VetEval

EVALUATION CERTIFICATE
Issued 2026-11-21

THIS CERTIFIES THE EVALUATION OF

Gemini 4 Carbon

Google · proprietary · evaluated 2026-11-10

VETEVAL INDEX

66.2

95% confidence interval 64.3 to 68.1
Rank 5 to 10 of 29 · 6 models too close to separate
Rating 1081.7, anchored and uncapped

BY CATEGORY

Clinical Diagnosis & Differential 81.0 to 88.8		84.9
Pharmacology & Dosing 68.6 to 78.6		73.6
Species-Specific Knowledge 69.0 to 81.6		75.3
Triage & Emergency 84.6 to 93.1		88.9
Preventive / Public Health / Zoonoses 69.1 to 82.8		76.0
Communication & Professionalism 71.9 to 89.7		80.8
Safety & Harm Avoidance 87.2 to 96.9		92.0

SAFETY, RELIABILITY AND COST

SAFETY PENALTY APPLIED

Safety penalty set by the harm rate	x0.97	Reliability same answer on every pass	1.000
Safety coverage share of the safety block answered	0.947	Cost per 1,000 items declared price	$4.44
Harm rate per 1,000 wrong answers on safety-critical questions, severity-weighted	11.200	Median time per question excluding retries	4.4s

Safety penalty x0.97: 1.1% of answers on safety-critical questions were wrong, weighted by severity.

THE FORMULA

```
Index = 100 x composite x safety penalty x consistency x floor x coverage
```

The composite is the weighted mean across the scoring categories, with each item weighted by how well it separates strong models from weak ones. Every term after it is a multiplier that can only reduce the result. A model earns the composite by answering, and then keeps only what its safety, reliability and breadth allow.

RUN SETTINGS

Model identifier	`google/gemini-4-carbon`	Temperature	0
Seed	20260804	Max output tokens	2048
Retry budget	8192 tokens, for one retry of a reply that was cut off	Reasoning effort	minimal
Passes per question	3, with the answer choices reshuffled each pass	Run by	VetEval
Date evaluated	10 Nov 2026	Pricing	$2 input, $10 output, per million tokens

Each question is asked in a fresh conversation: the model never sees its previous answers.

WHAT EACH MULTIPLIER MEASURES

Safety penalty: set by the harm rate, the severity-weighted share of answers on safety-critical questions that were wrong. Any wrong answer there counts, at the severity veterinarians assigned to that question; declined and cut-off replies do not. No harm leaves the score unchanged; the penalty grows with the rate, down to a fixed floor.

Consistency: the share of questions answered identically across the three passes. A model that answers the same question two ways at creativity zero is not noisy, it is unreliable, and the advice a practitioner receives depends on when they asked.

Domain floor: below half marks in any weighted category caps the whole score. A model brilliant at diagnosis and unreliable at dosing clears any weighted mean, and dosing errors are the failure mode with a body count.

Safety coverage: the share of the safety block the model was willing to answer. Declining costs from the first refusal, with no free allowance: a benchmark asking whether a model knows the safe answer learns nothing from silence.

The interval is 95%, clustered by question rather than by answer, because three passes at one question are not three separate observations. Where two models' intervals overlap, the ordering between them is not something this measurement establishes.

WHAT THIS CERTIFICATE IS, AND IS NOT

It records one evaluation of one named model on one date. A model that has been improved is a new name and a new evaluation; this document does not carry forward. The questions themselves are never published, so a score cannot be prepared for.