SG-6 — Moduli Stabilization (GUT Gate 6): Per-Gate Closure-Attack Dossier — Fable_Version rendered package. Rendered from `DOSSIER_SG6_MODULI_STABILIZATION_CLOSURE_ATTACK.md`; frozen technical content unchanged by rendering.

SG-6 — Moduli Stabilization (GUT Gate 6): Per-Gate Closure-Attack Dossier

What this is. The closure-attack packet for **SG-6 (Moduli stabilization / Gate 6)** on the **frozen 13D K_6 branch ONLY**. It recaps how our geometry controls the downstream-used moduli (concise; the full derivation lives on the published site), verifies where the status actually stands, names every open residual, and lays out the attack plan to push the gate further. It is **not** a rival comparison (that is `BATTLE_GATES/`), **not** the one-line status ledger, **not** a reprint of the manuscript.

Binding discipline (carry verbatim). PROMOTIONS:0. Frozen branch `dcc66f1b2685` / `a5b1e6f9d951` READ-ONLY. Honest throughout: the gate certifies **admissibility-restriction + LOCAL stability at the chamber-center witness + phenomenological sufficiency** for Gates 1–10 — **global stabilization is an explicit NON-claim** (F.9.4 / F.10.1); the **positive-definite moduli-mass Hessian is DIAGNOSTIC-ONLY** (F.9.5, the named most-natural attack surface); witnesses fix the moduli **WITHIN the already-selected chamber** (given-E, given-the-selected-geometry). The acute soft spot is **$\theta_{H^\star} \approx 2.46\times10^{-14}$** (the Hosotani phase), **READ from the chamber one-loop V_{Hos} minimum**, carrying **~85% of the electroweak hierarchy** via $v_{EW} = \theta_{H^\star}/(2\pi R_\gamma)$ — a value extracted from a minimum, not independently derived. Closure paths must be **non-target-loaded** (the κ^3/π kill-test: a proposed axiom counts only if it would be written WITHOUT knowing the target). **given-E ≠ derivation of E.** AXIOM-CLOSED ≠ proven; selection ≠ derivation; dissolved ≠ solved. **The broader four-term V(σ) discharge lives only in downstream Paper IV (EXTERNAL; its own caveats) and changes NOTHING in this gate's certificate.**

0. Header & Verdict

Field	Value
Gate id	SG-6 — Moduli stabilization (the Weyl-rigid chamber-center witness)
GUT manuscript gate	Gate 6 (§6.6; narrative §5.5; certificate `certificates/G06_stabilization/` ; authority **Appendix F**)
Status label (binding)	**PARTIAL**
Manuscript card status	*Claimed certificate pass under declared admissibility and moduli-control assumptions* (the manuscript's own conditional phrasing; PARTIAL is the honest scoped-GUT roll-up of it)
Frozen hashes it rides	Branch `dcc66f1b2685` / manifest meta `a5b1e6f9d951` . Moduli primitives: three Cartan radii of K_6 R1.2 `634438ce0776` / `2381d472c62e` / `0e8b8dba2cf0` ; chamber-center $\vec{u} = (1,1,1)$ (A1.2); modular fixed point $\tau = \omega$ `03b30a9c931a` ; Higgs winding $n_H = 1$ + cycle γ (A1.12); finite chamber determinant $\eta_{BK} = 0.009721281516312$ `84e94518d3f5` ; RG-transport rule `f531205a9159` ; comparison scale M_Z `a6852c7a6b00` ; threshold spectrum $(\delta_1, \delta_2, \delta_3) = (+4.8424, -3.1112, -1.7313)$; species normalizations $N_{u,d,e}$ `20dc4e0b8220` ; chamber angle θ_F `1ff57f48d45a` ; spin-$\mathbb{C}$ index / family count $\chi(K_6, E) = -3$ (R1.4) `0fd19c9ae0c1` . **θ_H⋆ (Hosotani minimum) and M_R: no SG-6 hash — read/uncomputed.**
Target anchor(s)	**No new anchor.** Witness ledger terminates on **spectrum-E** (given-E, AXIOM-CLOSED) and the **$\hbar$-footing only** (μ_cell, SCHEME-ANCHORED/BLOCKED); **global stabilization = explicit NON-claim** (terminates on no number). The one acute open leg, **θ_H⋆** (read from the V_Hos minimum, ~85% of the EW hierarchy), must reproduce **v_EW = 246 GeV** — here **measured-but-irreducible**; θ_H⋆-as-derived stays **OPEN**. Hardening move = the **θ_H⋆ / Hessian derivation**, not a new datum. See §A target-anchor block. PROMOTIONS: 0.

0.1 Abstract — established / open / what would close it

Established (given-E, given the selected & frozen chamber). Gate 6 is the *curve-not-a-number* discipline made operational (§6.6 "$y = R^2$" toy): **every modulus an output reads must carry a structural witness of a declared TYPE** — Weyl-rigid chamber for K_6's three Cartan shape moduli (chamber-center $\vec{u} = (1,1,1)$), modular fixed point $\tau = \omega$ for the F^+ Cartan-torus, integer winding $n_H = 1$ for the Wilson-line Higgs, discrete $\mathbb{Z}_2/\mathbb{Z}_6$ topology for the orbifold/identification, plus boundary-condition pins for the S^2/S^1_Y radii and the threshold spectrum — **not a scalar potential whose minimum was placed at the answer.** The witness-type **predicate** tests the *type* of the witness, so "a potential we minimized at R_0" **fails by construction**. The deliverables are the **12-row witness ledger (F.3)**, a **no-tachyon check (F.4)**, a **witness-coverage lint**, and the five-concept claim-type ledger (F.9.1 / F.10.1) that explicitly separates what is proved from what is not. This is genuinely strong discipline and is the program's honest answer to the "hidden-knob" attack.

Open (the honest deductions). (1) **Global stabilization is an explicit NON-claim** (F.9.4 row 1 / F.10.1) — outside-chamber branches are *rejected by admissibility, not stabilized*. (2) The **positive-definite moduli-mass Hessian is DIAGNOSTIC-ONLY** (F.9.5); a precise positive-definite Hessian is "out-of-scope refinement" — and the **named most-natural attack surface** is F.9.7(2): *show the chamber-center witness is a saddle, not a minimum*. (3) The acute soft spot is $\theta_H\star \approx 2.46\times10^{-14}$ — the Hosotani-phase minimum that carries **~85% of the EW hierarchy** via $v_{\rm EW} = \theta_H^\star/(2\pi R_\gamma)$ — **read from a one-loop $V_{\rm Hos}$ minimum, not derived-small** (frozen $\eta_{BK} = 0.009721$ gives only

$\sqrt{\eta_{BK}}/2\pi \sim 10^{-2}$, ~12 orders short). (4) The witnesses fix the moduli **WITHIN the already-selected chamber** (given-E): the gate certifies the selected geometry is internally controlled, **not** that the witnesses determine the geometry. (5) Stabilization is a **SHARED-OPEN** hard problem — rivals (string/M/F/NCG/lattice) also struggle (Battle verdict: TIE/SHARED-OPEN, **not** a Fable win).

What would close it (the ladder). DERIVED-CLOSED would require either (a) the **four-term V(σ) discharge** turning the breathing/volume modulus from witness into a *derived* positive-Hessian minimum with no new knob, or (b) a **target-blind derivation of θ_H⋆** (the small Hosotani phase from chamber data, not read from V_{Hos}), or (c) a **positive-definite Hessian certificate** at $(1,1,1)$. The realistic ceiling per residual is **AXIOM-CLOSED**: name one explicit, target-blind posit (the modular-fixed-point uniqueness; the uniform operational cell value Δ_0 that pins the loop scale; the second-dimensionful-anchor irreducibility of v_{EW}). The honest expected outcome on the two genuine-physics fronts (shape-doublet Hessian, θ_H⋆) is **sharper-OPEN** — and on the shape doublet, on current Λ-free perturbative geometry, **a SADDLE** (a falsifier-grade negative, not a closure).

0.2 Website source-of-truth (link, don't duplicate)

The full construction, the witness ledger, and the claim-type discipline are the **published manuscript**, Paper I:

- **GUT.html §6.6** (Gate-6 card) — https://physics.magflowmeters.com/articles/GUT.html
- **§5.5** narrative module; the $y = R^2$ **curve-not-a-number** discipline; **Appendix F** (Stabilization — the binding authority): **F.1/F.2** witness structure, **F.3** the 12-row moduli ledger, **F.4** tachyon check, **F.5** sensitivity, **F.9** Stabilization-vs-Admissibility five-concept distinction (**F.9.4** NON-claim, **F.9.5** Hessian diagnostic, **F.9.7** falsification paths), **F.10** claim-type ledger + required-vocabulary substitution rule; **Appendix CR module CR6** reader companion. Same public article.
- Downstream cross-reference (does **not** touch this gate's certificate): the four-term $V(\sigma)$ discharge in **Paper IV, TOE.html** — https://physics.magflowmeters.com/articles/TOE.html (EXTERNAL; its own caveats). The θ_H⋆ / v_{EW} leg is the Gate-8 Hosotani object (§6.8 / Appendix H) on which SG-6's acute soft spot rides.

This dossier recaps only what is needed to attack; the site controls all common material.

---

1. How OUR geometry controls the moduli — the closure claim (concise)

> Full derivation: **GUT.html §6.6 + §5.5 + Appendix F (F.1–F.10) + CR6**. This section is the attack-grade recap, not the derivation.

1.1 The mechanism end-to-end

Gate 6 answers a single plain fact: **a predictive compactification cannot have its hidden shape and size drifting freely — every modulus any later gate reads must be pinned before a**

number is compared (§5.5; the $y = R^2$ toy, §6 "Curve-not-a-number"). The Fable mechanism is **not** "write a scalar potential, minimize it, read off the value." It is a **witness-type predicate**: each downstream-used modulus must carry a witness drawn from a declared set of *structural* mechanisms, and the predicate tests the **type** of the witness, so a tuned potential whose minimum was *placed* at the answer fails by construction (R1.6 forbids exactly this).

The load-bearing factors, named so a reader can attack each:

1. **The Weyl-rigid K_6 chamber** (the three Cartan shape moduli $\vec{u} = (u_1, u_2, u_3)$ of $K_6 = SU(3)/T^2$). Witness type = **admissibility / no-runaway on the spin-$\mathbb{C}$ bundle data**: off-chamber moduli fail the A0 admissibility check and the branch is *eliminated by the selector* (Appendix B1.6), not stabilized. The chamber is $\vec{u} \in [1/2, 3/2]^3$ with center pinned at $(1, 1, 1)$. *(Attack handle: the Weyl group S_3 permutes (u_1, u_2, u_3), so $(1, 1, 1)$ is the S_3-fixed point — automatically a critical point of any S_3-invariant functional, but the Hessian sign is the open object; see §3 R3.)*
2. **The modular fixed point $\tau = \omega = e^{2\pi i/3}$** (`03b30a9c931a`) for the F^+ Cartan-torus modulus. Witness type = **residual modular symmetry forces the order-three fixed point**; small perturbations off ω acquire a non-zero potential under the chamber's RG transport, restoring $\tau \to \omega$. *(Attack handle: this is the cleanest witness — a symmetry-protected fixed point, not a minimum read off a curve. But the manuscript's own F.2 wording ("produce a non-zero potential off the fixed point") leans on a potential it does not compute; the symmetry-uniqueness leg is the strengthening target, R7.)*
3. **Integer Wilson-line winding $n_H = 1$** on K_{gauge}, on the declared homology cycle γ. Witness type = **topological invariant** — $\delta n_H \notin \mathbb{Z}$ is forbidden; integer winding is the only allowed value. *(Attack handle: the winding count is protected, but the phase θ_H that the Wilson line settles at is a continuous modulus on the cycle — and that is θ_H⋆, the acute soft spot; see §3 R1.)*
4. **Discrete $\mathbb{Z}_2$ / $\mathbb{Z}_6$ topology.** The $S^1_Y/\mathbb{Z}_2$ orbifold action $y \mapsto -y$ and the global $\mathbb{Z}_6$ identification have **no continuous deformation parameter to drift** — fixed by topology, no witness-potential needed.
5. **Boundary-condition pins for the S^2/S^1_Y radii and the threshold spectrum.** R_{S^2} by KK / Wilson-line balance against the threshold-unification target M_U; $R_{S^1_Y}$ by the hypercharge gauge-coupling boundary condition at M_Z; the threshold vector $(\delta_1, \delta_2, \delta_3) = (+4.8424, -3.1112, -1.7313)$ from the KK tower. *(Attack handle: these are threshold-pinned, i.e. their witness is consistency with Gate 7 — they inherit Gate 7's open status, notably the δ-triple reproducibility gap; given-E, given the frozen scheme.)*
6. **The 12-row witness ledger (F.3) + no-tachyon check (F.4) + coverage lint.** Every modulus / Wilson-line / chamber coordinate / threshold spectrum that any Gate 1–10 reads is listed with its witness, pinned value, downstream consumer, and the effect of a small perturbation δ. The certificate JSON (F.6) reports `stabilized_object_count: 6` "potential/witness" classes covering the 12 ledger rows, `tachyon_check: no tachyon on the active branch`.

The pipeline is one discipline, not one computation:

```
every downstream-used modulus
  --witness-type predicate-->  {Weyl-rigid | modular fixed pt | integer winding | discrete topology | BC pin}
    --type test (NOT minimize-at-answer)-->  PASS (witnessed) / FAIL (tuned-or-floating)
      --12-row ledger + no-tachyon + coverage lint-->  Claimed certificate pass (given-E, in-chamber)
```

1.2 What "closed" means here, and its conditionality

"Closed" for SG-6 is **moduli-control by structural witness under declared assumptions** — strictly **not** global stabilization, **not** a proven positive-definite Hessian, **not** a derivation of the chamber. It is the conjunction of exactly three claim types (F.9.3 / F.10.3):

- **(i) Admissibility restriction** — non-chamber configurations rejected by $\mathcal{C}_{\text{admiss}}$ (Appendix C6). *Proven* (it is an $\oplus$-layer rulebook rule).
- **(ii) Local stability at the chamber-center witness $\vec{u} = (1,1,1)$** — small in-chamber perturbations return to the center. *Claimed* at tree-plus-one-loop within the declared scheme.
- **(iii) Phenomenological sufficiency** — the remaining moduli do not move any Gate 1–10 output beyond its published tolerance. *Claimed* under the Weyl-rigid chamber + freeze discipline.

It is conditional on:

- **given-E** — the SM chiral content and the family index $\chi = -3$ are supplied as input from SG-2/SG-3; the witnesses operate *on the selected bundle*, they do not derive it.
- **given-the-selected-geometry** — the witnesses fix the moduli **inside** the already-selected chamber. SG-6 certifies the selected geometry is internally controlled; it does **not** prove the witnesses determine the geometry. *(This is the load-bearing conditionality of the whole gate.)*

What the gate explicitly does NOT claim (F.9.4 / F.10.1, carried verbatim):

Concept (F.9.1)	Claimed?	Status
Admissibility restriction	**Yes**	proven (rulebook)
Local stability at $(1,1,1)$	**Yes**	claimed, tree+1-loop, declared scheme
Moduli mass generation (positive-definite Hessian of $V(\vec{u})$)	**Partial → Diagnostic only**	F.9.5; **the named attack surface (F.9.7-2)**
Global stabilization (no flat directions anywhere)	**NOT claimed**	F.9.4 row 1 explicit non-claim
Phenomenological sufficiency for Gates 1–10	**Yes**	claimed

So the precise statement of closure: **the gate certifies that every downstream-used modulus carries a structural witness (not a tuned potential), that small in-chamber perturbations are controlled at the chamber-center, and that the residual moduli do not spoil Gates 1–10 — given the SM content and the selected geometry. It does NOT prove a positive-definite Hessian, does NOT claim global stabilization, and the one modulus carrying the electroweak hierarchy (θ_H⋆) is READ from a minimum, not derived.**

---

2. Verify the status — is PARTIAL real?

This is the verification a skeptic would run. Each witness gets a grade — **hand-checkable** / **symbolic** / **machine-lane** — and an honest **reproduces?** flag.

2.1 The witness ledger

#	Witness	What it asserts	Grade	Reproduces?
W1	**Witness-type predicate** (R1.6: "no tuned potential whose minimum was chosen at the answer")	every downstream modulus carries a witness of declared TYPE; "minimize-at-answer" fails by construction	symbolic/audit	**Yes** — the predicate is a type-check, hand-auditable against F.3; it is the strongest, most defensible leg
W2	**12-row moduli ledger (F.3)**	every modulus/Wilson-line/chamber coord/threshold spectrum used by Gates 1–10 is listed with witness + pinned value + downstream consumer + δ-response	hand-checkable	**Yes** — the table is complete and self-consistent; coverage lint passes (no floating downstream modulus)
W3	**Weyl-rigid chamber + center $\vec{u} = (1,1,1)$**	admissibility rejects off-chamber; center is the S_3-fixed point	hand-checkable	**Yes for admissibility + critical-point**; the *Hessian sign* is NOT asserted here (it is the diagnostic)
W4	**Modular fixed point $\tau = \omega$**	residual modular symmetry forces the order-three fixed point	hand-checkable	**Yes (as a fixed point)**; the "non-zero potential off ω" claim is symbolic, not computed
W5	**Integer winding $n_H = 1$**	topological invariant; non-integer forbidden	hand-checkable	**Yes** — pure topology
W6	**Discrete $\mathbb{Z}_2/\mathbb{Z}_6$**	no continuous deformation parameter	hand-checkable	**Yes** — pure topology
W7	**No-tachyon check (F.4)**	no tachyonic mode on the active branch at the declared witnesses	symbolic	**Conditional** — asserted for the declared spectrum; a deformation away from a witness *could* produce one (then the certificate fails)
W8	**Phenomenological sufficiency**	residual moduli leave Gates 1–10 within tolerance	machine-lane	**AUDIT** — depends on the same per-gate harnesses (e.g. the Gate-7 threshold pipeline) that are themselves AUDIT
W9	**Positive-definite Hessian of $V(\vec{u})$**	$(1,1,1)$ is a minimum, not a saddle	machine-lane	**DIAGNOSTIC-ONLY (F.9.5)** — *not* computed in-gate; the named attack surface. **An out-of-gate Λ-free computation gives a SADDLE in the shape sector (see §3 R3).**

2.2 What reproduces, plainly

- **The witness-type discipline reproduces by inspection.** The predicate is a *type* test; a reader can walk F.3 row by row and confirm each witness is structural (admissibility / symmetry fixed-point / integer winding / discrete topology / BC pin) and not "a number we minimized to." This is the genuine, strong leg of the gate, and it is exactly the $y = R^2$ curve-not-a-number toy applied to every dial the certificates touch.
- **The discrete/topological witnesses reproduce by hand.** $n_H = 1$, the $\mathbb{Z}_2$ orbifold, the $\mathbb{Z}_6$ identification, the cycle γ — all are topological invariants with no continuous knob; nothing to reproduce numerically and nothing to drift.
- **The S_3-fixed-point structure reproduces by hand.** Because the Weyl group S_3 permutes (u_1, u_2, u_3), the center $(1, 1, 1)$ is automatically a **critical point** of any S_3-invariant functional ($\partial V = 0$ for free). This upgrades the manuscript's "fixed by fiat" wording to "fixed by symmetry **as a critical point**" — a genuine, free, target-blind structural fact (TEST 2).

2.3 What does NOT yet reproduce (honest gaps in the status)

- **The positive-definite Hessian does NOT reproduce — it is diagnostic-only, and the one Λ-free piece that CAN be computed gives a SADDLE.** F.9.5 explicitly declines the Hessian computation as "out-of-scope refinement." When the chamber Hessian is decomposed by symmetry into $\mathbf{3} = \mathbf{1} \oplus \mathbf{2}$ (TEST 2), the **shape doublet** (S_3-doublet, traceless) is computable Λ-free from the flag-manifold scalar curvature — and the result is $\partial^2 R/\partial\varepsilon^2|_{(1,1,1)} = -1 < 0$, **a saddle** (the structure-constant triangle term $-5/2$ overwhelms the diagonal $+2$). The **breathing/volume singlet** rides the uncomputed c_{loop} (the 13D FRG-2 wall). So the F.9.7(2) saddle attack surface is not merely *open* — on current perturbative geometry it **leans toward saddle**, rescuable only by a Casimir term whose net sign needs the absent admissible-rep multiplicity table (TEST_CASIMIR_RESCUE).
- **θ_H⋆ does NOT reproduce as a derivation — it is read from a minimum.** The Hosotani phase $\theta_H^\star \approx 2.46 \times 10^{-14}$, which carries ~85% of the EW hierarchy via $v_{\text{EW}} = \theta_H^\star/(2\pi R_\gamma)$, is **extracted from the one-loop V_{Hos} minimum**, not computed-small from chamber data. The frozen $\eta_{BK} = 0.009721$ gives only $\sqrt{\eta_{BK}}/2\pi \sim 10^{-2}$ — ~12 orders short. The lightness-of-v leg leans on a read modulus.
- **Global stabilization does NOT reproduce — it is not claimed.** F.9.4 row 1 says so explicitly; outside-chamber branches are *rejected by admissibility, not stabilized*. There is no all-loop positive-definite Hessian and no claim of "no flat directions anywhere."
- **Phenomenological sufficiency inherits downstream AUDIT.** The "residual moduli don't spoil Gates 1–10" claim is only as machine-real as the downstream harnesses it relies on — and the Gate-7 threshold pipeline (which consumes the SG-6 radii/spectrum) is itself AUDIT (δ-triple injected reals, `reproduce_all.py` absent). SG-6's sufficiency claim is therefore declared-not-independently-verified at the numerical level.

2.4 Why PARTIAL (not higher, not lower)

- **Not DERIVED-GIVEN-E** (the SG-2/SG-3 tier): SG-6 is not a rigid integer/representation recovery. It has a diagnostic-only Hessian (with a Λ-free saddle in the shape sector), a read-from-minimum θ_H⋆ carrying the hierarchy, an explicit global-stabilization non-claim, and downstream-inherited AUDIT — too many open residuals for the DERIVED tier.
- **Not OPEN/DECLARED-FROZEN:** the witness-type discipline is real, strong, and non-trivial; the discrete/topological witnesses and the S_3-critical-point structure are genuine target-blind facts; the claim-type ledger honestly fences what is and is not proved. The gate is well clear of "unfixed knobs."
- **PARTIAL is exactly right:** genuine structural moduli-control under declared assumptions, a diagnostic-only Hessian named as the most-natural attack surface, an explicit non-claim on global stabilization, and a read-from-minimum acute soft spot. This matches the SCOPED_GUT ledger SG-6 line verbatim.

---

3. The attack surface — what to attack (ranked by leverage)

Every open residual as a named object with its precise obstruction. **Leverage** = how much closing it moves the gate (and, for SG-6, how much of the electroweak hierarchy / how much of the certificate it touches).

3.1 The residual register

ID	Residual (named object)	Precise obstruction	Status	Leverage
R1	**θ_H⋆ ≈ 2.46×10⁻¹⁴ read from the V_{Hos} minimum**	The Hosotani phase carries **~85% of the EW hierarchy** ($v_{\mathrm{EW}} = \theta_H^\star/(2\pi R_\gamma)$) yet is **read from a one-loop minimum, not derived-small**; $\sqrt{\eta_{BK}}/2\pi \sim 10^{-2}$ is ~12 orders short. The lightness-of-v leg leans on a read modulus.	OPEN (acute soft spot)	**HIGHEST** — it carries the bulk of the hierarchy and is the gate's single most attackable physics object
R2	**Global stabilization is an explicit NON-claim**	F.9.4 row 1 / F.10.1: outside-chamber configurations are *rejected by admissibility, not stabilized*; no all-loop positive-definite Hessian; no "no flat directions anywhere."	DISCLOSED (explicit non-claim)	**MEDIUM** — honesty/claim-boundary; the SHARED-OPEN hard problem; closing it is a known no-go-adjacent program
R3	**Positive-definite Hessian is DIAGNOSTIC-ONLY (the named attack surface)**	F.9.5 / F.9.7(2): "show the chamber-center witness is a saddle, not a minimum." The Λ-free shape-doublet Hessian = **−1 (SADDLE)** on perturbative geometry; the breathing singlet rides the uncomputed c_{loop}.	OPEN — **saddle-leaning**	**HIGH** — this is the explicitly-named falsification path; the shape sector is *already* computed to a saddle
R4	**Witnesses fix moduli WITHIN the selected chamber (given-E)**	The gate certifies internal control of the *selected* geometry; it does not prove the witnesses *determine* the geometry. Selection ≠ derivation; given-E ≠ derivation of E.	OPEN (inherited from SG-1/SG-3)	**MEDIUM** — conditions every SG-6 claim; shared with SG-1
R5	**c_{loop} / breathing-mode singlet = the 13D FRG-2 β-vector wall**	The volume/breathing direction's Hessian sign rides the uncomputed c_{loop} (the $e^{-6\sigma}$ coefficient = a_6 of the σ-fluctuation determinant). c_{loop} is `BLOCKED_MISSING_FRG_MATCHING`; even a finite cell-sum leaves one log-scheme factor μ_{cell} free.	OPEN / BLOCKED_INPUTS	**MEDIUM-HIGH** — shared deep wall with Gap-04 / hierarchy / a6; closing it decides the singlet Hessian sign
R6	**Casimir net sign needs the absent admissible-rep multiplicity table**	The shape-doublet saddle can be rescued by a graded Casimir term; per-sector sign is LOCKED (bosons stabilize, fermions destabilize), but the **net sign** needs the multiplicity table (ABSENT) and a magnitude $\mu_{\mathrm{cell}} \cdot q \gtrsim 0.5$.	BLOCKED_INPUTS / NEEDS-TUNING	**MEDIUM** — the only route that could flip R3 from saddle to minimum; gated on a missing owner artifact
R7	**$\tau = \omega$ "non-zero potential off the fixed point" leans on an uncomputed potential**	F.2 asserts perturbations off ω "produce a non-zero potential," but the potential is not computed; the leg should rest on **modular-symmetry uniqueness** of ω, which is not proven in-gate.	OPEN	**LOW-MEDIUM** — the cleanest witness; strengthening it from "potential" to "symmetry-protected" is a contained group-theory win

ID	Residual (named object)	Precise obstruction	Status	Leverage
R8	**No-tachyon check is conditional on the declared spectrum**	F.4: no tachyon "on the active branch"; a deformation away from a witness could produce one. The check is asserted for the frozen spectrum, not proven stable under the deformation classes the gate exposes.	DISCLOSED (conditional)	**LOW** — disclosed; the deformation classes are bounded by admissibility
R9	**Phenomenological sufficiency inherits downstream AUDIT (Gate-7 radii/spectrum)**	The "residual moduli don't spoil Gates 1–10" claim relies on the Gate-7 threshold pipeline (which consumes the SG-6 radii + threshold spectrum), which is itself AUDIT (δ-triple injected, `reproduce_all.py` absent).	AUDIT (inherited)	**LOW-MEDIUM** — required for the sufficiency leg to be machine-real; the artifact is owned by SG-7

3.2 Leverage ranking (attack order)

1. **R1** (θ_H⋆ read from a minimum) — carries ~85% of the hierarchy; the single highest-value physics target.
2. **R3** (the diagnostic Hessian — the named saddle attack surface) — already computed to a saddle in the shape sector; the explicit F.9.7 falsification path.
3. **R5** (the c_{loop} / breathing-mode singlet wall) — decides the singlet Hessian sign; the shared deep 13D FRG-2 wall.
4. **R6** (the Casimir rescue) — the only route that could flip R3; gated on the missing multiplicity table.
5. **R4** (given-E / within-chamber) — conditions everything; honesty residual shared with SG-1/SG-3.
6. **R2 / R7 / R8 / R9** — disclosed non-claim / cleanest-witness strengthening / conditional tachyon / inherited AUDIT.

> **The cardinal honest point.** R1 and R5 are where the "85% hierarchy" and the "singlet Hessian" live, and **both reduce to the same single object: μ_{cell}, the spectral value of the uniform operational cell Δ_0** (the log-scheme / cell-scale factor of the 6-D a_6 determinant). The decisive, already-run firewall verdict (CLOSURE_CAMPAIGN §3) is that **μ_{cell} has no v-independent readout** — its only available anchor is $\partial_\sigma V = 0$, which **IS** the hierarchy — so anchoring μ_{cell} there to "predict" v is **circular by construction (the κ^3/π signature)**. Any closure path that pins μ_{cell} from the EW hierarchy has **RELOCATED**, not removed, the input. The κ^3/π kill-test applies in full force here. R2/R7/R8/R9 are the *non-physics-flip* residuals (claim-boundary / symmetry-strengthening / conditional-check / inherited reproducibility); closing them improves honesty and machine-reality but **does not derive the hierarchy or flip the Hessian**.

---

4. THE ATTACK PLAN — closure paths (the core)

For each residual: the **technique** (closure playbook T1 axiom-floor / T4 no-go / T5 firewall / T6 all-operator-conditional / T7 eliminative / T10 selector), the **named axiom it could reduce to** (stated so it would be written WITHOUT the target value — the κ^3/π kill-test), the **specialist target** (theorem to

hand off) or the **owner artifact** (CSV/ruling/computation) needed, the **math to attempt**, and the **success ladder** (DERIVED-CLOSED rare → AXIOM-CLOSED likely → sharper-OPEN → REFUTED).

4.1 R1 — θ_H⋆ ≈ 2.46×10⁻¹⁴ read from the $V_{\rm Hos}$ minimum (the acute soft spot)

Why first. This single modulus carries **~85% of the electroweak hierarchy**. If θ_H⋆ could be *derived* target-blind from chamber data instead of read off the one-loop Hosotani potential, the lightness-of-v leg would stop leaning on a read number and SG-6 would deliver the hierarchy's dominant factor. If it provably cannot, the gate's honest standing tightens to "the hierarchy is a second irreducible anchor."

Technique: T5 (no-target-loading firewall) + T7 (eliminative) — is θ_H⋆ a derivation or a relocation? This is exactly the SCALE firewall, pointed at the Hosotani phase.

The two routes, and which the firewall kills:

Route	Statement	Firewall verdict
Path A — exact v via $\mu_{\rm cell}$	derive the loop scale $\mu_{\rm cell}$ from $\{\hbar, M_U/M_{\rm Pl}, \text{frozen geometry}\}$ at a fixed UV reference, then read $\theta_H^\star$ off $V_{\rm Hos}$	**CIRCULAR / RELOCATION** — $\mu_{\rm cell}$ must be a *mass*; the only σ-independent mass is $M_{\rm Pl}$ (an anchor), the rest is dimensionless shape; turning shape into a mass needs an inverse length, and the only one available is $R_{K_6}(\sigma)$ — σ-carrying. Buckingham-π: a *second* mass cannot be built from $\{M_{\rm Pl}, \hbar, \text{dimensionless geometry}\}$. $\mu_{\rm cell} \to v$ is invertible-by-construction (exponential) ⇒ $\mu_{\rm cell}$ IS the knob.
Path B — structural "huge" via transmutation	bound $\theta_H^\star$ small by a dimensional-transmutation exponent $t_* = 2\pi/(b\alpha)$	**WRONG MECHANISM** — θ_H⋆ is the **location of a stationary point of a periodic potential**, NOT a running coupling, so there is *no* transmutation exponent to bound. The field-count bound governs only $\ln(M_{\rm Pl}/M_U)$ (the frozen ~15%), never the ~85% in θ_H⋆. (Bonus genuine sub-result: Gate-8 Hosotani protection *does* remove the Higgs-mass quadratic destabilization, but only to ~10^{14} GeV — 12 orders short.)

Named axiom it could reduce to (κ³/π-clean). The honest reduction is **not** a closure but a precise statement of what R1 reduces to once the two routes are killed:

> **AXIOM-VEW-SECOND-ANCHOR** (target-blind form): *"the electroweak scale $v_{\rm EW}$ is a second independent dimensionful anchor on the same footing as $M_{\rm Pl}$, irreducible in this geometry because no second mass can be built from $\{M_{\rm Pl}, \hbar, \text{dimensionless frozen geometry}\}$ without a σ-carrying length (Buckingham-π), and because the carrier θ_H⋆ is a periodic-minimum location, not a transmutation exponent."*

This is writable with **no** value of v or θ_H⋆ in sight — it is a *dimensional + mechanistic* statement about what kind of object the hierarchy is. **It PASSES the κ³/π kill-test** (it carries no reverse-engineered coefficient). It does **not** derive θ_H⋆; it names *why* θ_H⋆ must be read, not computed, in this structure.

Specialist target (to hand off). A **target-blind θ_H⋆ derivation theorem**: "Given only the frozen chamber data (cycle γ, $n_H = 1$, η_{BK}, R_γ) and a UV reference frozen *before and independently of* $v_{\rm obs}$,

compute $\theta_H^\star$ and check whether it returns $\approx 2.46 \times 10^{-14}$ with no v in the inputs." If a specialist can produce such a derivation, R1 closes (DERIVED) and the hierarchy's dominant factor is delivered; if the derivation provably requires $\partial_\sigma V = 0$ (the hierarchy itself), R1 is confirmed RELOCATION and AXIOM-VEW-SECOND-ANCHOR stands.

Math to attempt. (1) Write $V_{\text{Hos}}(\theta_H)$ from the frozen cycle/winding/η_{BK} data. (2) Locate its minimum θ_H⋆ **without** consulting v_{obs}. (3) Check whether the minimum lands at $\sim 10^{-14}$ from chamber data alone, or only after a scale is pinned by v. (4) Run the v-independence audit: does the threshold-weight swing across the σ-band ($\sim 34\times$, the documented forbidden count-only "35" stationarity) confirm μ_{cell} is the knob?

Success ladder. - DERIVED-CLOSED: θ_H⋆ computed target-blind from chamber data, returning $\sim 10^{-14}$ with no v input → the hierarchy's dominant factor delivered. **Unlikely** — the dimensional near-no-go (Buckingham-π) blocks a second mass from $\{M_{\text{Pl}}, \hbar, \text{shape}\}$. - **AXIOM-CLOSED (likely):** AXIOM-VEW-SECOND-ANCHOR named; θ_H⋆/v_{EW} conceded as a second irreducible dimensionful anchor with a *principled* (dimensional + mechanistic) reason, not "we failed to reduce it." This is the honest ceiling and is essentially where the SCALE frontier already stands. - sharper-OPEN: the derivation attempt returns fail-closed on μ_{cell} (the FRG-2 wall) without deciding circularity. - REFUTED: a chamber-only computation gives a θ_H⋆ that *misses* 10^{-14} → the lightness-of-v leg is a genuine prediction-vs-data failure.

Honest disposition: AXIOM-CLOSED at AXIOM-VEW-SECOND-ANCHOR; θ_H⋆ as "derived-small" stays OPEN/RELOCATION. The firewall has already confirmed Path A circular and Path B wrong-mechanism; the genuine banked sub-result is the Higgs-mass quadratic protection (proven, finite, 12-orders-short). Saying precisely "the hierarchy is a second anchor, here is the structural reason it must be" is the closure — **not** a higher claim.

4.2 R3 — The diagnostic Hessian: the named saddle attack surface (F.9.7-2)

Why second. F.9.5 grades the positive-definite moduli-mass Hessian DIAGNOSTIC-ONLY and F.9.7(2) names the single most-natural falsifier: *show the chamber-center witness is a saddle, not a minimum.* This is the explicitly-named attack surface, and it is **already partially executed against the gate.**

Technique: T7 (eliminative) by symmetry decomposition — split the Hessian, compute the computable piece. The Weyl group S_3 permutes (u_1, u_2, u_3), so the 3-dim perturbation space splits exactly into $\mathbf{3} = \mathbf{1} \oplus \mathbf{2}$ (TEST 2):

Mode	Irrep	Stability controlled by	Status
breathing / volume ($u_1 = u_2 = u_3$)	singlet $\mathbf{1}$	$V''(\sigma)$, the $e^{-6\sigma}$ term → c_{loop}	**= the FRG-2 wall (R5)**
shape ($u_1 - u_2$, ×2 degenerate)	doublet $\mathbf{2}$	flag-manifold curvature / Casimir	**Λ-free, c_{loop}-free, computable**

$(1,1,1)$ is a minimum **iff both eigenvalues > 0**, and they decouple by symmetry.

The Λ-free computation, already run (TEST_SHAPE_DOUBLET_STABILITY). The flag-manifold scalar curvature $R(u) = \sum_i 1/u_i - \frac{1}{2}T(u)$, $T(u) = \sum_k u_k/(u_i u_j)$, along the doublet ray

$u = (1+\varepsilon, 1-\varepsilon, 1)$ gives: - diagonal $\sum 1/u_i$: $\to +2$ to the ε^2 coefficient (convexity of $1/u$), - structure-constant triangle T: $\to -5/2$, - net $\partial^2 R/\partial\varepsilon^2|_{(1,1,1)} = +2 - 5/2 = \mathbf{-1} < 0 \to$ **SADDLE.**

The boundary (Gibbons–Hawking) term is small and mildly destabilizing; the Wilson-line term is exactly zero in the doublet direction. **On purely perturbative geometry, the shape doublet is a SADDLE** — the structure-constant coupling between the three root planes makes the equal-scale flag metric unstable to shape deformation. (This is the well-documented fact that $SU(3)/T^2$ carries three invariant Einstein metrics and the normal one is not the shape-sector extremizer.)

The only rescue, and its kill-test (TEST_CASIMIR_RESCUE). A graded Casimir / KK vacuum-energy term can flip the doublet sign. The per-sector sign is **LOCKED** (bosons stabilize, fermions destabilize, by convexity $g''(0) = 2(w_1^2 + w_2^2) > 0$ and $f'(m^2) > 0$, scheme-independent). The **net** sign is $\mathbf{sgn(boson\text{-}surplus - fermion\text{-}surplus)}$ over admissible (p, q) at $(1,1,1)$ — which **needs the admissible-rep multiplicity table (ABSENT from the corpus)** and a magnitude $\mu_{\text{cell}} \cdot q \gtrsim 0.5$. The frozen heat-kernel ledger (G.3.2a) *weakly* favors boson dominance (adjoint gauge+ghost $-2.49/-4.02$ vs matter $+0.79/+0.92$, ~3:1) → conjectured rescue — **but the opposing frozen fact is $\chi = -3$ (a low-lying fermionic chiral surplus that destabilizes)**, and banking boson-dominance without the multiplicity table is NEEDS-TUNING.

Named axiom it could reduce to (κ³/π-clean). Not a closure, but the precise reduction of the doublet question:

> **AXIOM-SHAPE-DOUBLET-CASIMIR** (target-blind form): *"the shape-doublet Hessian at the chamber center is $-1_{\text{curvature}} + q\,\mu_{\text{cell}}\,(+2_{\text{convexity}})$, with the convexity factor $+2$ and the per-sector sign locked by geometry; the net sign is fixed by the graded boson–fermion surplus of the admissible-rep spectrum."* Writable with **no** Λ and **no** target value — it references only the root-system curvature and the graded multiplicity. PASSES the kill-test as a *statement*; the test is then whether the multiplicity table *yields* a net positive.

Specialist target / owner artifact. (1) **Owner artifact:** the **admissible-rep multiplicity table** at $(1,1,1)$ + the regularized zeta (from `certificates/` + `reproduce_all.py` / the KK eigenvalue file R0.7) — this is the single missing input that decides the net Casimir sign. (2) **Specialist target:** "Given the multiplicity table, compute $\mathbf{sgn(boson\text{-}surplus - fermion\text{-}surplus)}$ and the magnitude $\mu_{\text{cell}} \cdot q$, and report whether the doublet net Hessian is > 0 (rescue) or ≤ 0 (saddle persists)."

Math to attempt. (1) The doublet curvature -1 is **done** (banked, target-blind). (2) Mount the multiplicity table; form the graded doublet-weighted sum $\sum (-1)^F d\,n\,(w_1^2 + w_2^2)$. (3) Decide the net sign. (4) Combine the singlet (R5) and doublet results for the full Hessian signature.

Success ladder. - DERIVED-CLOSED: both eigenvalues (singlet via R5, doublet via the multiplicity table) proven > 0 target-blind → the Hessian is positive-definite, F.9.7(2) defeated, the diagnostic upgrades to a derived minimum. **Gated on both the multiplicity table AND the FRG-2 c_{loop} — unlikely in one step.** - AXIOM-CLOSED: AXIOM-SHAPE-DOUBLET-CASIMIR named and the multiplicity table *yields* a net positive doublet → the shape sector reduced to one symmetry/spectral posit. **Best realistic outcome for the doublet.** - **sharper-OPEN / saddle-leaning (current standing):** doublet curvature is a saddle (-1); the rescue needs the absent table; the singlet rides the FRG-2 wall. The

honest current state is "saddle on perturbative geometry, conditional rescue pending owner artifacts." - REFUTED: multiplicity table computed and the net doublet stays negative → the chamber-center is a **saddle**, F.9.7(2) succeeds, and the moduli-mass-generation row of F.9.2 downgrades (per F.10.4-2, to *Open / not claimed*).

Honest disposition: sharper-OPEN, saddle-leaning. The named falsification path is the *closest to firing* of any SG-6 residual: the one Λ-free piece that can be computed is a saddle, and the rescue is gated on a missing multiplicity table and the μ_{cell} value (shared with R1/R5). **Do not bank the boson-dominance rescue without the table.**

4.3 R5 — c_{loop} / breathing-mode singlet = the 13D FRG-2 β-vector wall

The obstruction. The breathing/volume singlet eigenvalue rides $V''(\sigma)$ at the $e^{-6\sigma}$ term — i.e. c_{loop}, the coefficient of the σ-fluctuation one-loop determinant, equivalently the a_6 **Seeley–DeWitt heat-kernel coefficient** on the 6-manifold K_6. c_{loop} is `BLOCKED_MISSING_FRG_MATCHING` — the uncomputed 13D FRG-2 β-vector, the deep core shared with Gap-04 and a6.

Technique: T7 (eliminative) — is the wall the continuum or the scale? Replace the continuum mode integral by a **finite cell-sum** under the uniform operational cell law $N \leq B/\Delta_0$ (TEST_CLOOP_CELLSUM_ATTACK).

What the cell-sum genuinely buys, and what it does not. The cell-sum **does** kill the UV ($a \to 0$) divergence — the continuum-as-source-of-divergence assumption is genuinely false, and dropping it is banked (ROOT-1). But finiteness does **not** imply uniqueness: a finite supertrace of a log-running 6-D determinant still carries a **log-scheme degree of freedom** μ_{cell}. Three residual ambiguities survive: (3a) the cutoff is a mode-**COUNT**, but a_6 needs a mode-**SCALE**; (3b) the η-dependence the ledger flagged IS the scheme knob; (3c) supertrace ordering / sign is not fixed by a count. So:

$$c_{\text{loop}} = f(\text{frozen } K_6 \text{ spectrum},\ \chi = -3,\ e^{-6\sigma} \text{ scaling}) \times [\text{one log-scheme factor fixed by } \mu_{\text{cell}}].$$

Named axiom it could reduce to (κ³/π-clean). This is exactly the GRANULARITY root's residue:

> **AXIOM-UNIFORM-CELL-VALUE** (target-blind form): *"there exists a single uniform operational cell $\Delta_0 > 0$ on the operational distinguishability metric (existence is the GRANULARITY root, `AXIOM-COSTFLOOR`; the spectral VALUE of Δ_0 is a FLOOR-VALUES-RESIDUE on the same footing as $\hbar, k_B$), and μ_{cell} is its spectral value at the K_6 scale."* Writable with no Λ and no hierarchy value. PASSES the kill-test as an *existence + value-residue* statement.

The κ³/π pivot (the firewall, already run). The single anchor that would pin μ_{cell} is **one measurement fixing Δ_0 in spectral units at the K_6 scale** — and the *only* such anchor available is the moduli-stabilization point $\partial_\sigma V = 0$. **But $\partial_\sigma V = 0$ IS the electroweak hierarchy**, so anchoring μ_{cell} there and then "predicting" the hierarchy is **target-loading — explicitly FORBIDDEN** (the CLOSURE_CAMPAIGN firewall designated $\partial_\sigma V = 0$ the forbidden anchor; the earlier TEST_CLOOP suggestion to use it is superseded). The honest anchor is Δ_0's spectral value as an *independent* floor-value input, on the footing of $\hbar$ — which the corpus does **not** supply.

Specialist target / owner artifact. (1) **Owner artifact:** the KK eigenvalue file / `reproduce_all.py` (R0.7) that would pin μ_{cell} *if* an independent floor-value for Δ_0 existed. (2) **Specialist target (the deep one):** "Compute the 13D FRG-2 β-vector and hence $c_{\text{loop}} = \mathrm{tr}[a_6]$ of the σ-fluctuation determinant on the frozen K_6 branch, target-blind." This is the shared Gap-04 / hierarchy / a6 wall.

Success ladder. - DERIVED-CLOSED: the FRG-2 β-vector computed → c_{loop} sign decided → the singlet Hessian eigenvalue fixed. **Blocked on the deep 13D FRG wall; not reachable here.** - AXIOM-CLOSED: AXIOM-UNIFORM-CELL-VALUE named, with μ_{cell} conceded as a floor-value residue (no new knob, no target). **Honest ceiling; already the GRANULARITY root's standing.** - **sharper-OPEN / BLOCKED_INPUTS (current standing):** the cell-sum lowers the wall (finite, single-parameter) but does not remove it; the one residual μ_{cell} has no v-independent readout. - REFUTED: the FRG computation gives a definite c_{loop} with the wrong sign → the breathing singlet is a saddle → chamber-center is not a minimum.

Honest disposition: BLOCKED_INPUTS, reducible to AXIOM-UNIFORM-CELL-VALUE. The breathing-singlet Hessian is the same FRG-2 wall as the hierarchy; the cell-sum converts "uncomputed continuum determinant" into "finite cell-sum + one named floor-value anchor," which closes nothing but names the bottom precisely. **The $\partial_\sigma V = 0$ anchor is forbidden** (it is the hierarchy).

4.4 R6 — The Casimir net sign (the only route that could flip R3)

This is the dependent of R3, broken out because it is the single decision that flips saddle→minimum. Fully covered as the rescue analysis under §4.2: per-sector sign LOCKED (bosons +, fermions −, scheme-independent); **net sign = boson-surplus − fermion-surplus over the absent admissible-rep multiplicity table**; magnitude threshold $\mu_{\text{cell}} \cdot q \gtrsim 0.5$.

Technique: owner artifact (machine-lane) + T7. Mount the multiplicity table + regularized zeta; form the graded doublet-weighted sum; decide the net sign and magnitude.

The κ^3/π guard. The frozen heat-kernel ledger weakly favors boson dominance (~3:1), but the opposing frozen fact is $\chi = -3$. **Banking "boson-dominated rescue" without the table is NEEDS-TUNING** — the conjecture $q > 0$ is exactly the kind of plausible-but-unverified sign that the kill-test forbids until the table is computed. The magnitude carries μ_{cell} (shared with R1/R5), which has no v-independent readout.

Success ladder. BLOCKED_INPUTS until the multiplicity table is mounted → then **VERIFIED-RESCUE** (net > 0, $\mu_{\text{cell}} \cdot q \gtrsim 0.5$ → minimum) or **SADDLE-CONFIRMED** (net ≤ 0 → F.9.7(2) fires). **Highest value-per-effort *physics* item if the table exists** — it converts the named saddle attack surface from "leaning saddle" into a decided verdict with no new tuning.

Honest disposition: BLOCKED_INPUTS / NEEDS-TUNING. The rescue is conditionally available but unbankable without the absent multiplicity table; do not assert boson dominance as a result.

4.5 R2 — Global stabilization is an explicit NON-claim

This is a claim-boundary residual, not a physics flip. Global stabilization (no flat directions anywhere in the active-branch moduli space) is **explicitly NOT claimed** (F.9.4 row 1 / F.10.1) — outside-chamber configurations are *rejected by admissibility, not stabilized*. The task is not to "solve" it but to keep the boundary honest and to state what *would* close it.

Technique: T6 (all-operator-conditional) via the manuscript's own F.10.2 vocabulary rule. The required-vocabulary substitution rule already bans "full stabilization" / "no flat directions" / "all moduli stabilized" in body text without the chamber-restriction qualifier. Compliance is mechanical.

Named principle (κ³/π-clean): AXIOM-CHAMBER-RESTRICTION — *"outside-chamber configurations are rejected by the admissibility rulebook $\mathcal{C}_{\text{admiss}}$ (an ⊕-layer rule), not dynamically stabilized; the stabilization claim is scoped to the moduli used downstream by Gates 1–10."* This is the F.9.4/F.10.3 statement; it carries no target value and is the honest scope wall.

What would close it (and why it is SHARED-OPEN). Global stabilization is the textbook hard problem on which **all** rivals struggle (Battle verdict: TIE/SHARED-OPEN — *not* a Fable win). The downstream four-term $V(\sigma)$ discharge (Paper IV) is the only object that even attempts the volume modulus globally, and it is EXTERNAL with its own caveats and "promotions: zero" discipline; it changes **nothing** in this gate's certificate (the F.8 downstream-corpus pointer says so explicitly).

Success ladder. AXIOM-CLOSED is not applicable as a *flip* (no axiom closes global stabilization here). The endpoint is **DISCLOSED-CONSISTENT**: the non-claim is explicit, the vocabulary rule is enforced, and the SHARED-OPEN status is stated honestly. **The cleanest, lowest-risk honesty item.**

Honest disposition: DISCLOSED non-claim; keep it disclosed. Do not let any downstream cross-reference (Paper IV) be cited to upgrade this gate's certificate.

4.6 R4 — Witnesses fix moduli WITHIN the selected chamber (given-E)

The obstruction. The witnesses certify that the *selected* geometry is internally controlled; they do **not** prove the witnesses *determine* the geometry. The chamber, the bundle E, and the family index $\chi = -3$ are inputs (from SG-1/SG-3). Selection ≠ derivation; given-E ≠ derivation of E.

Technique: T1 (axiom-floor), shared with SG-1/SG-3. The honest reduction names the inheritance:

> **AXIOM-GIVEN-E-CHAMBER** (κ³/π-clean): *"the stabilization witnesses operate on the selected 13D K_6 branch and the supplied SM chiral content E; SG-6 certifies internal moduli-control of that selected geometry, not a cross-geometry determination."* Writable with no target value — it is a scope statement inherited from SG-1 (DECLARED-FROZEN, SHAPE selected-not-forced-absolute, ~9–10 injected reals beyond the 4 anchors) and SG-3 (DERIVED-GIVEN-E).

Specialist target. None unique to SG-6 — this is the SG-1 SHAPE-minimality / SG-3 given-E conditionality flowing downstream. The contained SG-6 strengthening is R7 (the modular-fixed-point uniqueness), which is the one witness that *could* be promoted from "selected" to "symmetry-forced."

Success ladder. AXIOM-CLOSED (AXIOM-GIVEN-E-CHAMBER named) is essentially already standing via the SG-1/SG-3 labels. DERIVED-CLOSED is not available at SG-6 (it would require closing SG-1's SHAPE-forcedness, which is SELECTED-not-forced-absolute). **Honest disposition: AXIOM-CLOSED at the inherited scope axiom; the within-chamber conditionality is the load-bearing caveat of the gate.**

4.7 R7 — $\tau = \omega$: strengthen "non-zero potential" to symmetry-uniqueness

The obstruction. F.2 asserts that perturbations off $\tau = \omega$ "produce a non-zero potential" restoring $\tau \to \omega$, but the potential is **not computed**. The leg should rest on **modular-symmetry uniqueness** of ω — a fixed point of an order-three modular symmetry is *symmetry-protected*, not tuned — which is stronger than "a potential we did not compute."

Technique: T1 (axiom-floor) — name the symmetry, prove the uniqueness.

> **AXIOM-MODULAR-FIXED-POINT** (κ³/π-clean): *"the F^+ Cartan-torus modulus sits at the order-three modular fixed point $\tau = \omega = e^{2\pi i/3}$, the unique modular-symmetric / Weyl-rigid point of the chamber."* Writable with no flavor or hierarchy number — a symmetry statement. **Stronger than "read from a minimum"** because an order-three fixed point is symmetry-protected.

Specialist target. Hand the moduli specialist: "Show $\tau = \omega$ is the **unique** fixed point of the F^+ Cartan-torus modular group, independent of any one-loop potential." A clean, bounded group-theory claim.

Math to attempt. Enumerate the fixed points of the chamber's modular group on the upper half-plane; confirm ω is the order-three fixed point and characterize uniqueness within the chamber.

Success ladder. AXIOM-CLOSED (AXIOM-MODULAR-FIXED-POINT named; symmetry-protected) is the realistic endpoint; DERIVED-CLOSED (uniqueness theorem proven) closes R7 outright and converts the τ-witness from "potential we did not compute" to "symmetry-forced point." sharper-OPEN if the fixed point is non-unique. **Honest disposition: AXIOM-CLOSED likely, DERIVED if uniqueness is proven; this is the one witness that can be genuinely strengthened by a contained computation.** (Note: this also strengthens SG-8's inherited $\tau = \omega$ soft spot.)

4.8 R8 — No-tachyon check is conditional on the declared spectrum

The obstruction. F.4 asserts no tachyonic mode "on the active branch" at the declared witnesses, but a deformation away from a witness could produce one (and then the certificate fails). The check is for the frozen spectrum, not proven stable under the deformation classes the gate exposes it to.

Technique: T6 (all-operator-conditional) within the admissibility-bounded deformation class. The deformations the gate must survive are bounded by admissibility (off-chamber ⇒ branch eliminated). The honest statement is conditional, and the F.5 sensitivity analysis already bounds the in-chamber responses.

Named principle (κ³/π-clean): AXIOM-NO-TACHYON-IN-CHAMBER — *"on the active branch at the declared witnesses, and under admissibility-bounded in-chamber deformations, no tachyonic mode appears."* Carries no target value; it is the F.4 statement scoped to the admissible deformation class.

Success ladder. DISCLOSED-CONSISTENT (already substantially in place via F.4 + F.5). The residual is a consistency sweep over the admissibility-bounded deformations, not a closure. (Note: R3's saddle finding is the *spatial* (Hessian) analog of the tachyon check — a saddle direction is a flat/negative mass direction in field space; R8 and R3 are linked, and a confirmed shape-doublet saddle would put pressure on R8 in that direction.)

Honest disposition: DISCLOSED-CONSISTENT, conditional on the declared spectrum; dependent on R3.

4.9 R9 — Phenomenological sufficiency inherits the downstream AUDIT (Gate-7 radii/spectrum)

The obstruction. The "residual moduli don't spoil Gates 1–10" claim relies on the downstream harnesses that consume the SG-6 radii and threshold spectrum — chiefly the **Gate-7 threshold pipeline**, which is itself AUDIT (the three δ are injected reals quoted to five sig figs but not regenerated; the G.3.2a row formulas do not reproduce the decimals; `reproduce_all.py` + the ledger CSVs are ABSENT). SG-6's sufficiency leg is only as machine-real as those harnesses.

Technique: owner artifact (machine-lane), not an axiom. A fail-closed reproducibility task, not a physics closure.

Owner artifact needed. (1) Mount the Gate-7 ledger CSVs (`appendix_F_heat_kernel_ledger.csv`, `appendix_F_threshold_outputs.csv`). (2) Run `reproduce_all.py` against the frozen radii + spectrum, target-blind. (3) Confirm the threshold vector + unification residual regenerate. (4) Confirm a small δ on the SG-6 radii moves the Gate-7 output within the F.3 δ-response column.

Math to attempt. None new — execution + verification. Mechanical and fail-closed: if a perturbation of an SG-6 radius moves a Gate 1–10 output beyond tolerance and is not bounded by F.3, phenomenological sufficiency fails (F.10.4-3) and the offending downstream gate downgrades.

Success ladder. BLOCKED_INPUTS until the Gate-7 CSVs/script are mounted → then **VERIFIED** (sufficiency machine-real) or **REFUTED** (a perturbation breaches tolerance → downgrade). **Value-per-effort high, but the artifact is owned by SG-7** — closing SG-7's δ-harness AUDIT closes this too.

Honest disposition: AUDIT, inherited from SG-7; closes when SG-7's reproduction harness is mounted.

4.10 Attack-plan roll-up

Residual	Technique	Named axiom (κ³/π-clean)	Specialist target / owner artifact	Realistic endpoint
R1 θ_H⋆ read from minimum	T5 + T7	AXIOM-VEW-SECOND-ANCHOR	target-blind θ_H⋆ derivation; v-independence audit	AXIOM-CLOSED (2nd anchor, principled); derived-small OPEN/RELOCATION
R3 diagnostic Hessian (saddle surface)	T7 (symmetry split)	AXIOM-SHAPE-DOUBLET-CASIMIR	admissible-rep multiplicity table + zeta	sharper-OPEN, **saddle-leaning** (-1 curvature)
R5 c_{loop} / breathing singlet	T7 (cell-sum)	AXIOM-UNIFORM-CELL-VALUE	13D FRG-2 β-vector; KK eigenvalue file	BLOCKED_INPUTS → AXIOM-CLOSED (floor-value residue)
R6 Casimir net sign (R3 flip)	owner artifact + T7	(uses AXIOM-SHAPE-DOUBLET-CASIMIR)	multiplicity table	BLOCKED_INPUTS / NEEDS-TUNING
R2 global-stabilization non-claim	T6 (vocab rule)	AXIOM-CHAMBER-RESTRICTION	F.10.2 vocabulary sweep	DISCLOSED-CONSISTENT (SHARED-OPEN)
R4 given-E / within-chamber	T1	AXIOM-GIVEN-E-CHAMBER	(inherited SG-1/SG-3)	AXIOM-CLOSED (scope axiom)
R7 $\tau = \omega$ strengthening	T1	AXIOM-MODULAR-FIXED-POINT	modular-fixed-point uniqueness	AXIOM-CLOSED → DERIVED if unique
R8 no-tachyon conditional	T6	AXIOM-NO-TACHYON-IN-CHAMBER	in-chamber deformation sweep	DISCLOSED-CONSISTENT (dep. R3)
R9 sufficiency AUDIT	machine-lane	(none)	mount Gate-7 CSVs + `reproduce_all.py`	BLOCKED_INPUTS → VERIFIED

REDUCE-vs-RELOCATE verdict on the plan. The plan does **not** turn one hard problem into three harder ones. Crucially, the two genuine-physics fronts collapse onto **one shared object**: μ_{cell}, the spectral value of the uniform operational cell Δ_0. R1 (the hierarchy via θ_H⋆), R3/R6 (the doublet Casimir magnitude), and R5 (the breathing-singlet c_{loop}) **all** reduce to "what is μ_{cell}, and can it be read v-independently?" — and the already-run firewall answers: **no v-independent readout exists; the only anchor ($\partial_\sigma V = 0$) IS the hierarchy.** So the plan *reduces* the gate to a single named floor-value residue plus a SADDLE-or-rescue verdict gated on one absent multiplicity table — it does not multiply the difficulty. Each path either (a) names a single target-blind axiom that pays a debt in plain sight (R1, R2, R4, R5, R7, R8), (b) is a bounded owner-artifact computation that decides a sign (R3/R6 the multiplicity table; R9 the Gate-7 CSVs), and every μ_{cell}-touching path carries the explicit κ³/π kill-test (the $\partial_\sigma V = 0$ anchor is forbidden) so a circular closure cannot be banked. The honest expected outcome of a full campaign: **~4 AXIOM-CLOSED (R1 as 2nd-anchor, R4, R5 to the cell residue, R7), ~2 DISCLOSED-CONSISTENT (R2, R8), 1 sharper-OPEN/saddle-leaning (R3), and 2 BLOCKED_INPUTS pending owner artifacts (R6 multiplicity table, R9 Gate-7 CSVs).** No DERIVED-CLOSED is promised; the gate would move from PARTIAL-with-diagnostic-Hessian to **PARTIAL-with-a-named-axiom-floor + a decided-or-saddle Hessian verdict** — a real

honesty/structural gain, **not** a promotion. **The single most likely *negative* — and the most valuable one — is R3 firing: the shape-doublet saddle, if the Casimir rescue fails, REFUTES the moduli-mass-generation row and downgrades it per F.10.4-2. That is the gate working as designed.**

A. Anchoring & Hardening Map

This section runs SG-6 through **our internal honesty methodology** — the same hardening method that produced the live **Gaps & Walls Register**: **classify each residual as a *gap* (a route exists; what is missing is a finished computation or a measured input) or a *wall* (the route itself is the problem) → hunt the implicit assumption it hides → name the measured-invariant truth it must terminate on → assign the most conservative defensible disposition** — exactly that method (the framework, the gap-vs-wall split, and the disposition vocabulary are defined in the linked Register). The discipline is honest bookkeeping, not closure: *no theorem breaks a wall*, and the ceiling on every line below is **serious candidate / NOT validated**. The only existing measured anchors a residual may terminate on are **{$\hbar$, M_Pl, spectrum-E, α_i(M_Z), y_t, |V_us|}**; anything else must name a NEW invariant, hinge on a SCHEME-anchor, be a COMPUTATION-DEBT, or have NO witness → OPEN. Disposition words are used exactly as the Register defines them: **DERIVED / DERIVED-GIVEN-E / AXIOM-CLOSED / DISSOLVED / SCHEME-ANCHORED / OPEN / BLOCKED / measured-but-irreducible**.

A.1 Per-residual anchoring map (one row per §3.1 residual)

Residual (§3.1)	Gap or Wall (+kind)	Measured-invariant truth it must terminate on	Honest disposition	What would HARDEN it (concrete next step)
R1 — θ_H⋆ ≈ 2.46×10⁻¹⁴ read from the V_Hos minimum	**WALL** (SCALE / hierarchy wall; the only anchor *is* the answer)	**No existing anchor reaches it; needs a NEW named invariant.** Carries ~85% of the EW hierarchy via v_EW = θ_H⋆/(2πR_γ); v_EW is a *second dimensionful anchor* that cannot be built from {M_Pl, ℏ, dimensionless frozen geometry} without a σ-carrying length (Buckingham-π). It is **measured-but-irreducible** — v_EW=246 GeV is a datum, not a deduction here.	**AXIOM-CLOSED** at **AXIOM-VEW-SECOND-ANCHOR**; θ_H⋆ as "derived-small" stays **OPEN/RELOCATION** (Path A circular, Path B wrong-mechanism — firewall already run).	Hand a specialist the *target-blind θ_H⋆ derivation theorem*: compute θ_H⋆ from {γ, n_H=1, η_BK, R_γ} + a UV reference frozen **before** v_obs; if it provably needs ∂_σV=0 (the hierarchy), R1 is confirmed RELOCATION and the axiom stands.
R2 — Global stabilization is an explicit NON-claim	**WALL** (textbook SHARED-OPEN; rivals also struggle)	**Terminates on nothing measurable here** — it is a *claim-boundary* statement, not a number. No anchor; disclosure only.	**DISSOLVED-as-non-claim → DISCLOSED-CONSISTENT.** Not AXIOM-CLOSED as a flip; no axiom closes global stabilization. Outside-chamber configs are rejected by admissibility, *not stabilized*.	Mechanical: enforce the F.10.2 vocabulary rule (ban "full stabilization"/"no flat directions" without the chamber qualifier) and keep any Paper-IV V(σ) cross-reference from being cited to upgrade this certificate (F.8).
R3 — Positive-definite Hessian is DIAGNOSTIC-ONLY (named saddle attack surface)	**GAP** (computation-debt; one Λ-free piece already done)	**SCHEME-anchored + COMPUTATION-DEBT.** The Λ-free shape-doublet curvature is already **−1 (SADDLE)** target-blind; the breathing singlet rides the uncomputed c_loop. No existing measured anchor; the rescue magnitude carries μ_cell (a scheme object).	**OPEN — saddle-leaning** (sharper-OPEN). The named F.9.7(2) falsifier is the *closest to firing* of any SG-6 residual.	Mount the **admissible-rep multiplicity table** (ABSENT) + regularized zeta; compute sgn(boson−fermion surplus) and the magnitude μ_cell·q≳0.5 → decide rescue vs SADDLE-CONFIRMED. **Do NOT bank boson-dominance without the table.**
R4 — Witnesses fix moduli WITHIN the	**WALL** (inherited SHAPE/given-E	**Terminates on spectrum-E** — the SM chiral content + χ=−3	**AXIOM-CLOSED** at **AXIOM-GIVEN-E-CHAMBER** (already	Nothing closes it at SG-6 (it is SG-1 SHAPE-forcedness,

Residual (§3.1)	Gap or Wall (+kind)	Measured-invariant truth it must terminate on	Honest disposition	What would HARDEN it (concrete next step)
selected chamber (given-E)	wall from SG-1/SG-3)	are supplied as input; witnesses operate *on* E, they do not derive it. given-E ≠ derivation of E.	standing via SG-1 DECLARED-FROZEN / SG-3 DERIVED-GIVEN-E). This is the load-bearing conditionality of the whole gate.	SELECTED-not-forced-absolute). The only contained SG-6 strengthening is R7 (promote one witness from "selected" to "symmetry-forced").
R5 — c_loop / breathing-mode singlet = the 13D FRG-2 β-vector wall	**WALL** (GRANULARITY-root residue; shared deep FRG-2 wall w/ Gap-04, a6)	**Terminates on a NEW invariant** — the spectral VALUE of the uniform operational cell Δ_0 (μ_cell), a floor-value residue *on the footing of* $\hbar$ but **not supplied by the corpus**. The only available anchor (∂_σV=0) IS the hierarchy → **forbidden** (circular).	**BLOCKED_INPUTS**, reducible to **AXIOM-CLOSED** at **AXIOM-UNIFORM-CELL-VALUE**. The cell-sum kills the continuum divergence (DISSOLVED leg) but leaves one log-scheme factor μ_cell → **SCHEME-ANCHORED** residue.	The deep specialist target: compute the **13D FRG-2 β-vector** → c_loop = tr[a_6] of the σ-fluctuation determinant, target-blind. Short of that, supply an *independent* floor-value for Δ_0 (footing of $\hbar$) — **never** anchor μ_cell at ∂_σV=0.
R6 — Casimir net sign (the only route that could flip R3)	**GAP** (BLOCKED-INPUTS; one machine-lane decision)	**SCHEME-anchored** (per-sector sign LOCKED by geometry; net sign = graded boson−fermion surplus over the **ABSENT** multiplicity table; magnitude carries μ_cell, no v-independent readout).	**BLOCKED_INPUTS / NEEDS-TUNING.** The rescue is conditionally available but **unbankable** — banking boson-dominance (~3:1 heat-kernel) without the table is the κ^3/π-forbidden plausible-sign. $\chi=-3$ is the opposing frozen fact.	Mount the multiplicity table → form $\Sigma(-1)^F d\cdot n\cdot(w_1^2+w_2^2)$ → VERIFIED-RESCUE (net>0, μ_cell·q≳0.5) or SADDLE-CONFIRMED (net≤0 → F.9.7(2) fires). Highest physics value-per-effort *if* the table exists.
R7 — τ=ω "non-zero potential" leans on an uncomputed potential	**GAP** (bounded group-theory; cleanest witness)	**Terminates on a symmetry invariant** — order-three modular fixed-point uniqueness, *not* a measured number. Symmetry-protected, not read from a curve.	**AXIOM-CLOSED** at **AXIOM-MODULAR-FIXED-POINT**, → **DERIVED** if uniqueness is proven. The one witness genuinely upgradable by a contained computation.	Prove τ=ω is the **unique** fixed point of the F^+ Cartan-torus modular group, independent of any one-loop potential — converts the leg from "potential we did not compute" to "symmetry-forced point." (Also hardens SG-8's inherited τ=ω soft spot.)
R8 — No-tachyon check is conditional on the declared spectrum	**GAP** (consistency sweep; admissibility-bounded)	**No new anchor** — conditional on the frozen spectrum + admissibility-bounded in-chamber deformations; F.5	**DISCLOSED-CONSISTENT** at **AXIOM-NO-TACHYON-IN-CHAMBER**; dependent on R3 (a confirmed shape-doublet saddle is the field-space analog of	Run the consistency sweep over the admissibility-bounded deformation class; couple the verdict to R3's Hessian-sign outcome (saddle ⇒

Residual (§3.1)	Gap or Wall (+kind)	Measured-invariant truth it must terminate on	Honest disposition	What would HARDEN it (concrete next step)
		already bounds in-chamber responses.	a tachyon and would put pressure here).	pressure on the no-tachyon leg in that direction).
R9 — Phenomenological sufficiency inherits downstream AUDIT (Gate-7 radii/spectrum)	**GAP** (COMPUTATION-DEBT; owned by SG-7)	**Hinges on a SCHEME-anchor inherited from SG-7** — Gate-7 is now **Diagnostic-only**, its threshold vector (+4.8424, −3.1112, −1.7313) **fitted-to-target, not formula-derived**; signs geometric, magnitudes SCHEME-ANCHORED. SG-6's sufficiency is only as machine-real as that harness.	**AUDIT**, inherited from SG-7 (**SCHEME-ANCHORED / BLOCKED_INPUTS**). Closes when SG-7's reproduction harness is mounted; cannot be more real than its upstream.	Mount the Gate-7 ledger CSVs + `reproduce_all.py` (ABSENT); regenerate the threshold vector + unification residual target-blind; confirm a small δ on the SG-6 radii moves Gate-7 output within the F.3 δ-response column. (Owned by SG-7 — closing SG-7's δ-harness AUDIT closes this.)

A.2 Gate-level rollup

- **Overall disposition.** SG-6 is **PARTIAL → conservatively OPEN** (matching the live Register's SG-6 line). The genuine, strong content is the **witness-type predicate** (a *type* test that makes "a potential we minimized at the answer" fail by construction) plus the discrete/topological witnesses and the S_3-critical-point structure — these are real, target-blind facts, not promotions. The open surface terminates as **3 AXIOM-CLOSED** (R1 as a second irreducible anchor, R4 the given-E scope, R7 the modular fixed point — likely DERIVED), **1 reducible-to- AXIOM-CLOSED but BLOCKED** (R5 → AXIOM-UNIFORM-CELL-VALUE, gated on the FRG-2 wall), **2 DISCLOSED-CONSISTENT** (R2 global-stabilization non-claim, R8 conditional no-tachyon), **1 OPEN/saddle-leaning** (R3 — the named F.9.7(2) falsifier, *already* −1 on Λ-free geometry), and **2 BLOCKED_INPUTS pending owner artifacts** (R6 the multiplicity table, R9 the SG-7 reproduction harness). **No DERIVED-CLOSED is promised; PROMOTIONS: 0.**
- **Anchors it depends on. spectrum-E** (R4: the gate certifies internal moduli-control *given* the supplied chiral content + $\chi=-3$, never derives it) and ***ħ*-footing only** (R5: μ_cell is a floor-value residue *on the footing of ħ*, but its spectral value is **not supplied** — it is a NEW invariant, not the measured *ħ* itself). It depends on **no** clean termination on M_Pl / $\alpha\_i(M\_Z)$ / y_t / $|V\_us|$; the one residual that touches a *number* (R1, the hierarchy) terminates on **measured-but-irreducible v_EW**, not on an existing anchor. R9 inherits SG-7's **SCHEME-anchor**, not a measured one.
- **Single highest-leverage hardening move. Mount the admissible-rep multiplicity table** (currently ABSENT). It is the one owner artifact that decides the net Casimir sign and thereby flips the named saddle attack surface (R3/R6) from "leaning saddle" to a *decided verdict* — VERIFIED-RESCUE (minimum) or SADDLE-CONFIRMED (F.9.7(2) fires, the moduli-mass-generation row downgrades per F.10.4-2). It is the highest physics value-per-effort step and the only one that

converts a diagnostic into a falsifiable result **without new tuning**. The deeper twin wall (R1/R5, the hierarchy and the breathing singlet) collapses onto a **single object — μ_cell** — which the already-run firewall confirms has **no v-independent readout** (its only anchor ∂_σV=0 IS the hierarchy; the κ^3/π kill-test fires), so that front hardens only to a *named axiom floor*, never to a derivation.

🎯 Target anchor(s) for this gate

The body of SG-6 carries **no new measured anchor** and terminates on nothing of its own: the chamber witnesses fix the moduli **WITHIN the already-selected chamber, given-E** (witness-*type* test, **not** a tuned potential minimized at the answer) — so the witness ledger terminates on **spectrum-E** (R4, **AXIOM-CLOSED** at AXIOM-GIVEN-E-CHAMBER) and on the ***ħ*-footing only** (R5: μ_cell is a NEW invariant on the footing of *ħ*, not the measured *ħ* — **SCHEME-ANCHORED / BLOCKED**), and **global stabilization is an explicit NON-claim** (R2, DISCLOSED-CONSISTENT — it terminates on no number). The **one acute open leg is $\theta_{H^\star} \approx 2.46\times10^{-14}$**, *read from the chamber one-loop V_Hos minimum* and carrying **~85% of the EW hierarchy** via v_EW = θ_H⋆/(2πR_γ): its target anchor is **v_EW = 246 GeV**, which here is **measured-but-irreducible** (a second dimensionful anchor, not deducible from {M_Pl, *ħ*, dimensionless frozen geometry}) — θ_H⋆ as "derived-small" stays **OPEN/RELOCATION**, AXIOM-CLOSED at AXIOM-VEW-SECOND-ANCHOR. **No new anchor is introduced; PROMOTIONS: 0.** The hardening move is **not** a new datum but the **θ_H⋆ / Hessian derivation** — a target-blind θ_H⋆ from {γ, n_H=1, η_BK, R_γ} + a pre-v_obs UV reference, plus the positive-definite Hessian decision at (1,1,1) gated on the ABSENT admissible-rep multiplicity table (R3/R6).

---

5. References & source map

5.1 Website source-of-truth (common material — link, don't duplicate)

- **Paper I, GUT.html** — https://physics.magflowmeters.com/articles/GUT.html
- **§6.6** Gate-6 card (binding status: *Claimed certificate pass under declared admissibility and moduli-control assumptions*); **§6.12** falsification map (Gate 6 row: *Diagnostic only* if Hessian falsified / *Open* if phenomenological sufficiency falsified); **§6.13** certificate summary.
- **§5.5** narrative module; the **$y = R^2$ curve-not-a-number** discipline (§6 opening of the gate index).
- **Appendix F** (Stabilization — the binding authority for Gate 6): **F.1** stabilization claim; **F.2** potential/witness structure; **F.3** the 12-row moduli ledger; **F.4** spectrum + tachyon check; **F.5** sensitivity; **F.6** certificate JSON; **F.9** Stabilization-vs-Admissibility (F.9.1 five concepts, F.9.2 what is proved, **F.9.4 the NON-claim**, **F.9.5 the Hessian diagnostic**, F.9.6 required vocabulary, **F.9.7 falsification paths**); **F.10** claim-type ledger (F.10.1 ledger, **F.10.2 vocabulary substitution rule**, F.10.3 binding statement, F.10.4 reviewer falsification path).
- **Appendix CR module CR6** — reader companion (explanatory only; mints no status).

- **§6.8 / Appendix H** — the Gate-8 Hosotani object on which R1 (θ_H⋆ / v_{EW}) rides; **Appendix G** — the Gate-7 threshold pipeline on which R9 (sufficiency) depends.
- **Paper IV, TOE.html** — https://physics.magflowmeters.com/articles/TOE.html (the four-term $V(\sigma)$ discharge; EXTERNAL; cross-reference only — changes nothing in this gate's certificate per F.8).

This dossier recaps only what is needed to attack; the site controls all common material.

5.2 Corpus locations (authoritative inputs to this dossier)

Source	Path	Role
Per-gate dossier spec	`…/rendered/TOE/PER_GATE_DOSSIER_SPEC.md`	structure (sections 0–5)
SG-6 status line	`…/rendered/TOE/SCOPED_GUT_PER_GATE_STATUS_LEDGER.md`	the PARTIAL label + honest caveats (carried verbatim)
Λ-anchored moduli test	`…/rendered/TOE/TEST1_LAMBDA_ANCHORED_MODULI_STABILIZATION_2026-06-23.md`	Step A/B/C verdict; chamber V(u) is fiat-not-minimized; the four-term V(σ) is σ-only; counting/under-determination wall
Λ-free chamber potential (S_3 split)	`…/rendered/TOE/TEST2_LAMBDA_FREE_CHAMBER_POTENTIAL_2026-06-23.md`	the $\mathbf{3} = \mathbf{1} \oplus \mathbf{2}$ decomposition; $(1,1,1)$ = S_3-critical-point free; singlet=wall, doublet=computable
Shape-doublet Hessian	`…/rendered/TOE/TEST_SHAPE_DOUBLET_STABILITY_2026-06-23.md`	the Λ-free doublet curvature Hessian $= -\mathbf{1}$ (SADDLE); structure-constant triangle is the destabilizer
Casimir rescue test	`…/rendered/TOE/TEST_CASIMIR_RESCUE_SHAPE_DOUBLET_2026-06-23.md`	per-sector sign LOCKED; net sign needs the ABSENT multiplicity table; $\mu_{\text{cell}} \cdot q \gtrsim 0.5$
Cell-sum c_loop attack	`…/rendered/TOE/TEST_CLOOP_CELLSUM_ATTACK_2026-06-23.md`	c_loop = finite cell-sum + one log-scheme factor μ_{cell}; uniform-Δ_0 residue
SCALE/hierarchy final verdict	`…/rendered/TOE/SCALE_HIERARCHY_FINAL_VERDICT_2026-06-23.md`	θ_H⋆ ≈ 2.46e-14 carries 85% hierarchy; Path A circular (Buckingham-π), Path B wrong-mechanism; v_{EW} a 2nd anchor; $\partial_\sigma V = 0$ forbidden
Closure campaign (R1)	`…/rendered/TOE/CLOSURE_CAMPAIGN_RESULT_2026-06-24.md`	the decisive SCALE-firewall verdict (μ_cell no v-independent readout); κ³/π kill-test; AXIOM_CLOSED ≠ proven
Closure campaign (R2)	`…/rendered/TOE/CLOSURE_CAMPAIGN_ROUND2_2026-06-24.md`	demotion-on-verify norm; relabel-fail kill-test discipline
Axiom ledger	`…/rendered/TOE/AXIOM_LEDGER.md`	ROOT-1 `AXIOM-COSTFLOOR` (existence not value) + FLOOR-VALUES-RESIDUE (Δ_0 value); ROOT-2 SHAPE (selection ≠ derivation)
GUT manuscript	`…/rendered/GUT/GUT.md`	§5.5 / §6.6 (Gate-6 card); Appendix F (F.1–F.10); R1.6 witness predicate; F.3 12-row ledger

5.3 Frozen-object hash index (load-bearing; READ-ONLY)

Branch `dcc66f1b2685` · manifest meta `a5b1e6f9d951` · three Cartan radii `634438ce0776` / `2381d472c62e` / `0e8b8dba2cf0` · chamber-center $\vec{u} = (1, 1, 1)$ (A1.2) · $\tau = \omega$ `03b30a9c931a` · η_{BK} `84e94518d3f5` · RG transport `f531205a9159` · M_Z `a6852c7a6b00` · $N_{u,d,e}$ `20dc4e0b8220` · θ_F `1ff57f48d45a` · spin-$\mathbb{C}$ index / $\chi = -3$ (R1.4) `0fd19c9ae0c1` · threshold vector $(\delta_1, \delta_2, \delta_3) = (+4.8424, -3.1112, -1.7313)$. **θ_H⋆ (Hosotani minimum), μ_{cell} / Δ_0 spectral value, c_{loop}, the admissible-rep multiplicity table, and M_R: no SG-6 hash — read / uncomputed / ABSENT.**

Closing honest statement

SG-6 is **PARTIAL**. Its genuine, defensible content is real and strong: the **witness-type predicate** — every downstream-used modulus carries a structural witness of declared TYPE (Weyl-rigid chamber / modular fixed point / integer winding / discrete topology / boundary-condition pin), with "a potential we minimized at the answer" failing **by construction** — backed by the **12-row witness ledger**, a **no-tachyon check**, a **coverage lint**, and the honest **claim-type ledger** that fences exactly what is proved. Its open surface is equally clear: **global stabilization is an explicit NON-claim** (F.9.4/F.10.1); the **positive-definite Hessian is DIAGNOSTIC-ONLY** (F.9.5) — and on current Λ-free perturbative geometry the **shape-doublet Hessian is a SADDLE** (-1, the named F.9.7(2) attack surface leaning toward firing, rescuable only by a Casimir term whose net sign needs an ABSENT multiplicity table); the acute soft spot is **θ_H⋆ ≈ 2.46×10⁻¹⁴**, carrying **~85% of the EW hierarchy**, **read from a one-loop minimum, not derived-small**; and the witnesses fix the moduli **within the selected chamber** (given-E). The attack plan reduces these to named, target-blind axioms and bounded owner-artifact computations — and crucially collapses the two genuine-physics fronts (the hierarchy and the Hessian sign) onto **one shared object** μ_{cell}, which the already-run firewall confirms has **no v-independent readout** ($\partial_\sigma V = 0$, its only anchor, **IS** the hierarchy — the forbidden, circular anchor; the κ³/π kill-test fires). The realistic ceiling is a named axiom floor (AXIOM-VEW-SECOND-ANCHOR, AXIOM-UNIFORM-CELL-VALUE, AXIOM-MODULAR-FIXED-POINT) plus a **decided-or-saddle Hessian verdict** — **not** a promotion; and the most valuable likely outcome is a *negative*: the shape-doublet saddle firing F.9.7(2) if the Casimir rescue fails. **PROMOTIONS:0; frozen branch `dcc66f1b2685` / `a5b1e6f9d951` READ-ONLY; given-E ≠ derivation of E; selection ≠ derivation; dissolved ≠ solved; nothing applied, nothing deployed.**

Dossier built 2026-06-24. Our geometry (13D K6 branch) only. Common material referenced to the published website source-of-truth, not duplicated.