What HMTc Is

A Neutral Primer for Counsel — What the Program Certifies, What a Licensee Signs, and What It Does Not Claim

Karen Pendergrass
Standards Architect and Chief Executive Officer, Paleo Certified, Inc. dba The Paleo Foundation
June 2026

This document describes the Heavy Metal Tested & Certified (HMTc) program plainly and completely, for counsel and other professional readers who need an accurate picture before relying on the program or recommending it to a client. It is deliberately not a sales document. It states what the program is, how it works, what a licensee actually agrees to, and — with equal care — what the program does not do, where its commercial and structural conflicts lie, and how those are disclosed and mitigated. It is the first of two documents: this primer describes the program; a companion brief, *The Evidence a Jury Will See*, makes the case for why a brand would enroll and how the program bears on heavy-metal litigation. Read this one first. Nothing here should be a surprise later; that is the point of writing it this way.

What HMTc is, in one paragraph

HMTc is a voluntary, fee-based, third-party certification program for finished infant and child food products, operated by Paleo Certified, Inc., a Florida corporation doing business as The Paleo Foundation. It tests certified products against per-category concentration limits for toxic heavy metals, on an ongoing surveillance basis, and authorizes a certification mark for products that meet those limits. The 2026 Edition of its Program Manual is the program's first bound publication (published May 18, 2026; DOI 10.5281/zenodo.20270512)[1]. Infant and child food is the first of a planned twenty-three-category framework; it is the only category certified today. The program is new, and this primer treats its youth as a fact to disclose, not a thing to manage.

What it certifies, and against what

The unit of certification is a finished product (a specific SKU), not an ingredient, a supplier, or a brand in the abstract. Each product is evaluated against per-category action levels for eight priority heavy metals, measured through a ten-analyte laboratory panel. The eight metals are classified into two tiers by toxicology, and the classification carries governance consequences.

Tier 1 — lead, cadmium, inorganic arsenic, and methylmercury — are metals that are carcinogenic, neurotoxic, or bioaccumulative with no established safe threshold in children. There is zero tolerance for exceedance: a confirmed result over the limit escalates immediately. Tier 2 — nickel, tin, aluminum, and chromium — have established tolerable intakes or shorter biological half-lives, and a limited transitional allowance (up to 150 percent of the limit) is available to them under a defined status. The ten-analyte panel is the eight total-metal measurements plus three speciated forms (inorganic arsenic, methylmercury, and hexavalent chromium) required where the product or a reflex trigger calls for them. Limits are expressed as concentrations (parts per billion, as sold) rather than per serving, which prevents compliance through serving-size manipulation. Fifteen specific product chapters — the infant-formula, cereal, purée, juice, and snack subcategories — carry the actual numeric limits.

Action levels, not safety thresholds

This is the most important conceptual point in the program, and the one most easily misunderstood. HMTc limits are action levels, not safety thresholds. They are risk-management triggers anchored to occurrence data and regulatory benchmarks. Exceeding one obliges sourcing, process, and quality interventions; it does not, by itself, mean a product is unsafe or that any consumer was harmed. The program states this expressly and requires it to be reflected in all program communications and mark usage. The corollary is that certification does not assert safety, and the mark is forbidden from implying it (see below).

Two related commitments follow. The program operates under an ALARA paradigm — As Low As Reasonably Achievable — which it treats as an enforceable obligation rather than aspirational language: a brand that certifies and then plateaus is not in compliance with the program's intent. And the program carries an explicit legal-compliance disclaimer: certification is not legal advice and does not replace a licensee's obligation to comply with every applicable regulation. A regulatory limit violation in any market of sale is a hard-stop condition regardless of certification status, and the operator commits to cooperating with regulators rather than shielding licensees from them.

How the limits are set

The methodology is fixed and applies identically across categories. For each product subcategory, the program pools the published-literature occurrence distribution for each analyte and sets the per-analyte limit at a defined percentile of that distribution (the 97th percentile for a cleaner subcategory, the 45th for a more contaminated one, a calibration chosen to deliver an intended rate of day-one certifiability). That percentile is then capped so that it can never exceed the lowest applicable government maximum in the brand's markets of sale; the final limit is the lower of the two. Where a limit is set tighter than either the literature or the regulatory floor, the program records the reason — precautionary, market-ratcheting, feasibility-driven, or regulatory-alignment — rather than leaving the gap unexplained.

HMTc limits generally sit below FDA action levels, and the reason is quantitative rather than rhetorical: a single day's consumption at the FDA action level across four common categories yields roughly 10.4 µg of lead per day, about 4.7 times the FDA's own Interim Reference Level of 2.2 µg/day for young children[1]. Full compliance with each single-product federal floor still produces aggregate infant exposure several times over the health-based benchmark, because a child eats across categories. The program calibrates across categories so that aggregate intake sits below the health benchmark, not merely below each single-product floor.

The standards-setting methodology is published in principle (in the Program Manual and in three peer-track preprints[1, 2]); the specific software engine that performs the pooling arithmetic is held as a trade

secret. The program commits to a controlled-terms audit process: regulators, litigation counsel, licensees auditing their own category, and approved academics may request the underlying aggregate distribution data and methodology documentation, with a thirty-business-day response commitment (regulators without an NDA; others under a confidentiality agreement that preserves the engine's trade-secret status while making inputs and outputs verifiable).

How the literature base is built – and the AI disclosure

The limits rest on the Heavy Metal Index (heavymetalindex.com), a curated synthesis of the heavy-metals-in-food literature that the Foundation operates as a property architecturally separate from the certification program. The separation is one-way and deliberate: the certification program cites the Index; the Index does not endorse the certification program. The intent is that the evidence base is not the certifier's self-published justification but an independent reference the certifier happens to operate, with every numeric claim traceable to a specific source.

Counsel should understand, clearly and without it being buried, how that literature base is produced. The Index is built using two artificial-intelligence systems — OpenAI's Codex and Anthropic's Claude Code — operating as parallel literature-processing instruments. They read the same corpus of peer-reviewed studies and government reports and produce the per-source pages that constitute the evidence layer. The two are drawn from different vendors on purpose, so that each independently audits the other's work against the original source on a five-point protocol (numerical fidelity, controlled-vocabulary validation, arsenic and mercury speciation discipline, the brand-data firewall, and the wiki/certification separation), with verdicts and corrections logged to version history. Neither system sets certification thresholds or decides pass/fail. Those determinations sit on the human-controlled side of the firewall, under a named principal investigator (the Standards Architect) with sole authority over them. The program analogizes the dual-instrument arrangement to the dual-reviewer protocol of a Cochrane systematic review, and discloses it at the front of its manual rather than in a footnote, on the reasoning that concealment would be the failure mode and that AI literature-processing is increasingly ordinary in systematic-review work while AI threshold-setting would be novel and hard to defend. A reader who is going to vouch for the program should be comfortable with this architecture, because a sophisticated adversary will raise it; the program's position is that the literature processing is AI-assisted and cross-verified, the standard-setting judgment is human, and the whole chain is auditable from version history.

The certification lifecycle: what a licensee actually does

Certification is a continuing relationship, not a one-time test. A brand establishes a three-lot baseline, receives a status, and then submits to ongoing surveillance.

The statuses are five. A is full compliance; B is transitional compliance (the Tier 2 variance); C is probation under heightened surveillance; D is suspension; E is revocation. Surveillance testing runs on a schedule tied to how much of a product a child consumes relative to body weight — monthly for the highest-exposure categories such as infant formula and cereal. Samples are drawn by a party independent of the brand's production team, analyzed by laboratories accredited to ISO/IEC 17025[3], moved under documented chain of custody (a documented breach invalidates the result), and recorded as structured data. Pass/fail decisions incorporate the laboratory's measurement uncertainty on a two-sided rule: a lot passes only if the result plus uncertainty is at or below the limit, fails only if the result minus uncertainty exceeds it, and is otherwise borderline and resolved by confirmatory testing before any status is assigned. The program prohibits "testing into compliance" — undocumented resampling, lab-shopping, or selective reporting is a data-integrity violation that costs a brand its status — so that the test record represents

the program's actual findings rather than a curated subset. A confirmed exceedance triggers corrective and preventive action on defined timelines. Testing fees are paid by the brand directly to its accredited laboratory; the operator does not mark them up.

The Confidential Remediation Track, and its honest limits

A brand whose baseline does not yet qualify for Status A or B — but whose lots still meet applicable legal limits — may enroll confidentially through the Confidential Remediation Track. Enrollment and test data are treated as confidential business information; the certification mark is withheld until at least one product qualifies; and the engagement is framed as a good-faith quality-improvement initiative. The intended value is a documented, timestamped, third-party-verified record of improvement that exists before any litigation or inquiry.

Counsel must understand the limits of that protection, which the program states plainly rather than glossing. Track data is held by the operator as a service provider and is potentially compellable under subpoena. The Track does not create attorney-client privilege or work-product protection over the testing data. The record is admissible *for* the brand and discoverable *by* a plaintiff, both at once. The operator commits to notifying the licensee of any subpoena and cooperating with its defense, but it expressly disclaims being a substitute for a counsel-directed quality program that may carry stronger privilege in a specific matter, and it directs brands to consult counsel about how the Track interacts with their broader discovery posture. This is the single most important thing for a recommending lawyer to internalize: the Track changes the *content* of the discoverable record (from "no program" to "documented good-faith improvement"), not its *discoverability*.

The mark, and what it may not claim

Only Status A and Status B products may bear the HMTc Certification Mark; products on probation, in suspension, or in the Remediation Track may not. The mark is governed by a published style guide and by hard limits on claims: a licensee may not state or imply "heavy metal free," "no toxic metals," "100% safe," or the absence of risk. The program states that the mark "does not certify safety" — it certifies that a product has been independently tested and meets defined action levels under ongoing surveillance. An anti-arbitrage rule prevents a brand from dropping the mark during a bad period and resuming it on a single lucky lot (re-entry requires three consecutive passing lots and closure of any open corrective action), and brands control whether their certified products appear in any public registry. The prohibition on safety claims is not incidental; it is what keeps the mark itself from becoming a misrepresentation a plaintiff could attack.

Governance, the operator's conflicts, and how they are mitigated

Two structural conflicts are inherent in a fee-funded certifier, and the program discloses both rather than leaving them to be discovered.

The first is institutional: a brand appealing a status determination is, at first instance, appealing to the entity whose revenue depends on the outcome. The program's answer is a two-tier appeals process. Tier 1 is internal; Tier 2 is an independent Appeals Committee of three external reviewers (a toxicologist or food-safety scientist, a quality professional, and a legal adviser), none employed by or contracted to the Foundation, structured so its funding does not depend on case outcomes. Tier 2 reversals bind the operator. The program treats disclosure of the conflict as necessary but not sufficient, and the independent committee as the actual mitigation.

The second is personal: the Standards Architect and the Chief Executive Officer are the same person (Karen Pendergrass). The program discloses this on the byline and in the colophon. The structural mitigations offered are the independent Appeals Committee above and the firewall separating the literature evidence layer from the threshold-setting. A reader weighing the program should weigh the disclosure and the mitigations together, and form an independent view of whether they are adequate; the program's position is that they are, and it does not ask the reader to take that on faith.

Beyond conflicts, governance includes a Technical Committee (advisory, with authority to recommend metal-tier reclassification), data-integrity hard stops, document-precedence rules, and the regulatory-escalation rule already noted (a regulatory violation overrides certification status). The methodology evolves on a defined cycle through the Standards Ratchet, below.

The Standards Ratchet: what the commitment becomes over time

HMTc limits are designed to tighten over time as a category's certified market improves, and never to loosen. Tightening fires only on defined triggers measured over two consecutive annual reviews, is capped at a 20 percent reduction per event, and transitions over 12 to 18 months with existing certifications grandfathered through the transition. The program documents anti-gaming provisions (the trigger pool combines Status A and B data so a brand cannot withhold its best results; any single licensee is capped at 25 percent of the trigger pool; a brand that exits and re-enrolls within 24 months is held to the current standard, not the one it left), and a forecast-accountability clause that publishes prior forecasts against actual outcomes and triggers a methodology review when they diverge by more than 25 percent. For a licensee, the practical meaning is that certification is a commitment to a moving floor, with the movement bounded, forecast, and announced in advance.

What a licensee signs

The license agreement is a separate, bilaterally executed document; the Program Manual summarizes its structure. License fees are revenue-tiered (bands for brands under $5M, $5–50M, $50–500M, and enterprise terms negotiated above $500M), and are charged per brand-tier rather than per SKU, so breadth of certification is not penalized; specific schedules are provided on request and are not published in the manual. Testing costs are separate and paid to the laboratory directly. Terms run one to three years with defined renewal and a 30-day termination right for either party (immediate for cause), and a 180-day inventory sell-off period except in for-cause termination. Confidentiality is mutual, with brand data treated as confidential business information that the operator may use only in aggregated, anonymized form, surviving termination for seven years. The agreement grants the licensee an audit right over the methodology behind its own category (the controlled-terms process, 30-business-day response). Disputes go to mediation and then AAA commercial arbitration seated in San Diego under California law — a choice the program ties to the fact that the governing heavy-metal case law (MDL 3101), AB 899, and Proposition 65 are all Californian; non-U.S. licensees may negotiate alternative forums. The operator commits in the agreement to regulatory cooperation and states that it will not shield a licensee from regulatory obligations.

Where HMTc sits in the landscape

HMTc is a private, voluntary program. It is not a government approval and does not replace regulatory compliance or a brand's own quality program. It is complementary to USDA Organic and the Non-GMO Project, which do not address heavy metals at all (organic produce can still carry soil- and water-derived metals). It overlaps with, and differs from, certifications and testing efforts that do address heavy

metals — Clean Label Project, EWG's scoring, Consumer Reports' periodic testing — chiefly in offering staged compliance with a confidential remediation pathway, a measurement-uncertainty decision rule, a mandatory tightening ratchet, cross-brand supplier intelligence, and a public, auditable evidence corpus. The program's own recognition chapter makes this comparison on the merits and explicitly declines to claim it is uniformly superior; it positions itself as the heavy-metal-specific, industry-facing certification framework, not a substitute for the others.

What HMTc is not (worth reading twice)

A neutral account has to be as clear about the limits as the features. HMTc is:

- Not a safety guarantee. It certifies action-level compliance under surveillance. The mark may not imply safety, and the program says so itself.
- Not regulatory approval, and not a substitute for compliance. A regulatory violation is a hard stop regardless of HMTc status.
- Not legal advice, and not immunity. It is documentation and an evidentiary posture. The Confidential Remediation Track's data is discoverable and unprivileged.
- Not a finished, road-tested institution yet. The 2026 Edition is the first; infant and child food is the only certified category so far; limits are calibrated against the global literature (regional calibration is on the roadmap, not in force); and the ratchet that tightens limits will run on lot data the program has only begun to collect.
- Not the work of an AI black box, and not free of human conflict. The literature layer is AI-processed with cross-vendor verification and a human principal investigator; thresholds are human-set. The operator is fee-funded and the founder holds a dual role; the mitigations are an independent binding appeals body and the evidence-base firewall.
- Not yet a registered mark. The HMTc certification mark is pending registration with the USPTO (Serial No. 99307306) and in parallel with the EU IP Office; it is not yet registered. The EU corporate entity that would carry EU database-right protection for the corpus is on the roadmap, not yet incorporated; current corpus protection rests on U.S. copyright and trade-secret law.

None of these is hidden in the program's own documentation. Several are stated on the front pages of the manual. A recommending lawyer should know each of them before speaking, and should be able to state them as readily as the program's strengths.

The program's own adversarial read

One feature is itself worth knowing, because it is unusual and bears on credibility. The Program Manual contains a Program-Level Adversarial Read that enumerates fourteen categories of attack a hostile expert, regulator, competitor, or opposing counsel would advance — on the methodology's reproducibility, the toxicological framing, the institutional conflict, the discoverability of Track data, the limits of anonymization, the choice of arbitration forum, and more — and answers each on the merits, citing where in the manual the answer lives and noting where the audit drove a revision. Each Part 3 product chapter carries its own adversarial read as well. For a lawyer deciding whether to stand behind the program, that chapter (Part 5.2) is the most efficient stress test available, and it is offered, not hidden.

In sum

HMTc is a young but unusually well-documented voluntary certification program: a literature-anchored, percentile-derived, regulatory-capped set of action levels for eight metals in infant and child food, applied

through independent, ISO 17025-accredited, measurement-uncertainty-aware lot surveillance, with a confidential pathway for brands not yet qualifying, a mark disciplined against overclaiming, a tightening ratchet, disclosed AI involvement in the literature layer with human-held thresholds, disclosed operator conflicts with structural mitigations, and an evidence base operated separately from the certifier. That is what the Foundation has built and what a licensee signs up for. The companion brief takes the next step and argues why a brand should.

Karen Pendergrass is the Standards Architect of the Heavy Metal Tested & Certified program and the founder and Chief Executive Officer of Paleo Certified, Inc. dba The Paleo Foundation. She can be reached at karen@paleofoundation.com. *This primer is descriptive and is not legal advice. Where it characterizes the program, the controlling text is the 2026 Program Manual and the executed license agreement.*

References

[1] K. Pendergrass, *Heavy Metal Tested & Certified (HMTc) Infant and Child Foods Program Manual*, 2026 Edition, Paleo Certified, Inc. dba The Paleo Foundation, 2026. doi: 10.5281/zenodo.20270512. Cited throughout: action levels and disclaimers (Part 2.1); analyte tiering and aggregate exposure (Part 2.2); status framework (Part 2.3); Confidential Remediation Track and discovery dynamics (Part 2.4); mark rules (Part 2.5); decision rules (Part 2.6); standards ratchet (Part 2.10); appeals and conflict mitigation (Part 2.12); methodological / AI disclosure (Part 3 methodological disclosure and Part 5.7); program-level adversarial read (Part 5.2); recognition framework (Part 5.5); licensing terms (Part 5.6).

[2] Peer-track preprints: K. Pendergrass, *The Counterproductive Consequences of Public Exposé Testing*, doi: 10.5281/zenodo.19470572 *The Cost of Operating Without Credible Third-Party Heavy-Metal Certification*, doi: 10.5281/zenodo.18903738 *Certification as a Framework for Reducing Heavy Metal Exposure in Infant and Child Foods*, doi: 10.5281/zenodo.18905821.

[3] International Organization for Standardization, ISO/IEC 17025:2017, General Requirements for the Competence of Testing and Calibration Laboratories.