Program structure

Karen Pendergrass · ORCID 0000-0002-2348-7259

Karen Pendergrass. "Program structure". https://heavymetalcertified.com/program/structure. Retrieved 2026-09-20.

Analyte tiering

Eight priority metals, tested as ten analytes because arsenic and chromium are speciated to the toxic form; methylmercury is a conditional eleventh, speciated by reflex on trigger and bounded by the total-mercury limit. Each metal is assigned to a tier by its toxicological profile, and the tier determines whether any exceedance can ever be tolerated.

Tier 1 PB · CD · IAS · MEHG

Carcinogenic with no established safe threshold, neurotoxic with no established safe threshold, or bioaccumulative with long biological half-lives in children. Any confirmed result above 100% of the limit escalates immediately.

Overage allowed — none

Tier 2 NI · SN · AL · CR

Established tolerable daily intakes, shorter biological half-lives, or primarily acute rather than chronic toxicity. Controllable through sourcing and process. An exceedance still bars the mark and enters Remediation, at an elevated rather than immediate hard-stop reaction; the tier sets the reaction, never a permission to exceed a limit.

Overage allowed — none

Speciation upgrades chromium to Tier 1 treatment when a credible hexavalent pathway exists. Cr(III) is an essential nutrient and is assumed in food; Cr(VI) is an IARC Group 1 carcinogen by ingestion.

How a limit is set

The default limit for any analyte is the strictest maximum level set by a credible government regulator, the EU, Codex, the FDA, WHO/JECFA, FSANZ, or Health Canada, converted to the row's native basis. It is a lookup, not a computation: it carries the Tier 2 metals directly, and where the law splits a limit by formulation the standard adopts the law's own split and cites it. The percentile machinery below runs only where the program sets a number itself: the four Tier 1 toxics, floored against the government maximum so the published value is the stricter of the two, and the cells no government regulates.

For those occurrence-set cells, subcategories are paired. Within a pair, one form carries a structurally higher contamination burden than the other: rice cereal against non-rice, soy formula against cow-milk, root vegetables against above-ground. Pooling the pair would raise the clean limit and flatter the dirty one in the same move, so they are separated.

P97 Clean subcategory

The lower-burden form within a pair: non-rice cereal, cow-milk formula, above-ground vegetables. Where the program sets this cell from occurrence rather than a government limit, the limit sits at the 97th percentile of the pooled distribution.

P45 Dirty subcategory

The structurally higher-burden form: rice cereal, soy formula, root vegetables, fish-containing. A far tighter percentile, because the point is to separate the better operators within a hard category rather than bless the category.

P97 Independent row

A subcategory with no within-pair partner defaults to P97.

> The percentile is a calibration choice, disclosed as one. A brand certifies a row of ten analytes at once, so per-analyte limits set at a raw P90 would produce a joint pass rate near 35%. For the occurrence-set cells only, the Tier-1 toxics floored against the government maximum, and the cells no government regulates, P97 and P45 are the values that deliver the program's stated row-level outcomes: roughly 90% of clean-subcategory brands and 40% of dirty-subcategory brands pass, confirmed at 94% and 42% against FDA per-sample lot data. A subcategory whose every published limit binds on a government maximum has no percentile-set cell, and this calibration does not touch it: infant formula powder is one, with all eight of its published values set by government lookup.

WHY WE PUBLISH THE PASS RATE

We are not aware of another certifier that publishes this number, and disclosing it invites an obvious reading: that the threshold was set to a commercial outcome rather than a health one. The reading is worth answering directly, because the alternative, not saying it, does not make the calibration disappear. It only makes it undiscoverable.

Every percentile-based standard in food safety embeds a pass rate. Codex maximum levels, EU regulation 2023/915 and FDA action levels are all set against achievability, using occurrence data and explicit statements about what industry can meet. The choice is never whether a pass rate exists; it is whether the body setting the limit tells you what it is.

The health-based floor is separate and binds first: no published value ever exceeds the lowest applicable regulatory ceiling in that basis. The calibration operates only in the space beneath that ceiling, where the literature alone does not determine a single number. Where the ceiling binds, the rationale tag reads `regulatory-alignment` and where an occurrence pool exists beneath a binding ceiling the pooled value is shown so the gap is visible. On infant formula powder no cell is occurrence-set: every published value binds directly on a government maximum: lead on the EU's 20 ppb.

Native basis

Every standard is published in the form the product is actually sold in: powder as placed on market, ready-to-feed as consumed, dry cereal as sold, purée as consumed. A brand ships and tests in one form, and the mark matches that form.

Sources reporting in another basis are converted inward during pooling, using documented factors: never the reverse. Where a cross-basis equivalent is shown next to a published value it is a reading convenience for audiences working in other units, not a second standard a brand may choose between.

Publication gates

A per-analyte value publishes only when it clears every gate below. A cell that fails any one of them stays unpublished rather than shipping with a caveat.

Gate 01 Sufficient pool

For a cell the program sets from occurrence: the Tier 1 toxics and the analytes no government regulates: enough contributing sources for the pooled distribution to mean something, with no single survey exceeding 50% relative share.

Gate 02 Native basis resolved

Every contributing source converted into the row's native basis using a documented factor, with the conversion recorded.

Gate 03 Regulatory ceiling loaded

The strictest applicable government maximum identified and expressed in the row's basis. For most analytes it is the standard; for a Tier 1 toxic it is the floor the occurrence value is capped against.

Gate 04 Analytical feasibility

The value must be measurable. Where the limit of quantitation binds above the pooled value, the limit is feasibility-driven and says so.

Gate 05

Rationale tagged

Exactly one of literature-baseline, regulatory-alignment, feasibility-driven or precautionary. An untagged value does not publish.

GATE NAMES NOT QUOTED VERBATIM

These are the program's stated publication conditions, drawn from the manual's methodology sections. The manual refers to "publication gates" without a single numbered list, so the names here are descriptive rather than quoted.

The four statuses

Defined in Program Manual Part 2. A status attaches to a product, not to a brand, and it is visible on the certificate. Certified is the only status that carries the mark.

STATUS	MEANING	MARK USABLE	CONDITION
Certified	Full compliance	yes	At or below every value in the Master Limit Table, all ten analytes, native basis, confirmed under the measurement-uncertainty rule.
Remediation	Not certified: applicant or corrective	no	Any analyte above its published limit, or a new applicant not yet at 100%. Not a softer certification and not an allowed overage: no mark, not listed as certified. Corrective action and intensified surveillance apply: a Tier 1 exceedance draws the most severe reaction, a Tier 2 exceedance an elevated one. The Confidential Remediation Track is the confidential form of this status and is likewise unmarked.
Suspension	Certification halted: hard stop	no	A government legal-limit violation, a data-integrity failure, or failure to execute testing, CAPA, or directives. The mark is removed from all units and marketing. Reinstatement requires a completed CAPA with verified effectiveness, a run of consecutive compliant lots set by the Program Operator, and restored data integrity: not a fresh baseline.
Revocation	Permanent termination	no	Willful misconduct, data fraud, concealment of a known exceedance, or persistent unremediated suspension. Re-entry by full re-application only.

The mark is never displayed above the published limit, on any metal, at any margin. There is no transitional band and no allowed overage: a product above any published limit is not Certified and carries no mark. Tiering sets how hard and fast the program reacts to an exceedance, not permission to exceed it: a Tier 1 exceedance draws the most severe reaction, a Tier 2 exceedance an elevated one.

Lot surveillance

Certification is continuous. Every SKU entering the program is tested across a minimum of three production lots to establish a baseline before any status is assigned at all: a single passing lot certifies nothing.

Baseline before any status	`3 lots minimum`
Qualifying prior accredited data	`credits up to 2 lots`
Certified	`routine protocol`
Remediation: Tier 2 exceedance	`elevated protocol`
Remediation: Tier 1 exceedance	`every-lot, hold-and-release`
Speciation	`on reflex trigger`
Chain of custody	`documented per lot`

A brand that already runs a mature, accredited lot-testing program does not repeat work it has already done. Up to two of the three baseline lots may be credited to qualifying prior results, and the baseline is completed with a single new full-panel lot that the program sizes and runs. This is the expedited baseline. Prior results qualify only where they come from an ISO/IEC 17025-accredited laboratory, are quantified at or below the program's reporting limit, were produced within the preceding twelve months, are supplied as structured data under chain of custody, are expressed in or convertible to the product's native basis, and are credited only for the analytes they actually measured. The completing lot is sampled independently and run across the full panel, including every analyte the prior data did not cover, so no analyte enters the baseline on a brand's word alone; where a single analyte rests on that one completing lot, the next two routine surveillance lots confirm it. The credit is available at ordinary program entry, not on the confidential remediation track.

Surveillance intensity follows status: routine while Certified; in Remediation it intensifies with the tier of the exceedance: elevated for a Tier 2 exceedance, every-lot with mandatory hold-and-release for a Tier 1 exceedance. Reflex testing triggers add speciation when a result or a formulation change indicates a credible pathway, and sampling follows a documented chain of custody so a result can survive challenge years later.

> The readiness and corrective-action pathways exist so a brand that finds a problem is better off reporting it than concealing it. While a self-reported item is open and on schedule, its register visibility and retailer disclosure follow the published status rules; a released lot confirmed above its applicable limit on product represented as certified is never shielded from a required register status change or from retailer, withdrawal, or recall action. A brand that conceals a known exceedance loses its certificate.

Two remediation pathways

When a lot falls short, a brand chooses how the correction is handled. Both routes require the same corrective action and the same re-testing; what differs is who sees it, and what the brand gets in return for that visibility.

Confidential track DEFAULT

The finding, the affected lots and the corrective action stay between the brand and the program while the item is open and on schedule. Register visibility, retailer disclosure, and the effect on the certificate all follow the published status rules, according to whether the affected lot was released and whether the product is currently represented as certified; a released lot confirmed above its applicable limit is never shielded from a required register status change or from retailer, withdrawal, or recall action.

WHAT THE BRAND GETS

Visible in register	per status rules
Disclosed to retailers	per status rules
Certificate effect	per status rules
Corrective action	required
Re-testing	required

Public track ELECTIVE

The brand elects to publish the finding and its remediation on its own brand page. The record shows what was found, what was done and when it closed: a documented history of catching and fixing a problem rather than an unexplained gap.

WHAT THE BRAND GETS

Visible in register	yes, with outcome
Citable by the brand	yes
Retailer questionnaire	answers itself
Corrective action	required
Re-testing	required

> Neither route is a way out of the corrective action. The choice is about disclosure, and it exists because a program that punished self-reporting would simply stop being told anything: which would make the mark worth less, not more.

Readiness and corrective action →

Governance

The program is operated and governed by the Institute of Contaminant Standards (ICS), which sets how limits are established without testing products or selling marks itself.

The Institute of Contaminant Standards

Governs how limits are set. It does not test products, sell marks, or publish threshold values or brand-level data itself: that separation is what lets it arbitrate.

Technical Committee

Toxicology, food-safety and analytical-chemistry experts with documented independence from licensees. Reviews methodology and recommends metal reclassification between tiers. It reports recommendations; it does not render binding decisions.

Independent Appeals Committee

Holds binding review authority on contested certification decisions. The Program Operator retains final authority on classification and methodology, subject to that review.

Substantive changes to a published limit run through the Standards Ratchet Mechanism, which requires a stated trigger and caps how far a limit may move in a single step. Limits tighten as evidence and analytical feasibility improve.