GenLayer Tokenomics

Design & Rationale

GenLayer Foundation

October 2026

This document is for technical discussion only and does not constitute an offer to sell or solicitation of an offer to buy any GEN token or other security in any jurisdiction. It is not legal, tax, financial, or investment advice, and no one should rely on it in making any purchase or investment decision. It contains forward-looking statements about protocol design, parameters, and governance; these reflect current intentions, are subject to change, and may not be realized. Participating in the GenLayer network, including staking, delegation, and contract deployment, involves risks, among them the loss of staked tokens through slashing or penalties, software defects, parameter changes adopted through governance, and regulatory change. GEN may not be available to persons in certain jurisdictions, and readers are responsible for complying with the laws that apply to them.

Abstract

GenLayer is a synthetic jurisdiction: an autonomous digital court in which Intelligent Contracts interpret code and natural language, read external data, and reach decisions without routine human adjudication. Validators running their own AI configurations decide in sampled panels by majority vote, with appeals to larger panels.

This document specifies the GEN token that funds and secures that court: its supply, allocation, vesting, issuance, rewards, staking, and penalties. The genesis supply and the issuance ceiling are committed and will be locked progressively until they are immutable on-chain (Section 9); all other parameters are set by governance.

Contents

Abstract **1**

1 Genesis Supply, Allocation, and Vesting **3**
1.1 Staking status of allocations 3
1.2 Unlock path 3

2 Inflation Schedule **4**
2.1 Effective inflation rate 4

3 Time-Unit Pricing **5**

4 Reward Routing **5**

5 Developer Incentives **5**

6 Staking Model **6**
6.1 Delegation 6

7 Weight Function **6**

8 Penalties, Slashing & Appeal Economics **7**
8.1 Appeal flow 7
8.2 Validator bans 7

9 Governance **8**

A Round Sizes **8**

B References with Links **9**

1 Genesis Supply, Allocation, and Vesting

Total genesis supply: $S_0 = 4.20\,\text{B}$ GEN, committed (Section 9). It is allocated across investors, **GenLayer Labs** (the development company), the **GenLayer Foundation** (the ecosystem steward), the community, and the DeepThought DAO. Every schedule below is a step function: an optional unlock at TGE, an optional one-step cliff unlock, then equal semi-annual tranches.

Bucket	GEN (M)	Share	TGE	Cliff (month 12)	Tranches	Vested
Investors	558.30	13.2929%	0%	33.3%	4× ≈16.7%	Month 36
Early contributors	800.68	19.0638%	0%	25%	6× 12.5%	Month 48
GenLayer Labs	741.02	17.6433%	0%	33.3%	4× ≈16.7%	Month 36
GenLayer Foundation	840.00	20.0000%	100%	–	–	TGE
Community	1,050.00	25.0000%				
Growth	819.00	19.5000%	25%	–	8× 9.375%	Month 48
Early ecosystem	231.00	5.5000%	100%	–	–	TGE
DeepThought DAO	210.00	5.0000%	0%	month 48†	semi-annual†	–

† Locked until DAO deployment; equal semi-annual tranches follow the month-48 cliff.

Table 1: Genesis allocation and vesting (October 2026). Tranches are equal and semi-annual; "4× ≈16.7%" means four tranches of ≈16.7% each.

1.1 Staking status of allocations

Vesting restricts transfer, not staking. Every allocation except the DeepThought DAO's 210 M can be staked or delegated from genesis, vested or not; staking Growth tokens leaves them reserved for ecosystem programs and does not change their vesting.

Self-stake commitment. GenLayer Labs and the GenLayer Foundation have committed to self-staking at most 10% of their allocations, team allocations included, with the remainder of any staked balance delegated exclusively to independent community validators.

1.2 Unlock path

The 33.3% cliffs are exact one-third unlocks. The DAO's tranche schedule after month 48 is not yet specified, so its row below marks eventual full vesting rather than a date. At TGE, the Foundation allocation (840 M), the Growth bucket's first 25% (204.75 M), and the early-ecosystem allocation (231 M) unlock; everything else is locked.

Month	Newly unlocked (M)	Cumulative (M)	% of S_0
0 (TGE)	1,275.8	1,275.8	30.4%
6	76.8	1,352.5	32.2%
12	710.1	2,062.6	49.1%
18	393.4	2,456.0	58.5%
24	393.4	2,849.4	67.8%
30	393.4	3,242.8	77.2%
36	393.4	3,636.3	86.6%
42	176.9	3,813.1	90.8%
48	176.9	3,990.0	95.0%
>48	210.0 (DAO)	4,200.0	100%

Table 2: Transferable supply over time (excludes issuance; figures rounded to 0.1 M).

TGE is the largest unlock (1,275.75 M GEN, 30.375% of S_0); month 12 is the next largest

($\approx 16.9\%$), when the investor, early-contributor, and Labs cliffs coincide. By month 48, 95% is transferable.

Team allocations inside the Foundation bucket (252 M) carry private lockups that mirror the early-contributor schedule, so the effective float at TGE is $\approx$1,023.75 M GEN (24.4% of S_0).

2 Inflation Schedule

GenLayer mints new GEN using an **inverse-sigmoid curve** defined on the genesis supply S_0:

$$r(d) = r_\infty + \frac{r_0 - r_\infty}{1 + \exp\big(k\,(d - d_{\text{mid}})\big)}$$

Symbol	Value	Description	Status
r_0	0.09	Upper asymptote; $r(0) \approx 8.9344\%$	committed
r_∞	0.04	4% tail rate as $d \to \infty$	committed
d_{mid}	1,440 days ($\approx$3.95 y)	midpoint: $r(d_{\text{mid}}) = 6.5\%$	committed
k	0.003 per day	steepness coefficient	committed

Table 3: Parameters of the scheduled issuance rate. The schedule as a whole is the issuance ceiling. "Committed" means fixed in the contracts and changeable only through a contract upgrade, locked progressively toward on-chain immutability (Section 9).

2.1 Effective inflation rate

Daily issuance converts the annual rate on a 360-day year and always applies it to the fixed genesis base:

$$M_d = S_0 \left((1 + r(d))^{1/360} - 1\right).$$

Nominal issuance is therefore predictable, and inflation measured on current supply falls as supply grows:

$$R_{\text{infl}}(d) = \frac{365\,M_d}{S_d}, \qquad S_d = S_0 + \textstyle\sum_{u<d} M_u.$$

Year	$r(d)$, nominal	R_{infl}, effective (on current supply)
0	8.93%	8.68%
2	8.47%	7.04%
4	6.43%	4.79%
6	4.48%	3.12%
10	4.01%	2.51%

Table 4: Gross scheduled issuance path. Reporting years contain 365 days; the daily-rate conversion uses 360. Figures are rounded and exclude burns.

The schedule is a ceiling: realized supply growth can only be lower, through penalty burns and the burned developer-inflation surplus (Section 5). Gross issuance over the first 365 days is $\approx$362.42 M GEN, roughly 28.4% of the 1,275.75 M protocol-transferable day-one float. Of scheduled issuance, 90% is allocated to staking participants and up to 10% to developers; unused developer issuance is burned.

3 Time-Unit Pricing

GenVM execution is funded through time budgets and paid in GEN. GenLayer settles on an underlying *base chain*; the *base-chain fee* is the fee that chain charges for posting GenLayer data. All other fees are *over-layer fees*, paid for execution and consensus on GenLayer itself.

- **Time Unit (TU).** 1 TU represents one second of allocated execution time. Users fund a time budget for each *consensus slot*, meaning one validator seat (leader or validator) in one consensus round; fees are based on these allocations, not measured runtime.
- `GENPerTimeUnit`. The GEN price of one TU. A single parameter, set and adjusted only by governance; there is no automatic adjustment.
- **Escrow.** Senders pre-escrow GEN for execution and for any appeal rounds and rotations (re-runs with a new leader) they choose to pre-pay. Result storage and receipt posting are priced separately; like the base-chain fee, they are cost reimbursements outside the routing keys.
- **Settlement.** Budgets for used consensus slots are routed per Section 4; unused slots are refunded to whoever funded them.

4 Reward Routing

Fee keys apply to the over-layer fees consumed in each epoch, after refunds; inflation keys apply to scheduled issuance. The DeepThought DAO, a planned AI-governed DAO, is funded only from fees, so its income tracks network usage, never issuance.

Reward source	Staker pool	Validator pool	Developers	DeepThought DAO
Fees	70%	15%	10%	5%
Inflation	75%	15%	$\leq$ 10%*	0%

Table 5: Reward routing. *The developer share of inflation is capped by usage (Section 5). The fee-side validator and DAO shares are governance-adjustable; the table shows launch values.

Staking inflation is allocated to validator positions by weight (Section 7); fee rewards go to the validators that performed the work. In both streams, each position's staker reward is then split between operator and delegators (Section 6). The validator pools pay for node operation and go to the operator's self-stake, not to delegators: each operator receives a share of the fees it collects, and validator inflation is divided evenly among active validators to cover fixed costs.

The active set is the top N eligible validators by weight, with $N = 1{,}543$ at launch; at each epoch's update, a heavier candidate displaces the lowest-weight active validator.

5 Developer Incentives

Developers earn from both rows of the Developers column in Table 5, and both depend on the fees their contracts generate:

- **Fee share.** 10% of every fee paid to a contract goes to its developer.
- **Inflation top-up.** Each epoch, up to 10% of inflation is divided among developers in proportion to the fees their contracts generated. No developer receives more than its fee share for that epoch, so the top-up can at most match it; whatever the cap leaves unclaimed is burned.

Each contract has its own transferable **Developer NFT**, and both streams accrue to whoever holds it; transferring the NFT transfers the right to claim them.

Wash trading does not pay: a developer who generates transactions on its own contract recovers at most 20% of the fees through these two streams, and the rest goes to pools shared with every other staker and validator.

6 Staking Model

Staking parameters are denominated in GEN and governance-adjustable.

Item	Value	Purpose
Validator self-stake minimum	42,000 GEN	Skin in the game; Sybil cost
Delegator minimum	42 GEN	Prevents micro-dust attacks
Active-set size	1,543 validators	Top validators by weight; also the electorate of the final post-appeal round
Reward epoch	1 day	Balances payout latency and compounding
Unbonding delay	7 epochs (7 days), self and delegated	Shields the network during exit

Table 6: Staking parameters.

A Verifiable Random Function samples each committee from the active set in proportion to weight, without replacement within a draw. A successful appeal is followed by a single, final round drawn from the whole active set (Appendix A).

6.1 Delegation

Delegators stake at least 42 GEN with any validator; rewards compound automatically, and there are no commissions.

A position's staker rewards R are split by the weight each side contributes. With W the position's weight (Section 7) and $W_0 = \sqrt{0.6 \times \text{Self Stake}}$ its self-only weight,

$$R_{\text{operator}} = R\frac{W_0}{W}, \qquad R_{\text{delegators}} = R - R_{\text{operator}},$$

and delegators share their part pro rata. The operator's share of staking inflation therefore depends on its own self-stake and total network weight, not on how much delegation it attracts. Delegation still pays the operator through fees: more weight means more committee seats.

Delegators earn without running infrastructure or meeting the validator minimum, and bear only the delegated share of a slash. An unstaked position earns through the current epoch and is withdrawable after the 7-epoch unbonding period, during which it earns nothing.

7 Weight Function

The weight used for validator selection and position-level staker-pool allocation is

$$W = \sqrt{0.6 \times \text{Self Stake} + 0.4 \times \text{Delegated Stake}}\,.$$

Self-stake enters with a higher coefficient than delegation and is slashed at a higher rate. The square root dampens concentration: doubling stake raises weight by about 41%. Splitting stake across identities would raise total weight, but each identity needs the 42,000 GEN minimum, its own infrastructure, separate slashing exposure, and a place among the top N by weight.

8 Penalties, Slashing & Appeal Economics

Only idleness and deterministic violations are slashed as a percentage of principal; ordinary model error is not. A validator on the minority side of a final result forfeits its fee for the task and pays a fee-sized penalty (equal to that fee), burned from the operator's self-stake, never from delegated stake. A slash at base rate b removes $0.8\,b$ of the operator's self-stake and $0.2\,b$ of the delegated stake.

Trigger	Base rate / penalty & destination	Rate weights (self/del.)
Minority-aligned result (votes with the losing side after any appeal)	Fee forfeited, plus a fee-sized penalty burned (debited from the operator's self-stake); no percentage slash	—
Idleness / timeout	1% base rate; 1% of the slashed amount to the caller, remainder burned	80% / 20%
Deterministic violation (provable fraud)	Base rates: leader 5%, other signers 1%; 1% of the slashed amount to the caller, remainder burned	80% / 20%
Failed appeal	100% of appeal bond, paid to the majority-aligned validators of the round that decides the appeal; a failed leader-timeout appeal pays 50% to the leader and refunds 50%	

Table 7: Penalty schedule.

Votes use commit–reveal, so no validator sees another's vote before committing its own. No penalty goes to a discretionary treasury. Slashes take effect two epochs after the offence and are capped at a 10% base rate per epoch of offence (at most 8% of self-stake and 2% of delegation); stake withdrawn after an offence remains liable while it unbonds.

8.1 Appeal flow

- **Window**: a governance-set duration that opens when a result enters Accepted or Undetermined and shrinks with each successive appeal.
- **Bond**: the appellant posts a bond B and separately pays the fees for the appeal's work, unless the sender pre-funded the appeal. A successful appellant receives $2.5B$: the bond back plus a $1.5B$ reward.
- **Escalation**: while appeals fail, round sizes follow the schedule in Appendix A, limited by the validators not yet used in the transaction.
- **Successful appeal**: the jury reaches a different outcome. A single, final round drawn from the whole active set follows, and its decision cannot be appealed.
- **Failed appeal**: the bond is paid to the majority-aligned validators of the round that decides the appeal, except that a failed leader-timeout appeal pays 50% to the leader and refunds 50%.

8.2 Validator bans

Three idleness strikes within an epoch suspend a validator for the current and following epoch. A leader accused of forging a result is quarantined until the on-chain tribunal rules: removed if it finds forgery, reinstated if not.

9 Governance

The issuance schedule and the genesis supply are *committed*: fixed in the contracts with no governance setter, changeable only through a contract upgrade, and to be locked progressively until immutable on-chain. All other parameters (staking minimums, reward shares, penalties, the appeal depth and window, `GENPerTimeUnit`) are set by on-chain governance.

Until token voting begins, an operational authority makes parameter changes and a separate emergency authority acts only in a crisis.

A Round Sizes

Consensus alternates normal rounds and appeal juries. The nominal round sizes are

$$S(N) = 3 \cdot 2^{\lfloor (N-1)/2 \rfloor + 1} + (-1)^N, \qquad N = 1, \ldots, 18,$$

odd N for normal rounds and even N for appeals (Table 8). Sizes are odd, so there are no ties, and roughly double every two rounds, so each escalation costs more than the last. Size 5 is the smallest panel that cannot be captured by two seats. A jury is limited to the validators not yet used in the transaction, so realized sizes can fall below the schedule.

A successful appeal ends the sequence. The final round's committee is the active set at the time the transaction was activated, less every validator that has already led a normal round of the transaction; its quorum is a strict majority of that active set ($\lfloor 1543/2 \rfloor + 1 = 772$ votes when it is full), and its decision cannot be appealed.

Round N	Validators $S(N)$	Type
1	5	Normal
2	7	Appeal
3	11	Normal
4	13	Appeal
5	23	Normal
6	25	Appeal
7	47	Normal
8	49	Appeal
9	95	Normal
10	97	Appeal
11	191	Normal
12	193	Appeal
13	383	Normal
14	385	Appeal
15	767	Normal
16	769	Appeal
17	1535	Normal
18	1537	Appeal

Table 8: Nominal round-size schedule.

B References with Links

Ref	Name (click to open)	Main takeaway for GenLayer
A	Amihud & Mendelson 1986: "Asset Pricing and the Bid-Ask Spread"	Shows that transaction costs and liquidity affect asset pricing; background for publishing a transparent, predictable issuance schedule.
B	Bonello 2018: "Conflict: Store of Value vs Utility"	Argues that store-of-value properties (scarcity, predictability) and the flexibility a utility platform needs pull against each other; background for prioritizing network usage over scarcity.
C	Buterin & Griffith 2017: "Casper the Friendly Finality Gadget"	Burn-based slashing offers strong deterrence without cosmetic deflation; mirrored in GenLayer penalties.
D	Catalini & Gans 2018: "Initial Coin Offerings and the Value of Crypto Tokens"	Analyzes token-based financing and monetary commitments; background for tying developer rewards to platform usage.
E	Chen & Micali 2019: "Algorand"	Studies cryptographic committee selection and Byzantine agreement; background for sampled-committee design.
F	Cong, He & Tang 2025: "The Tokenomics of Staking"	Models how staking interacts with platform activity; background for the staking-reward design.
G	Cosmos SDK Docs: x/mint module	Production example of parameter-controlled inflation in proof-of-stake chains (x/mint adjusts inflation toward a target bonded ratio rather than following a time-based schedule).
H	Electric Capital: Developer Report 2024	Tracks developer activity across blockchain ecosystems; background for funding developer incentives.
I	Helius: "Solana Validator Economics: A Primer"	Reports hardware costs of roughly \$350–\$470/month plus about 1 SOL/day in vote fees for a Solana validator; benchmarks validator cost assumptions.
J	Lo & Medda 2020: "Assets on the blockchain: An empirical study of Tokenomics"	Empirical study of how token design features relate to market outcomes; background for tokenomics design choices.
K	MakerDAO Docs: Emergency Shutdown Module	A pre-defined emergency mechanism, triggered when enough MKR is deposited, that can shut the system down quickly; precedent for a crisis-response path separate from routine governance votes.
L	Merton 1987: "A Simple Model of Capital Market Equilibrium with Incomplete Information"	Models markets in which investors know about only some assets; background for broad, transparent token distribution.
M	Polkadot OpenGov Referendum 1139 (2024)	A 2024 non-binding "Wish for Change" vote that passed, signalling support for lowering Polkadot's inflation; shows that governance-mutable issuance attracts change proposals, the contrast motivating GenLayer's committed path to an immutable ceiling.
N	Buterin 2021: "Moving beyond coin voting governance"	Discusses the limitations of coin voting and alternatives that dampen large holders; background for square-root damping of stake weight.