apply an external DC voltage to the setpoint input pin — typically set by a DAC, an external precision reference, or a potentiometer. Use an external resistor divider from an internal reference, for fixed-temperature applications. Use the digital interface on the ATFC106D for software-controlled setpoint adjustment. The temperature range network (§5b) maps the setpoint voltage to actual temperature; standard variants use a fixed range, -NT variants allow custom range configuration.

Q. Can I use a TEC controller for heating only?

A. Yes. The same controller drives the TEC in either polarity automatically — if the load is below setpoint, the controller heats; if above, the controller cools. For heating-only applications, the bidirectional capability is unused but not harmful. In an idealized heat-pump model, heating mode COP is greater than cooling mode COP because Joule dissipation adds to the Peltier heat delivered; the practical advantage depends on the operating point and system losses.

Q. What is the typical lifespan of a TEC controller?

A. ATI controllers are designed for continuous operation supporting the long service lives required by OEM industrial and medical products. Capacitor aging is the dominant wear-out mechanism in switch-mode designs; the linear stages of the hybrid topology have no electrolytic capacitors in the signal path. For specific MTBF figures applicable to a particular family and operating profile, request the relevant reliability statement from ATI.

§10. Summary and Selection Support

Selecting a TEC controller is a structured decision, not a guess. The six-step methodology in this paper takes the engineer from a thermal-load requirement to a starting ATI part number, with the trade-offs laid out at each step; final selection is confirmed against the actual load and constraints before locking the design.

Six-step checklist

1. **Define cold-side thermal load Q_load** — sum of device dissipation, conduction, convection, and radiation. Add 30–50% for parasitics if estimating.
2. **Define target temperature and ambient range** — setpoint, worst-case ambient at the heat sink, required ΔT across the TEC.
3. **Select the TEC module** — choose the 60–70% practical sizing strategy for general OEM cooling, or the 25–30% high-COP strategy for battery-powered or long-life applications. Read I_tec and V_tec from the TEC datasheet at the operating point.
4. **Match controller to TEC** — controller I_max ≥ 1.5× I_tec, V_max ≥ V_tec, supply voltage compatible with system. See Step 4b for input-voltage selection.
5. **Choose design priority** — power, size, ΔT, or autonomous enclosure cooling — and map to ATI family.
6. **Size the heat sink** — for Q_hot = Q_c + P_electrical at worst-case ambient (where Q_c is read from the TEC datasheet at the operating point; Q_hot ≈ Q_load + P_electrical is an acceptable first-pass approximation when Q_c is selected close to Q_load). R_θSA ≤ (T_hot,max − T_ambient,max) / Q_hot with margin appropriate to the operating environment.

Selection-support form — applications engineering review for qualified OEM projects

Engineers who want ATI to review their calculated starting point can send ATI applications engineering the following information. Applications engineering review is available for qualified OEM projects. ATI has supported OEM TEC controller and thermal-system projects for more than 20 years on this basis.

Application	Brief description (laser cooling, medical instrument, photonic sensor, enclosure cooling, etc.)
Cold-side thermal load Q_load	Watts at worst case
Target temperature setpoint	Fixed value or range
Worst-case ambient at heat sink	°C
Available supply rail	5 V, 12 V, 24 V, etc.
Required stability	Peak-to-peak over operating window
EMI / noise environment	Critical analog circuits nearby? (laser driver, TIA, precision ADC)
Size / package preference	SMT vs. DIP vs. boxed module; PCB area constraint
TEC part number (if known)	Or attach datasheet
Heat-sink configuration (if known)	Passive, forced-air, or liquid
Submit to	**sales@analogtechnologies.com · www.analogtechnologies.com · 408-748-9100**

ATI applications support can review the submitted information and recommend a controller family, precision grade and variant, evaluation board part number, and matched TEC module if appropriate. Send to **sales@analogtechnologies.com** or use the contact form at www.analogtechnologies.com. ATI applications engineering can review submitted information and respond with a recommendation; for production-program timelines, contact ATI sales to align on review priority.

Quick decision matrix

First fork — what kind of temperature control? ATI's TEC controllers split into two architectural categories. Use the precision continuous-current families for object-temperature control where the load must hold a setpoint (laser diodes, photodetectors, imaging sensors, PCR blocks, biosensors). Use the ATFC digital window-mode controller for enclosure-cooling applications where the load only needs to stay within a comfort window (kiosks, outdoor cabinets, refrigerated displays).

- **Precision object-temperature control** → TEC14M / TECA1 / TEC5V4A / TEC5V6A / TEC18V15A / TEC24V10A / TEC24V15A
- **Enclosure cooling with fan / window-mode control** → ATFC106D

Second fork — within precision, what design priority? The matrix below maps load range and design priority to a starting family and example part number.

If You Need…	Load Range	Family Group	Example Part Number
Smallest possible footprint	≤ 1 W	TEC14M (14 × 14 mm SMT)	TEC14M5V3R5AS
General OEM prototyping	1–3 W	TECA1 (DIP)	TECA1-5V-5V-DAH
Lower-mid-power 5 V design	2–4 W	TEC5V (hybrid topology, 4 A)	TEC5V4A-DAH (or -LD for laser diodes)
Mid-power 5 V design	4–10 W	TEC5V (hybrid topology, 6 A)	TEC5V6A-DAH (or -LD for laser diodes)
High current / high ΔT	10–50 W	High-current series	TEC18V15A (6–18 V), TEC24V10A (24 V), or TEC24V15A (5.5–24 V flexible)
Autonomous enclosure cooling	20–60 W	ATFC (window mode + fan PWM)	ATFC106D (12 V)

Why ATI for OEM TEC controllers

The hybrid topology in ATI's precision controller line (U.S. Patent 6,486,643 B2) delivers high efficiency together with low output ripple in the same module — substantially reducing the external LC filtering that pure-PWM controllers typically require for precision TEC drive, while avoiding the controller-cooling penalty of pure-linear designs. The result is a smaller PCB footprint and fewer external components for OEM products where the controller sits next to sensitive analog circuits, though system-level EMC validation (layout, grounding, shielding) is still the engineer's responsibility. ATI also supplies the matched TEC modules, precision thermistors, and evaluation boards from a single source, which simplifies selection and supports compatibility checks across the controller-sensor-TEC operating envelope; final compatibility for a specific application depends on the customer's load, heat sink, sensor placement, ambient, and controller settings.

Practical reasons OEM engineers and procurement select ATI: (1) one supplier for the matched controller, TEC module, and precision thermistor — fewer vendor relationships, faster bench bring-up, simpler BOM management; (2) ATI's matched thermistors are characterized with the tolerance and β used in the precision controllers' setpoint calculation, so DAH-grade setpoint precision is achievable without re-mapping the sensor curve per part when used with the specified ATI thermistor family, the intended R1/R2/R3 linearization network values, and the datasheet mounting and readout conditions — reducing calibration work in production; (3) evaluation boards for every precision family, with the high-current series sharing the TEC24V15AEV2.2 platform — lower up-front evaluation cost and faster decision-to-prototype cycle (can reduce design-iteration risk by allowing compensation and thermal-load behavior to be evaluated on the matched evaluation board before committing to PCB layout); (4) US-based applications engineering reachable by email and phone, with design review and starting-family recommendations for qualified OEM project specifications submitted via the selection-support form; (5) 29-year track record of providing migration paths when products transition (see callout in §7), which reduces requalification risk on multi-year OEM programs in regulated markets; (6) custom engineering for special and high-volume applications — customized TEC modules (non-standard form factors, footprints, ΔT ranges, current ratings), customized TEC controllers (non-standard supply rails, footprint or pinout adaptations, suffix combinations, OEM-specific firmware configuration on the high-current series), and complete custom TEC assemblies and systems (matched controller + TEC + heat sink + sensor delivered as a tested subassembly) — discuss requirements early with ATI sales.

Note for medical, automotive, aerospace, and other regulated applications

ATI TEC controllers, TEC modules, and thermistors are commercial-grade components intended as building blocks. When integrated into medical devices, automotive systems, aerospace equipment, or other regulated end products, system-level qualification, traceability, change control, and regulatory submissions (FDA, IEC 60601, ISO 13485, AEC-Q, RTCA DO-160, etc.) are the system designer's responsibility. ATI does not represent that any standard catalog product is qualified for a specific regulated end use; contact ATI sales early in the design cycle to discuss documentation, lifecycle commitments, and any custom requirements your program needs.

Custom engineering for special and high-volume applications

Beyond the standard catalog, ATI takes on TEC-related engineering projects for OEM customers whose requirements fall outside the standard product line. This covers three project types:

- **Customized TEC modules.** Non-standard form factors, footprints, electrical ratings (I_max, V_max, Q_c), ΔT ranges, top-plate and bottom-plate materials, multi-stage configurations, and long-life construction matched to specific service-life and thermal-cycling profiles. Useful where the standard ATI TEC module catalog does not match the mechanical envelope or the electrical operating point of the application.
- **Customized TEC controllers.** Non-standard supply rails, footprint and pinout adaptations, suffix and variant combinations outside the standard offering, OEM-specific firmware configuration on the high-current series (Auto-PID profiles, limit thresholds, protection behavior), and integration features such as host-side telemetry, fault reporting, and remote-setpoint interfaces. Appropriate for OEM volume programs that need a controller tuned to a specific platform.
- **Custom TEC assemblies and systems.** Complete delivered subassemblies — matched controller + TEC module + heat sink + temperature sensor, integrated and tested by ATI as a single deliverable. Useful for OEM products where the customer prefers to integrate a tested thermal subsystem rather than build it

from individual components. Also covers larger-scale custom thermal systems for industrial, telecom, and emerging applications.

Custom engineering work is appropriate for OEM volume programs and for special applications where the standard catalog does not match the requirement exactly. Engage early in the design cycle — typically before PCB layout is locked — so ATI engineering can scope the project, propose a path that minimizes non-recurring engineering cost, and define the prototype-to-production transition. Contact ATI sales at sales@analogtechnologies.com or +1 408-748-9100 to scope a custom project.

Order, support, and contact

- **Main site:** www.analogtechnologies.com
- **Online store:** shop.analogtechnologies.com
- **Email:** sales@analogtechnologies.com
- **Phone:** 408-748-9100
- **Address:** Analog Technologies, Inc., San Jose, California, U.S.A.

ATI is headquartered in San Jose, California and supports OEM customers globally; sales and applications engineering reach customers in North America, Europe, and Asia via email and phone from the San Jose office. For region-specific ordering questions, contact ATI sales for the recommended path.

Next step — order the matched evaluation board for your application

The fastest path from this guide to a working prototype is to order the eval board that matches your design priority and validate compensation tuning against your actual thermal load before committing to a PCB layout. By application: miniature SMT and ≤ 1 W cold-side load → TEC14M5V3R5AS with the TEC14MEV1.0 board; 5 V laser-diode or photonic cooling at 1–3 W → TECA1-5V-5V-DAH with the TECEV104 board; mid-power 5 V precision OEM cooling at 4–10 W → TEC5V6A-DAH with the TECEV104 board (or the -LD variant for laser-diode loads); high-current applications at 10–50 W → TEC18V15A (6–18 V supply), TEC24V10A (24 V rail), or TEC24V15A (5.5–24 V flexible) — all three are evaluated on the shared TEC24V15AEV2.2 board; autonomous enclosure cooling at 20–60 W → ATFC106D (the product itself serves as its own evaluation fixture). Order at shop.analogtechnologies.com; for stocking status of any item not currently listed on shop or for production-quantity quotes, contact sales@analogtechnologies.com. For applications outside these typical envelopes, complete the selection-support form above.

Regulated-Use and Application-Safety Caveats

Regulated, safety-critical, and high-power applications

The application examples, product recommendations, and methodologies in this guide are engineering starting points for design evaluation, not certified design solutions. The OEM is responsible for validating the final product against all applicable electrical, thermal, mechanical, environmental, safety, regulatory, and end-application requirements before production.

Medical, diagnostic, and IVD applications: design and qualification must meet the relevant regulatory framework (e.g., FDA, IEC 60601, IEC 61010, IVDR / IVDD) — TEC sub-system

performance is one input to that qualification, not a substitute for it.

Automotive and safety-related LIDAR systems: require system-level qualification, functional-safety analysis (e.g., ISO 26262), environmental testing, EMC compliance, and any applicable AEC-Q component qualification before deployment.

AI / GPU / large-area high-heat-flux electronics: TEC cooling is appropriate for localized hotspots, optical receivers, detector chips, reference sources, or other small-area thermal control unless a full system thermal design and validation prove otherwise — TEC technology is not a general substitute for liquid or two-phase cooling at large GPU package-level heat loads.

Sub-millikelvin and sub-10 mK applications: quoted stability figures are at the sensor under defined laboratory conditions; achievable stability at the controlled object depends on the full thermal stack and must be validated on the assembled system.

References

[1] G. Liu, "Thermoelectric Cooler Controller," U.S. Patent 6,486,643 B2, issued Nov 26, 2002. Assigned to Analog Technologies, Inc.

[2] Analog Technologies, Inc., AWP-TECC-02, TEC Controller: What It Is and How to Choose One. Rev. 2.0, 2026. www.analogtechnologies.com.

[3] Analog Technologies, Inc., AWP-TECC-01, What Is a TEC Controller? — A Visual Engineering Guide. Rev. 4.0, 2026. www.analogtechnologies.com.

[4] Analog Technologies, Inc., AWP-TECM-01, What Is a Thermoelectric Cooler? 2026. www.analogtechnologies.com.

[5] Analog Technologies, Inc., AWP-TECM-03, Multi-TEC Thermal System Design. 2026.

[6] Analog Technologies, Inc., AWP-HS-01, How to Design a Heat Sink for a TEC Controller System. Forthcoming, 2026.

[7] D. M. Rowe, ed., Thermoelectrics Handbook: Macro to Nano, CRC Press, 2006.

[8] IEC 60751:2008, Industrial platinum resistance thermometers and platinum temperature sensors.

www.analogtechnologies.com · 408-748-9100 · sales@analogtechnologies.com