


Engineering Audit: CC 2.1.220 Renewal & Opus 5 Adoption

Defect forensics, cascade analysis, and codebase realities.


[DATA_PACKET: 0xAEFD]



Defect Forensics

Isolating the **40% undercount** on Opus 5 routing.

[AUTO_VER: 2.1.220]



The Floor Cascade


Mapping the blast radius of **114 manual enforcements hidden** behind an automated version bump.

[CLAIM_ID]	[VERIFIED]	-> [STATUS]
[CLAIM_001:	[VERIFIED]	
[CLAIM_001:	SECURE	
[CLAIM_002:		[FAILED]
[CLAIM_003:		-> [FAILED]
[CLAIM_003:	OPTIMIZED	-> [FAILED]
[CLAIM_003:	COMPLIANT	-> [FAILED]
[CLAIM_004:	SCALABLE	-> [FAILED]

Codebase Realities

Refuting four inherited claims that **failed verification** at HEAD.

A silent routing failure mispriced the default model as Sonnet




Model Request	Resolved Route	Billed	Truth
claude-opus-5	FALLBACK to claude-sonnet-5	\$3 / \$15	
claude-opus-5[1m]	FALLBACK to claude-sonnet-5	\$3 / \$15	
claude-mythos-5	FALLBACK to claude-sonnet-5	\$3 / \$15	
			\$5 / \$25

Result: A flat 40% undercount on input and output for every session on the default model.


Identical historical pricing provided perfect camouflage

The Pricing Mask



Because the new model shares identical pricing with the old model, the top-line total costs didn't shift in reports.

The Alternative Universe



If the new model had different pricing, the fallback math would have produced an obviously incorrect total, triggering immediate alarm.

The test suite pinned the endpoints but ignored the route

The Pricing Table

Assertions verify every entry in `models.vocab.json` has correct values.

The Fallback

Assertions verify unknown models successfully route to Sonnet.

The Missing Route

No test asserted that a supported model actually avoids the fallback entirely.

Systemic Recurrence: This is the exact failure mode of `claude-fable-5` (70% underbilling). The previous fix added canaries ONLY for fable, leaving the class unclosed.

Closing the class requires a guardrail that can be deliberately broken

The Fix

```
const dynamicCanary = deriveExpectedTiers(fullIds);
dynamicCanary.forEach(tier => {
  assert.strictEqual(getPricing(), tier.expectedValue);
});
```

A new dynamic canary derives expected per-MTok tiers from each fullIds family and explicitly asserts getPricing() returns exactly that tier.


The Proof

FATAL ERROR: Expected tier \$10.00/MTok, but got \$0.00/MTok for model 'claude-opus-5'.
ASSERTION FAILED: Expected tier value mismatch.

A guardrail that cannot fail is theater.
ASSERTION FAILED:
at getPricing()...

Verification:
Removing the
'claude-opus-5'
pricing entry
now explicitly
triggers this
test failure,
proving the bug
class is
permanently
closed.

The illusion of a simple version bump to CC 2.1.220



Uncovering the unstamped manual cascade

114

Skill Floors

Enforced by ``SKILL_COMPAT_STRICT=1``.
A guaranteed hard CI failure until
moved. (83 pure floors, 31 with real
dependencies).

7

CI Pins

Installer workflows pinned at `<floor>`.
Would have kept quietly installing
below-floor CC builds in testing
environments.

19

Prose Sites



Source and documentation statements
(`'our floor is...'`) rendered instantly
false by the bump.

2

.latest Fixture Coupling

2 files, ~20 assertions requiring
identical movement to pass changelog
recovery tests.

Why a blanket find-and-replace would falsify historical records

Class A (Actionable Claims)	Class B (Historical Records)
Context: Statements about current state.	Context: Release headings, matrix rows, skill-eval logs.
We floor at 2.1.206.	CC 2.1.206 – external-worktree confirmation.
Action Taken: Rewritten (19 sites). 	Action Taken: Preserved. 

Execution Strategy: The regex sweep was explicitly gated on floor-claim phrasing appearing on the same line, reviewed prior to application to protect historical integrity.

Refuting inherited claims that failed verification at HEAD

Redline diagnostic ledger

Inherited Claim	True Status at HEAD
Nested spawning is off (Pattern 9 fails).	Restored to depth-3 in `2.1.219`. Documenting this would create false limitations.
500-line SKILL.md limit is unenforced.	Enforced at 520 lines. Authoring docs (500) and gate (520) disagree, but passing skills are legitimate.
`manifests/ork.json` derives the floor.	Has no `CC` version field; is not a stamp target.
`engines.claudeCode` guarantees fix.	`package.json` lacks this field. Truth lives strictly in `cc-support.json`.

Conscious Debt: Issues deliberately left undone

Compatibility Dedup

83/114 skill floors could be deleted entirely, but this requires changing test semantics ("must not declare" vs 'must match'). Opted for automated stamping instead.

Opus 4.8 Default Prose

~8 files still describe 4.8 as default behavior. Updating requires complex prose rewriting (Opus 5 effort ladders, prompt-cache minimums) – designated out of scope.

Doc Discrepancies

CONTRIBUTING-SKILLS.md states a 500-line limit while the system enforces 520. Identified but not reconciled.

DirectoryAdded Hook

Not wired. Requires verifying if the plugin validator rejects it, avoiding past MessageDisplay regressions.

Architectural takeaways for system reliability

