

Claims

SN111 v0 Settings and Operations Guide

Current Mainnet Workflow, Scoring, and On-Chain Parameters

Bittensor Subnet 111

Claims Team

Philipp Koellinger

Professor of Economics, Vrije University

Christian Roessler

Associate Professor of Economics, California State University-East Bay

Ogban Ugot

Engineer

Configuration snapshot: 29 August 2026

This guide describes the current mainnet configuration of CLAIMS as of the snapshot date. It explains what each setting controls and how a 50-paper batch moves from miner assignment through validation, Bronze generation, Silver adjudication, scoring, and on-chain weights. Values that are inherently dynamic, including registration cost and current occupancy, are identified as snapshots rather than fixed mechanism constants.

Abstract

CLAIMS converts scientific papers into structured records that link scientific claims to evidence, source locations, qualifiers, and provenance. This guide records the current v0 operating mechanism for Bittensor Subnet 111. It separates the stable v0 network envelope from the current **agent_v1** extraction and scoring pipeline; states the operative mainnet settings; describes miner selection, validator coordination, diagnostics, Bronze and Silver construction, failure ownership, scoring, and payout; and identifies the on-chain parameters that shape entry and timing. The guide is limited to active v0 behavior. Trusted Gold records, hidden synthetic challenges, miner judging, behavioral-family accounting, marginal-contribution rewards, and canonical production-graph ingestion are not active v0 mechanisms.

Keywords: scientific claim extraction; evidence records; Bittensor; Subnet 111; miner selection; validator; Bronze; Silver; adjudication; scoring.

Contents

1	What the Current System Does	1
1.1	Participant roles	1
1.2	Bittensor and runtime terms	2
1.3	The meaning of v0 and agent_v1	2
1.4	Active v0 boundary	3
2	Current Settings at a Glance	3
2.1	Application-level mainnet settings	3
2.2	On-chain and chain-facing settings	5
3	The Objects That Move Through the Workflow	6
3.1	Claims, evidence, and source grounding	6
3.2	The distinct roles of Bronze and Silver	7
4	One Complete Claims Scoring Cycle	7
4.1	Canonical assignment and collection	7
4.2	Submission and hydration	8
4.3	Diagnostics and Bronze	8
4.4	Pairing, adjudication, and Silver	8
4.5	Scoring, persistence, and weights	8
4.6	The three independent clocks	8
5	Miner Selection and Miner Operations	9

5.1	Eligibility and immutable assignment	9
5.2	The implemented adaptive selector	9
5.3	The miner’s production contract	10
5.4	Practical score implications	10
6	Validator Processing and Control Settings	10
6.1	Diagnostics	10
6.2	Bronze reference generation	10
6.3	Silver comparison and adjudication limits	11
6.4	Origin blindness and independent roles	11
7	From Silver to Scores and Payouts	12
7.1	The scored Silver set and importance weights	12
7.2	Coverage	12
7.3	Diagnostic quality	12
7.4	Adjudication quality	13
7.5	Paper and batch scores	13
7.6	Worked paper-level cases	14
7.7	Winner-takes-most weights and ties	14
8	Failure Ownership, Auditability, and Security	14
8.1	Failure ownership	15
8.2	Minimum reconstructable record	15
8.3	Security boundaries	15
9	Configuration Summary	16
9.1	Identity, scheduling, and selection	16
9.2	Diagnostics, Bronze, and scoring	16
9.3	Silver and persistence	17
10	What the v0 Result Means	18
11	Conclusion	19

1 What the Current System Does

CLAIMS produces reusable, machine-readable representations of scientific papers. A record preserves not only a statement, but also the evidence associated with it, the source span from which it was derived, and the qualifiers needed to interpret it. The mechanism uses Bittensor competition to obtain multiple independently produced candidate records and a validator-side process to compare, adjudicate, and score them. In the active v0 mechanism, every scoring paper receives a Bronze reference; miner and Bronze candidates are evaluated without origin labels; equivalent candidates are merged into a canonical Silver record; and batch performance is converted into concentrated miner weights.

1.1 Participant roles

Table 1. Roles in the active v0 workflow.

Role	Active responsibility
Miner	Receives the canonical paper batch, produces one structured agent_v1 artifact per paper, and returns a compact batch response before the deadline. A miner chooses its own models, tools, retrieval, parsing, and repair process.
Validator	Obtains the canonical assignment, queries assigned miners, runs structural and semantic diagnostics, creates or retrieves Bronze, runs origin-blind comparison and Silver adjudication, calculates scores, persists the audit record, and proposes or submits weights.
Claims backend	Fixes the canonical task, paper set, selection seed, UID-to-hotkey snapshot, miner assignment, collection lease, artifact records, and persisted scoring outputs for a batch.
Bittensor chain	Registers UIDs, exposes the metagraph and immunity state, grants validator permits, carries miner queries and responses, and accepts validator weights.
Reference pipeline	Produces the validator-owned Bronze candidate for each scoring paper. It is a comparison baseline, not a reward-competing miner and not ground truth.

1.2 Bittensor and runtime terms

Table 2. Terms used throughout the operating guide.

Term	Meaning in this guide
UID	Integer neuron slot on the subnet. A UID is associated with the currently registered hotkey and can later be reassigned through deregistration and registration.
Hotkey	Operational Bittensor key that identifies and signs for a miner or validator neuron.
Coldkey	Ownership and funding key associated with one or more hotkeys. It is not transmitted as part of a Claims task.
Axon	Network service endpoint through which a registered neuron receives Bittensor requests. A serving axon is required for adaptive miner eligibility.
Dendrite	Validator-side Bittensor client path used to send the task synapse to miner axons and receive compact responses.
Metagraph	Chain-derived snapshot of registered UIDs, keys, endpoints, permits, immunity, stakes, and incentive state for a subnet.
Synapse	Typed Bittensor request-and-response envelope. The Claims synapse carries batch metadata, paper assignments, version identifiers, and compact miner response records.
Agent harness	Runtime wrapper that launches a selected model and tool loop. <code>hermes-cli</code> is the harness selected in the mainnet validator configuration; the model remains operator-selected.
File-agent	Bounded validator workflow in which comparison, judging, canonicalization, or audit roles operate on structured files in a controlled workspace.

1.3 The meaning of v0 and agent_v1

Two version labels coexist because they describe different layers:

- **v0 is the mechanism and network envelope.** The task version remains `claims_task_v0`. The task protocol also identifies the current scoring version as `agent_v1_pass4_deterministic_v0`.
- **agent_v1 is the current extraction and validation pipeline.** It is the default miner pipeline and the canonical path for current structured agent artifacts. The older direct-model v0 miner and validator modules remain only for legacy compatibility and are not the normal current path.

Accordingly, “v0” in this guide never means that the legacy `miner/v0` implementation is active. It identifies the current incentive mechanism, while `agent_v1` identifies the implementation used inside that mechanism.

1.4 Active v0 boundary

The active v0 loop contains routine machine extraction, diagnostics, Bronze comparison, origin-blind machine adjudication, canonical Silver construction, and rank-based rewards. It does not contain maintained Gold labels, Anti-Gold challenges, hidden post-commit sampling, a miner-judge role, behavioral-family clustering, marginal-contribution payout, or production-grade canonical graph ingestion. Silver is therefore the strongest machine-curated record produced by the current process; it is not represented as verified truth.

The current backend also supports canonical assignments, a temporary collection lease, and follower reuse of completed batch artifacts. These controls reduce duplicate inference and make repeated validator execution consistent. They do not create the independent substantive validation architecture contemplated for later versions. V0 evaluation remains centrally coordinated through the Claims backend and the active validator pipeline.

2 Current Settings at a Glance

The tables in this section state the mainnet configuration and the stable on-chain parameters visible on 29 August 2026. Private provider credentials and model identifiers are intentionally absent because they are validator- or miner-selected implementation choices rather than common mechanism constants.

2.1 Application-level mainnet settings

Table 3. Core mainnet settings and their meanings.

Setting	Current value	Meaning
Bittensor network and sub-net	<code>finney</code> ; netuid 111	The validator and miners operate on the SN111 mainnet metagraph.
Claims backend partition	<code>mainnet</code>	Selects the mainnet Claims data partition. This is separate from the Bittensor chain endpoint.
Task version	<code>claims_task_v0</code>	Stable v0 request and response envelope.
Current miner pipeline	<code>agent_v1</code>	Canonical artifact-production path for new miner runs.
Current scoring version	<code>agent_v1_pass4_deterministic_v0</code>	Version identifier carried in the synapse for the active scoring path.
Papers per batch	50	Each selected miner receives the same canonical 50-paper assignment.
Selected miners	10	The adaptive selector targets ten eligible UIDs per batch.
Paper reuse	<code>false</code>	Mainnet selection excludes papers already assigned unless the profile is deliberately changed for a test.

Setting	Current value	Meaning
Miner-response deadline	3,600 seconds	One hour for the complete batch response.
Collection polling	5 seconds	Follower validators poll for batch completion at this interval.
Collection wait bound	3,900 seconds	The follower wait exceeds the one-hour miner deadline by five minutes.
Cycles per process run	4	<code>CLAIMS_MAX_STEPS=4</code> caps a standard invocation at four completed scoring cycles.
Delay between completed cycles	21,600 seconds	A six-hour delay is applied before another cycle begins. It is not a fixed six-hour wall-clock round: cycle execution time is additional.
Batch aggregation	Arithmetic mean	Eligible paper scores are averaged for each miner.
Payout mode	Winner-takes-most	Rank shares are 70/16/8/4/2 percent for the first five occupied slots when at least five miners score positively.
Weight submission	Enabled	<code>CLAIMS_AUDIT_ONLY=false</code> ; the normal validator persists the result and submits weights.
Validator pipeline	<code>auto</code>	Routes current artifacts to <code>agent_v1</code> while retaining legacy compatibility.
Semantic diagnostics	Enabled	Both the audit method and <code>agent_v1</code> validation mode are set to LLM-backed review.
Diagnostic threshold	0.0	Diagnostics remain score-bearing; the mainnet configuration does not impose a positive pass cutoff before the Silver path.
Silver workflow	Enabled; <code>file-agent</code>	Runs comparison, two anonymous direct judges, a conditional tiebreaker, canonicalization, and independent canonical audit.
Silver fallback	<code>none</code>	Failure of the configured <code>file-agent</code> workflow fails that paper rather than silently switching to a different adjudication system.

2.2 On-chain and chain-facing settings

Table 4. SN111 chain parameters and current capacity snapshot.

Parameter	Value	Operational meaning
Maximum UIDs	256	Total neuron capacity. The subnet displayed 256 active UIDs out of 256 at the snapshot, so occupancy was full.
Maximum validators	64	Maximum validator-permit capacity. This is not the number of miner slots.
Chain tempo	360 blocks	One on-chain epoch is approximately 72 minutes at 12-second blocks. This clock is independent of the Claims six-hour post-cycle delay.
Subnet immunity period	21,600 blocks	Approximately 72 hours at 12-second blocks. The validator profile reads this value from chain.
Immunity priority window	7,200 blocks	Claims-side selection control corresponding to approximately the final 24 hours of immunity.
Minimum registration burn	0.0500 TAO	Owner-set lower bound for the dynamic registration burn.
Maximum registration burn	1 TAO	Owner-set upper bound for the dynamic registration burn.
Maximum registrations per block	1	Chain limit on successful registrations in a block.
Displayed registration cost	0.0000 TAO	Dynamic snapshot display; distinct from the owner-set burn bounds and not a stable entry-price constant.
Burn increase multiplier	1.26	Raises the dynamic burn after registration activity.
Burn half-life	4,050 blocks	Approximately 13.5 hours at 12-second blocks.
Collateral lock share	0%	No registration collateral share is locked under the displayed setting.
Collateral drain ratio	1	Displayed collateral drain parameter.
Mechanism count and split	1 mechanism; 100%	SN111 exposes one active mechanism receiving the complete mechanism-level emission split.

The instantaneous registration burn is a dynamic chain-state value rather than a fixed mechanism constant. Entry calculations use the live chain quote together with the bounds and adjustment parameters above. Capacity, occupancy, and active validator counts can also change after the snapshot.

3 The Objects That Move Through the Workflow

The workflow distinguishes source material, candidate artifacts, reference content, adjudicated content, and findings. Treating these as separate objects prevents a baseline from being mistaken for truth and prevents a miner’s raw output from being mistaken for the final record.

Table 5. Principal v0 objects.

Object	Function
Paper source	The assigned document and its stable paper identifier, source URL, and source hash. Source locations support later grounding checks.
Task and batch	Immutable identifiers for the assignment window, paper set, selection seed, UID-to-hotkey snapshot, selected miners, network, protocol version, and scoring version.
Miner artifact	A structured <code>agent_v1</code> record containing claims, evidence, relations, source references, quotations or spans, and the metadata needed to validate cross-references and provenance.
Diagnostic finding	A structured issue attributed to a miner artifact, such as invalid JSON, schema failure, a missing layer, absent evidence, unresolved grounding, or a rigor concern. Findings carry severity and a machine-readable code.
Bronze record	The validator-owned reference extraction for the paper. It establishes an economical common baseline and enters comparison on the same origin-blind basis as miner candidates.
Adjudication case	A bounded comparison between potentially equivalent, conflicting, or additional candidates. Judge inputs use validator-owned short references rather than miner identity.
Silver unit	A canonical item produced after comparison and adjudication. It records equivalent candidate identifiers, whether it is required for completeness or an accepted improvement, its final importance class, and any associated findings.
Paper score breakdown	Coverage, diagnostic quality, adjudication quality, the multiplicative paper score, and failure metadata for one miner-paper pair.
Batch result	Per-miner mean scores, rank, ties, payout weights, missing papers, validator-failed papers, and the proposed on-chain weight vector.

3.1 Claims, evidence, and source grounding

A conforming artifact links scientific statements to the evidence that supports, qualifies, or contextualizes them. Source references are material because the validator must be able to find the cited content in the assigned paper. A grammatically plausible statement without a valid source anchor remains an extraction failure. Structural correctness and scientific grounding are evaluated separately: deterministic checks settle parseability, schema, cross-reference, and source-contract questions, while the semantic rigor pass evaluates whether the artifact’s interpretation is supported by the paper.

3.2 The distinct roles of Bronze and Silver

Bronze is produced or retrieved for every scoring paper. It gives all miners a common comparison object and prevents the mechanism from rewarding performance merely for falling below an economical known baseline. It remains fallible and can be corrected by valid miner-only content. Silver is assembled item by item. It retains supported baseline content, merges equivalent representations, accepts verified miner improvements, rejects unsupported additions, and records unresolved or failed cases explicitly. It does not select one whole miner artifact as the final answer. Multiple miners can receive coverage credit for the same canonical unit when their candidates belong to its equivalence group.

4 One Complete Claims Scoring Cycle

Figure 1 gives the operative sequence. Every stage is batch-scoped and tied to the canonical task and release identifiers.




Figure 1. The active v0 cycle from canonical assignment to proposed weights.

4.1 Canonical assignment and collection

The backend creates one immutable paper and miner assignment for the configured assignment window. The first validator atomically proposes the selected miner set. Later validators receive the same task ID, batch ID, papers, selection seed, and UID-to-hotkey snapshot. One validator holds a temporary collection lease and queries unfinished miners. Followers poll the batch record and reuse completed batch-scoped artifacts rather than repeating miner inference. An operator override cannot cause a validator to diverge from an already established canonical assignment.

Every selected miner remains assigned after the snapshot is frozen. The validator does not replace an offline, timed-out, invalid, or missing miner with a different UID. This preserves the ex ante selection rule and makes reliability part of the score.

4.2 Submission and hydration

The synapse carries the canonical batch metadata and paper list. Miners return one compact `articles[]` entry for each assigned paper and upload the full artifacts through the configured backend path. The validator binds each response to the assigned UID and hotkey, checks paper identifiers and hashes, hydrates the full artifact, and records missing or malformed entries.

The one-hour timeout applies to the complete batch response. It is not a per-paper allowance. A miner-side timeout or missing article becomes a zero for that paper. The 3,900-second collection wait is a follower-coordination bound and does not extend the miner's 3,600-second deadline.

4.3 Diagnostics and Bronze

For each available miner-paper artifact, the validator runs deterministic checks for JSON, schema, internal references, required layers, evidence records, and source references. A required model-backed rigor pass then evaluates semantic grounding and scientific interpretation. Diagnostic findings remain attached to the miner-paper pair and enter diagnostic quality.

The validator creates or retrieves a versioned Bronze artifact for the same paper using the reference-miner command. Bronze generation uses the current agent-v1 reference profile and the configured PDF reader. Cached Bronze content is tied to the reference profile and output root; it is not inferred from the miner submissions.

4.4 Pairing, adjudication, and Silver

The Silver pipeline projects Bronze and miner content into candidate objects and searches for potentially related pairs. The mainnet profile uses four neighbors per retrieval direction and disables the additional dense-pair expansion. This bounds comparison cost while preserving local semantic alternatives.

The file-workspace workflow then runs one comparator, two anonymous direct judges, and a conditional tiebreaker. The direct judges use distinct model roles. Deterministic consensus converts their structured votes into a decision. A canonicalizer assembles the resulting units, and an independent canonical-audit role checks the draft. The mainnet configuration specifies no silent fallback: a failed file-agent pipeline is a validator-side paper failure.

After canonicalization, the importance classifier labels units as central, supporting, or minor. This label affects coverage weight. Miner-provided importance labels do not control the score.

4.5 Scoring, persistence, and weights

The validator calculates each miner's coverage of scored Silver units, diagnostic quality, and adjudication quality. Their product is the paper score. Eligible paper scores are averaged by miner. The batch result then assigns rank shares under winner-takes-most, persists the complete score and failure record, and submits the proposed vector unless audit-only mode is active.

4.6 The three independent clocks

Three clocks must not be collapsed into one:

- the **miner-response deadline** is 3,600 seconds;

- the **Claims scheduler delay** is 21,600 seconds after a completed cycle and before another cycle begins;
- the **on-chain tempo** is 360 blocks, approximately 72 minutes at 12-second blocks.

A standard process invocation attempts four cycles, but the cycle execution time is additional to the six-hour delay. The settings therefore do not imply four rounds per calendar day and do not create a fixed half-hour production window inside a six-hour epoch.

5 Miner Selection and Miner Operations

5.1 Eligibility and immutable assignment

Adaptive candidates are registered UIDs with serving axons, excluding the validator's own hotkey. The selection state is UID- and hotkey-specific. A changed registration block resets the stored state for the UID. The mainnet configuration reads the chain immunity period and sets a 7,200-block priority window for under-vetted UIDs near the end of immunity.

After selection, no substitution occurs. Missing, invalid, offline, and timed-out miners remain in the assignment and receive zero for the affected papers. This rule prevents the realized competitor set from being improved after availability becomes known.

5.2 The implemented adaptive selector

The executable selector is identified as `adaptive_v1`. It assigns 20 percent of the sample to development, 20 percent to exploration, and the remainder to performance. With the current sample size of ten, the normal split is:

Lane	Slots	Purpose
Development	2	Evaluate UIDs with fewer than three recorded batches.
Performance	6	Favor miners with stronger smoothed historical scores while retaining a nonzero chance for every candidate.
Exploration	2	Uniform seeded sampling from the remaining eligible UIDs.

Development candidates are ordered first by historical batch count, then by the recorded last-scored timestamp, and then by a seed-dependent hash tie-break. If a miner has historical average score s_i based on n_i recorded batches and the subnet-wide historical mean is $\bar{s}$, its smoothed performance score is

$$\tilde{s}_i = \frac{3\bar{s} + n_i s_i}{3 + n_i}. \quad (1)$$

When no usable history exists, the network mean defaults to 0.5. Performance selection samples without replacement with weight

$$q_i = 0.05 + \tilde{s}_i^2. \quad (2)$$

The floor preserves a positive probability for historically weak miners, while squaring the smoothed score concentrates the performance lane toward stronger results. Exploration is uniform among remaining candidates. The selection seed makes both weighted and exploratory draws reproducible. If a lane cannot fill, the remaining positions are filled from the available pool through the performance sampler.

5.3 The miner’s production contract

A miner receives 50 papers and returns one article entry per paper before the one-hour deadline. Each full artifact contains the structured claim–evidence representation and resolvable source references. The task protocol standardizes inputs, output fields, identifiers, and transport; it does not prescribe the miner’s internal model, prompt, parser, retrieval system, or repair strategy.

The canonical pipeline is `agent_v1`. The command surface supports several agent harnesses and PDF extraction methods, with `pdf-inspector` as the current configuration default. Batch artifacts are uploaded outside the compact dendrite response when the backend URL is configured. A miner gains score from broad valid coverage and accepted improvements, but loses score through missing content, structural or grounding defects, and unsupported additions.

5.4 Practical score implications

The multiplicative score makes reliability and scientific precision complementary:

- a missing paper contributes zero to the batch mean;
- omitting a central Silver unit costs more coverage than omitting a minor unit;
- invalid additions reduce adjudication quality even when the valid portion is broad;
- schema and grounding failures cap diagnostic quality;
- a clean copy of Bronze can provide coverage, but accepted miner-only improvements can expand the scored record and distinguish the miner;
- returning more candidates is useful only when their expected Silver acceptance value exceeds their expected finding cost.

6 Validator Processing and Control Settings

6.1 Diagnostics

The current validator uses one diagnostic file-agent operation per paper to review all available miners when the diagnostic miner batch size is ten. It runs up to ten paper diagnostics concurrently, with two per-miner fallback workers. The rigor role is allowed up to 30 turns and 16,384 output tokens. Diagnostics report structured findings rather than a free-form grade; the final numeric deduction is deterministic.

6.2 Bronze reference generation

The reference command is `python-mclaims_reference_miner`. The mainnet configuration identifies `reference-agent-v1-hermes` as the reference profile, `hermes-cli` as the harness, and `pdf-inspector` as the reader. The model field is blank and must be supplied by the operator. Bronze is created on the validator side and is not a registered reward recipient.

6.3 Silver comparison and adjudication limits

Table 6. Principal Silver capacity and independence settings.

Control	Value	Effect
Pairing neighbors per direction	4	Limits embedding-retrieved candidate neighbors entering relation analysis.
Dense pair addition	0	Disables the additional dense comparison expansion.
Adjudication batch size	13	Maximum anonymous cases packed into a normal judge call before token-based splitting.
Estimated batch input bound	120,000 tokens	Splits oversized judge batches before execution.
Maximum adjudication tokens	81,920	Output budget for the adjudication role.
Maximum adjudication turns	30	Agent-loop turn bound.
Adjudication timeout	1,800 seconds	Per-execution wall-clock bound.
Maximum calls in flight	32	Global concurrent Silver model-call cap.
Adjudication workers	9	Worker bound for adjudication batches.
Distinct direct judges	true	Direct judge A and B use different model roles.
File-agent fallback	none	A failed current workflow is not silently replaced by the legacy graph workflow.
Paper-level Silver workers	50	Maximum concurrent paper pipelines.
Eligible claims per miner	1,000	Upper bound on miner candidates entering Silver for a paper.
Adjudication cases per paper	1,000	Upper bound on primary cases, with capacity shared across miners.
Diagnostic assessment filter	false	Diagnostic findings affect quality but do not automatically remove candidate claims from Silver.
Consolidation neighbors	4	Post-adjudication neighbor bound for duplicate consolidation.

The importance classifier uses a batch size of eight. Persistence writes ordinary Silver records in chunks of 50 and vote records in chunks of 150. These settings control throughput and storage; they do not change the scoring formulas.

6.4 Origin blindness and independent roles

Judge packets replace miner-owned identifiers with validator-owned short references and omit miner origin. Direct judges A and B are distinct, the tiebreaker is invoked conditionally, and the canonical draft is checked by an independent audit role. The mainnet configuration leaves all role-specific model identifiers blank. Operational independence therefore depends on the operator supplying genuinely distinct models for the roles marked distinct.

7 From Silver to Scores and Payouts

The paper score separates three failure margins. Missing final-record content lowers coverage. Structural, grounding, and rigor findings lower diagnostic quality. Unsupported extra content lowers adjudication quality. Their product prevents strength on one dimension from compensating for a zero on another.

7.1 The scored Silver set and importance weights

For paper p , the scored Silver set is

$$\mathcal{S}_p = \left\{ u : \begin{array}{l} u.\text{required_for_completeness} = \text{true} \\ \vee \quad u.\text{scoring_mode} = \text{accepted_improvement} \end{array} \right\}. \quad (3)$$

A surviving Bronze-containing unit is required for completeness. A valid miner-only unit is an accepted improvement. Units with neither role stay outside the coverage denominator.

Miner i covers unit u when at least one of its candidate identifiers belongs to the unit's equivalence group. Let $E_i(u)$ be the corresponding indicator. Final importance weights are

$$\omega(u) = \begin{cases} 0.70, & u \text{ is central,} \\ 0.30, & u \text{ is supporting,} \\ 0.10, & u \text{ is minor.} \end{cases} \quad (4)$$

7.2 Coverage

Coverage is the importance-weighted share of the scored record represented by the miner:

$$C_{ip} = \frac{\sum_{u \in \mathcal{S}_p} \omega(u) E_i(u)}{\sum_{u \in \mathcal{S}_p} \omega(u)}. \quad (5)$$

Accepted improvements enter the denominator and give coverage only to miners whose candidates belong to the accepted equivalence group. A miner can therefore receive credit for content missed by Bronze while the accepted improvement raises the completeness standard for all miners.

7.3 Diagnostic quality

Diagnostic findings use the following severity penalties:

$$\begin{aligned} \pi(\text{blocker}) &= 1.00, & \pi(\text{critical}) &= 0.25, & \pi(\text{major}) &= 0.10, \\ \pi(\text{minor}) &= 0.03, & \pi(\text{warning}) &= 0.01, & \pi(\text{suggestion}) &= 0. \end{aligned} \quad (6)$$

Certain finding codes also cap the maximum diagnostic score:

Table 7. Maximum diagnostic scores under specified failure conditions.

Condition	Maximum score
Blocker, including an unparseable invalid-JSON artifact	0.05
Schema validation failure	0.20
Missing logic or evidence layer	0.25
No evidence records	0.40
No source references, unresolved source span, or quotation absent from source	0.50
Rigor agent skipped	0.60
Rigor agent failed	0.30

Let F_{ip}^D be the diagnostic findings and $\bar{d}_{ip}$ the tightest applicable cap. The score is

$$D_{ip} = \min \left\{ \bar{d}_{ip}, \max \left[0, 1 - \sum_{f \in F_{ip}^D} \pi(f) \right] \right\}. \quad (7)$$

The implementation rounds the post-deduction diagnostic score to four decimals before applying the cap.

7.4 Adjudication quality

Let F_{ip}^A be non-completeness Silver findings attributed to miner i on paper p . Then

$$A_{ip} = \max \left\{ 0, 1 - \sum_{f \in F_{ip}^A} \pi(f) \right\}. \quad (8)$$

An invalid miner candidate normally generates a critical finding and a 0.25 deduction. Completeness findings stay outside this factor because the same omission has already lowered coverage.

7.5 Paper and batch scores

The paper score is

$$s_{ip} = C_{ip} D_{ip} A_{ip}. \quad (9)$$

If Silver contains no scored units and no invalid miner candidates, the validator forces the score to zero.

Let $\mathcal{P}_i$ contain miner i 's scoring-eligible papers. It includes explicit zeroes for miner-side misses and excludes validator-failed papers. The batch score is the arithmetic mean

$$S_i = \frac{1}{|\mathcal{P}_i|} \sum_{p \in \mathcal{P}_i} s_{ip}. \quad (10)$$

Median and minimum scores may be stored for reporting, but the active incentive score is the mean.

7.6 Worked paper-level cases

Suppose the scored Silver record contains a central required unit of weight 0.70, a supporting required unit of weight 0.30, a minor required unit of weight 0.10, and a supporting accepted improvement of weight 0.30. The denominator is 1.40.

Table 8. Illustrative paper scores under the active formula.

Case	Coverage and findings	C	D, A	s
Missed central unit	Covers all except central: 0.70/1.40; no findings	0.500	1.00, 1.00	0.500
Clean baseline coverage	Covers the three required units: 1.10/1.40; one warning	0.786	0.99, 1.00	0.778
Accepted improvement	Covers all four units; one minor diagnostic finding	1.000	0.97, 1.00	0.970
Invalid extra	Covers required units but adds one invalid candidate; no diagnostic finding	0.786	1.00, 0.75	0.589

7.7 Winner-takes-most weights and ties

For at least five positive scorers, the occupied rank-slot vector is

$$(0.70, 0.16, 0.08, 0.04, 0.02, 0, \dots). \quad (11)$$

The runner-up mass is 30 percent and uses geometric factors $1, 1/2, 1/4, 1/8$. When fewer runner-up slots are available, the 30 percent tail is renormalized over the available positive ranks. A sole positive scorer receives 100 percent.

Equal scores retain equal rank. Miners tied across several positions share the combined mass of the occupied slots. Miner identifier is used only to produce a deterministic display order; it does not break a score tie for payout. When all scores are zero, the payout function returns an all-zero vector rather than manufacturing a winner.

8 Failure Ownership, Auditability, and Security

8.1 Failure ownership

Table 9. How active v0 assigns operational failures.

Event	Score treatment	Reason
Miner missing, offline, timed out, malformed, or invalid	Zero for affected paper; paper remains in that miner's mean	Availability and submission validity are miner-controlled after immutable assignment.
Miner submits an unsupported extra	Adjudication-quality deduction	The invalid addition is miner-attributable and distinct from missing coverage.
Validator fails Bronze, Silver, canonical audit, or required persistence for a paper	Paper excluded neutrally for every miner	A common validator-side failure cannot alter relative miner performance.
Some papers fail for validator reasons	Remaining papers score; batch marked degraded	Preserves valid work while exposing the incomplete validator result.
No paper remains scoring-eligible	Cycle fails	No comparable batch score exists.
Existing canonical assignment conflicts with local override	Validator refuses divergence	Maintains one batch and miner set for the assignment window.
No positive miner score	All-zero payout result	The mechanism does not create a winner without positive performance.
Audit-only mode	Scores and proposed weights persist; no on-chain submission	Separates replay and inspection from live incentives.

8.2 Minimum reconstructable record

A complete batch record contains the task and release identifiers; paper set and source hashes; selection seed; UID-to-hotkey snapshot; selection lane and history fields; assigned and responding miners; artifact identifiers and validation status; diagnostic findings and caps; Bronze identifiers; comparison cases and anonymized votes; deterministic consensus; canonical Silver units and importance labels; paper-level score factors; excluded papers and failure stages; batch means and ties; payout weights; and the on-chain submission event or audit-only marker.

This record permits a reviewer to reconstruct a payout from the final Silver units and findings rather than trusting the final number alone. Dynamic model outputs need not be reproducible byte for byte for the deterministic scoring calculation to be reproducible from the persisted adjudication record.

8.3 Security boundaries

Miner artifacts are untrusted model input. The validator keeps system instructions, source content, miner-authored fields, and output schemas separated; validates source references and identifiers before semantic review; uses origin-blind judge packets; restricts tool access; and

records failures rather than allowing malformed content to alter the evaluation process. Bittensor hotkeys authenticate network participants, while backend task and artifact records bind outputs to the canonical batch.

The mainnet configuration does not specify the role-specific model set. Model independence and provider diversity are therefore operator-supplied properties and are recorded in the persisted release metadata rather than treated as common mechanism constants.

9 Configuration Summary

This section collects the principal environment controls from the current mainnet configuration. A command-line argument overrides the corresponding environment value. Blank credential and model fields are intentionally omitted because they have no common configured value.

9.1 Identity, scheduling, and selection

Table 10. Identity, scheduling, and selection controls.

Variable	Value	Function
BT_NETUID	111	Bittensor subnet.
BT_SUBTENSOR_NETWORK	finney	Bittensor mainnet endpoint selection.
CLAIMS_NETWORK	mainnet	Claims backend partition.
CLAIMS_BACKEND_URL	https://api.claims111.ai	Canonical batches, assignments, artifacts, and persistence.
CLAIMS_BATCH_SIZE	50	Papers per batch.
CLAIMS_ALLOW_PAPER_REUSE	false	Excludes prior assigned papers in normal mainnet selection.
CLAIMS_MINER_SELECTION_MODE	adaptive	Enables seeded adaptive UID selection.
CLAIMS_MINER_SAMPLE_SIZE	10	Selected UIDs.
CLAIMS_MINER_IMMUNITY_PERIOD_BLOCKS	0	Reads the current subnet immunity period from chain.
CLAIMS_MINER_IMMUNITY_PRIORITY_BLOCKS	7,200	Priority window for under-vetted immune UIDs.
CLAIMS_TIMEOUT	3,600	Complete miner-batch deadline in seconds.
CLAIMS_BATCH_COLLECTION_POLL_SECONDS	5	Follower polling interval.
CLAIMS_BATCH_COLLECTION_WAIT_SECONDS	3,900	Maximum follower wait.
CLAIMS_MAX_STEPS	4	Completed scoring cycles per process run.
CLAIMS_QUERY_INTERVAL	21,600	Delay after a completed cycle before another begins.
CLAIMS_AUDIT_ONLY	false	Enables live weight submission.

9.2 Diagnostics, Bronze, and scoring

Table 11. Diagnostics, reference, and scoring controls.

Variable	Value	Function
CLAIMS_VALIDATOR_PIPELINE	auto	Current agent-v1 routing with compatibility support.
CLAIMS_AUDIT_METHOD	11m	Model-backed semantic diagnostics.
CLAIMS_AGENT_V1_VALIDATION_MODE	11m	Agent-v1 semantic validation.
CLAIMS_AGENT_V1_THRESHOLD	0.0	No positive diagnostic cut-off before Silver processing.
CLAIMS_BATCH_SCORE_RULE	mean	Compatibility setting; Silver batch incentives use the mean.
CLAIMS_PAYOUT_MODE	winner-takes-most	Concentrated rank weights.
CLAIMS_RIGOR_HARNESS	hermes-cli	Configured harness for the rigor role.
CLAIMS_RIGOR_MAX_TURNS	30	Rigor agent turn bound.
SUBNET_CLAIMS_VALIDATOR_AGENT_MAX_TOKENS	16,384	Rigor output-token bound.
CLAIMS_DIAGNOSTIC_MINER_MAX_WORKERS	2	Per-miner fallback concurrency.
CLAIMS_DIAGNOSTIC_MAX_WORKERS	10	Concurrent paper diagnostics.
CLAIMS_DIAGNOSTIC_MINER_BATCH_SIZE	10	One paper diagnostic can review all ten miners.
CLAIMS_REFERENCE_MINER_COMMAND	python-mclaims_reference_miner	Validator-side Bronze generation.
CLAIMS_REFERENCE_MINER_HARNESS	hermes-cli	Reference-agent harness.
CLAIMS_REFERENCE_PROFILE_ID	reference-agent-v1-hermes	Reference profile identifier.
CLAIMS_REFERENCE_MINER_PDF_READER	pdf-inspector	Reference PDF reader.

9.3 Silver and persistence

Table 12. Silver workflow and persistence controls.

Variable	Value	Function
CLAIMS_SILVER_ENABLE	true	Activates Bronze-to-Silver comparison and scoring.
CLAIMS_SILVER_PAIRING_TOP_K	4	Candidate neighbors per retrieval direction.
CLAIMS_SILVER_PAIRING_MAX_DENSE_PAIRS	0	Disables dense-pair addition.
CLAIMS_SILVER_PAIRING_EMBEDDING_MODE	openrouter	Embedding-provider mode; model remains operator-selected.
CLAIMS_SILVER_ADJUDICATION_CLI_TIMEOUT	1,800	Judge execution timeout.
CLAIMS_SILVER_ADJUDICATION_CLI_MAX_TURNS	30	Judge turn bound.
CLAIMS_SILVER_ADJUDICATION_MAX_IN_FLIGHT	32	Global simultaneous Silver calls.

Variable	Value	Function
CLAIMS_SILVER_ADJUDICATION_MAX_TOKENS	81,920	Judge output-token bound.
CLAIMS_SILVER_ADJUDICATION_BATCH_INPUT_TOKENS	120,000	Estimated input bound before splitting.
CLAIMS_SILVER_ADJUDICATION_BATCH_SIZE	13	Cases packed into a normal judge call.
CLAIMS_SILVER_ADJUDICATION_MAX_WORKERS	9	Adjudication worker bound.
CLAIMS_SILVER_WORKFLOW_MODE	file-agent	Current comparison, judging, and canonicalization path.
CLAIMS_SILVER_FILE_AGENT_MAX_TURNS	30	File-agent turn bound.
CLAIMS_SILVER_FILE_AGENT_MAX_TOKENS	81,920	File-agent output-token bound.
CLAIMS_SILVER_FILE_AGENT_TIMEOUT	1,800	File-agent execution timeout.
CLAIMS_SILVER_FILE_AGENT_REQUIRE_DISTINCT_JUDGES	true	Requires distinct direct judge roles.
CLAIMS_SILVER_FILE_AGENT_FALLBACK	none	No legacy fallback after file-agent failure.
CLAIMS_SILVER_PAPER_MAX_WORKERS	50	Concurrent paper-level Silver pipelines.
CLAIMS_SILVER_IMPORTANCE_BATCH_SIZE	8	Units grouped per importance call.
CLAIMS_SILVER_IMPORTANCE_MAX_TOKENS	81,920	Importance-role output-token bound.
CLAIMS_SILVER_MAX_ELIGIBLE_CLAIMS_PER_MINER	1,000	Per-miner Silver candidate cap.
CLAIMS_SILVER_MAX_ADJUDICATION_CASES_PER_PAPER	1,000	Per-paper case cap.
CLAIMS_SILVER_FILTER_BY_ASSESSMENT	false	Diagnostics do not automatically filter candidates.
CLAIMS_SILVER_PERSIST_CHUNK_SIZE	50	Ordinary record write chunk.
CLAIMS_SILVER_PERSIST_VOTE_CHUNK_SIZE	150	Vote write chunk.
CLAIMS_SILVER_CONSOLIDATION_TOP_K	4	Post-adjudication consolidation neighbors.

10 What the v0 Result Means

A high v0 score means that a miner contributed broad, source-grounded content to the machine-adjudicated Silver record while avoiding structural defects and unsupported additions. It does not establish that every Silver unit is true, complete, or independently verified. Bronze and Silver can share model errors, and origin-blind machine adjudication can resolve a dispute incorrectly.

The active mechanism is therefore a competitive machine-curation system with an inspectable scoring record. It is not yet a full Bronze–Silver–Gold quality flywheel. In particular, no maintained Gold layer currently calibrates the judge, no hidden known-answer challenge measures reliability against independent truth, and no production ingestion status represents a record as human-verified. These limits define the correct interpretation of the current v0 output.

11 Conclusion

The current SN111 mechanism has a compact operating contract. A canonical 50-paper batch is assigned to ten adaptively selected miners. Miners have one hour to return agent-v1 artifacts. The validator runs deterministic and semantic diagnostics, creates a Bronze reference, constructs an origin-blind Silver record through two direct judges and a conditional tiebreaker, and scores miners by the product of coverage, diagnostic quality, and adjudication quality. Paper scores are averaged, and positive miners receive concentrated 70/16/8/4/2 rank weights with explicit tie sharing. The standard process attempts four cycles with a six-hour delay after each completed cycle. On chain, SN111 has 256 UID capacity, 64 validator capacity, a 360-block tempo, and a 21,600-block immunity period. These settings define the present v0 operating contract; Gold-based assurance and the broader production architecture remain outside it.