RESEARCH / REPORT 002

<< ON-CHAIN AUDIT · JUNE 2026 >>

The Validator Hierarchy

A 365-day audit of the top 25 Bittensor root validators, ranked by what nominators actually earned on root rather than what they advertise. The gap between best and worst is 4.41 points of yield. Take rate barely explains it.

FLOWSNIPER RESEARCH · JUNE 12, 2026 · 25 VALIDATORS · 129 SUBNETS · 365 DAYS · 860,642 SNAPSHOTS

If you stake TAO on Bittensor, you made a choice you probably didn't think hard about: which validator to delegate to. Most people pick on one number, the advertised take rate, the cut the validator keeps. Lower take, better deal. It feels obvious.

It is also mostly wrong. We pulled a full year of daily on-chain yield data for the 25 largest commercial validators on root, ranked them by what nominators actually earned after fees, and found that the validator you picked matters far more than the take rate suggests, and in a direction the take rate often hides.

> *The best and worst major validators are 4.41 points of yield apart. Take rate explains less than half of it.*

8.86%	4.45%	4.41pp	$44–$2,205
Best validator realized root APY (90d)	Worst validator realized root APY (90d)	Spread between best and worst	Annual cost of the worst validator, $1k to $50k staked

That last number is the one to sit with. Pick the worst-performing validator in this group instead of the best and you give up, every year, between $44 on a $1,000 stake and $2,205 on a $50,000 stake. Not to a fee you agreed to. To a difference in performance you were never shown, because nobody publishes it.

Use the buttons in the scorecard below to see the cost at your own stake size.

What we measured, and why it is not the take rate

Every commercial validator advertises a take rate: 0%, 9%, 18%. The pitch is simple. We keep X percent of the emissions, you keep the rest. Pick a low-take validator and you keep more. Pick a high-take validator and you keep less.

The problem is that take rate is only one of the things standing between you and your yield. A validator also has to actually perform: set good weights, stay online, earn its share of emissions from the network. A validator with a 0% take but mediocre performance can pay you less than a validator with a 9% take and excellent performance. The take rate is the part you see. The performance is the part you don't.

So we ignored the advertised numbers and measured the only thing that matters to a nominator: realized return. For each validator, the actual daily nominator APY on the root network, net of take, averaged across the window. This is what your stake actually grew by. It already includes the take. It already includes the performance. It is the bottom line.

DISCLOSURE

FlowSniper delegates by default to tao.bot (`5E2LP6E...d6ddd748`), which appears in this report at rank 2. We have no commercial relationship with tao.bot, no communication channel with their team, and they are likely unaware FlowSniper exists. We selected tao.bot as our default before this analysis began, based on the same publicly observable characteristics any nominator can verify: 0% take, broad subnet coverage, and consistent historical performance. Had the data ranked tao.bot poorly, we would have published it anyway and switched. That it ranks well is consistent with why we chose it, not a consequence of our choosing it.

⌖ Root staking versus subnet staking: what this ranks

Bittensor gives you two distinct decisions, and this report is about the first one.

Root staking means delegating TAO to a validator on the root network (netuid 0). Your stake stays denominated in TAO, you earn a TAO-denominated yield, and the only real variable is which validator you pick. This is the lower-risk, yield-focused choice most stakers default to, and it is what the scorecard in this report ranks. Because every commercial validator validates root identically, comparing them here is apples to apples: same network, same asset, same mechanics, only the operator differs.

Subnet staking means converting TAO into a subnet's alpha token and staking that. Here your return is dominated by the alpha token's price movement, not the validator's yield, which is exactly the trap our first report, The Yield Illusion, documented: the median subnet alpha staker lost money in TAO terms last month despite headline APYs above 50%. Validator choice still matters on subnets, but it is a second-order effect swamped by price.

> *This report ranks the root staking decision, where the validator is the whole game. If you are staking alpha on a subnet, read the Yield Illusion first, because the validator is the least of your risks there.*

Every figure below is root-only (netuid 0) unless explicitly stated. We pulled subnet-level yield for all 25 validators across 128 subnets as well, and the broad-performance picture is consistent with the root ranking, but we lead with root because it is the cleanest, most controlled, and most widely-used comparison. The cross-subnet breakdown is reserved for a follow-up.

⌖ The scorecard

Here are all 25 validators, ranked by realized root APY over the trailing 90 days. The `gap` column is how far behind the best validator each one fell, in percentage points of yield. The `cost/yr` column converts that gap into dollars on a $10,000 stake.

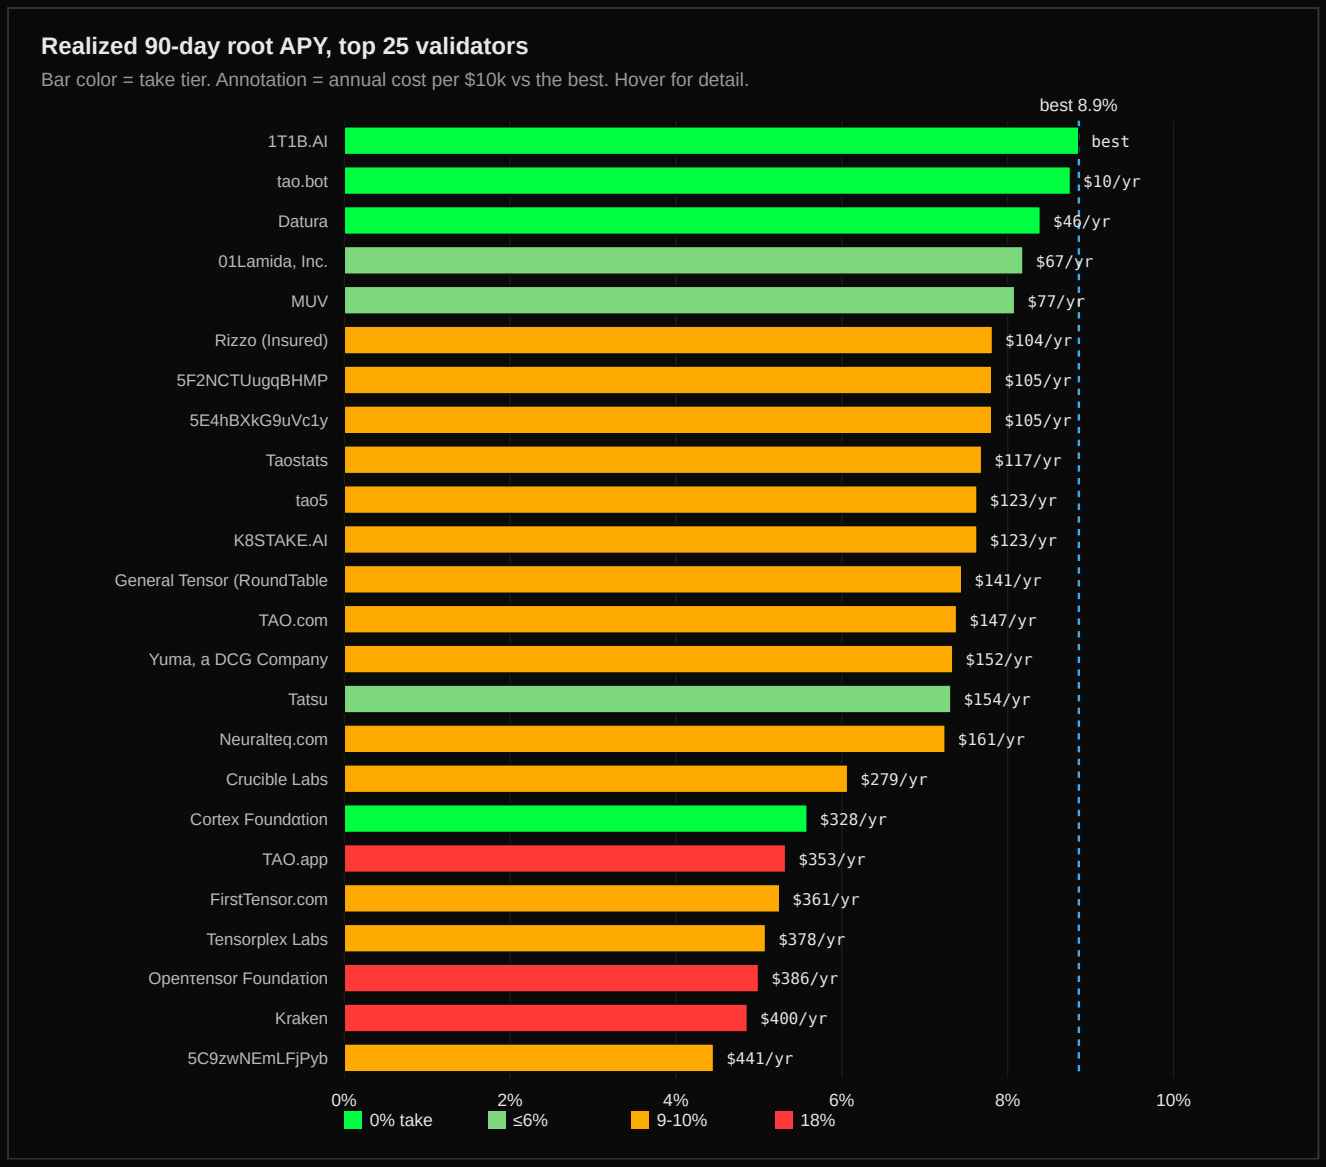

Figure 1. Realized 90-day root APY for the top 25 commercial validators. Bar color shows take rate tier. The annotation on each bar is the annual cost per $10k versus the best validator. Hover any bar for detail.

#	VALIDATOR	TAKE	90D ROOT APY	GAP (PP)	COST / $10K
1	1T1B.AI	0%	8.86%	+0.00	$0/yr
2	tao.bot	0%	8.76%	+0.10	$10/yr
3	Datura	0%	8.39%	+0.46	$46/yr
4	01Lamida, Inc.	2.1%	8.19%	+0.67	$67/yr
5	MUV	6%	8.08%	+0.77	$77/yr
6	Rizzo (Insured)	9%	7.82%	+1.04	$104/yr
7	5F2NCTUugqBHMP	9%	7.81%	+1.05	$105/yr
8	5E4hBXkG9uVc1y	9%	7.81%	+1.05	$105/yr
9	Taostats	9%	7.69%	+1.17	$117/yr
10	tao5	9%	7.63%	+1.23	$123/yr
11	K8STAKE.AI	10%	7.63%	+1.23	$123/yr
12	General Tensor (RoundTable21)	9%	7.45%	+1.41	$141/yr
13	TAO.com	9%	7.38%	+1.47	$147/yr
14	Yuma, a DCG Company	9%	7.34%	+1.52	$152/yr
15	Tatsu	5%	7.32%	+1.54	$154/yr
16	Neuralteq.com	9%	7.25%	+1.61	$161/yr
17	Crucible Labs	9%	6.07%	+2.79	$279/yr
18	Cortex Foundαtion	0%	5.58%	+3.28	$328/yr
19	TAO.app	18%	5.32%	+3.53	$353/yr
20	FirstTensor.com	9%	5.25%	+3.61	$361/yr
21	Tensorplex Labs	9%	5.08%	+3.78	$378/yr

#	VALIDATOR	TAKE	90D ROOT APY	GAP (PP)	COST / $10K
22	Opentensor Foundation	18%	5.00%	+3.86	$386/yr
23	Kraken	18%	4.86%	+4.00	$400/yr
24	5C9zwNEmLFjPyb	9%	4.45%	+4.41	$441/yr

Ary van der Touw is excluded from the headline ranking: at 28 days of history it has insufficient data for a 90-day window. It appears in the full dataset with that flag. All other validators have 89 days of root coverage in the 90-day window.

The same gap, at the stake sizes most nominators actually hold. The bottom five validators, and what staying with them costs you per year against the best in class:

VALIDATOR (BOTTOM 5)	90D ROOT APY	$1K	$5K	$10K	$50K
5C9zwNEmLFjPyb	4.45%	$44	$220	$441	$2,205
Kraken	4.86%	$40	$200	$400	$2,000
Opentensor Foundation	5.00%	$39	$193	$386	$1,930
Tensorplex Labs	5.08%	$38	$189	$378	$1,890
FirstTensor.com	5.25%	$36	$180	$361	$1,805

⌖ Take rate is a weak predictor of what you keep

If take rate told the whole story, the ranking above would sort cleanly by fee: every 0% validator on top, every 18% validator on the bottom, the 9% validators stacked neatly in the middle. It does not look like that.

Look at the validators charging the same 9% take. They span from 7.82% realized APY at the top to 4.45% at the bottom, a spread of 3.36 points. Same fee. A $337 per year difference per $10,000 staked, between validators that advertise the identical headline

number. The fee you were quoted told you almost nothing about what you would actually earn.

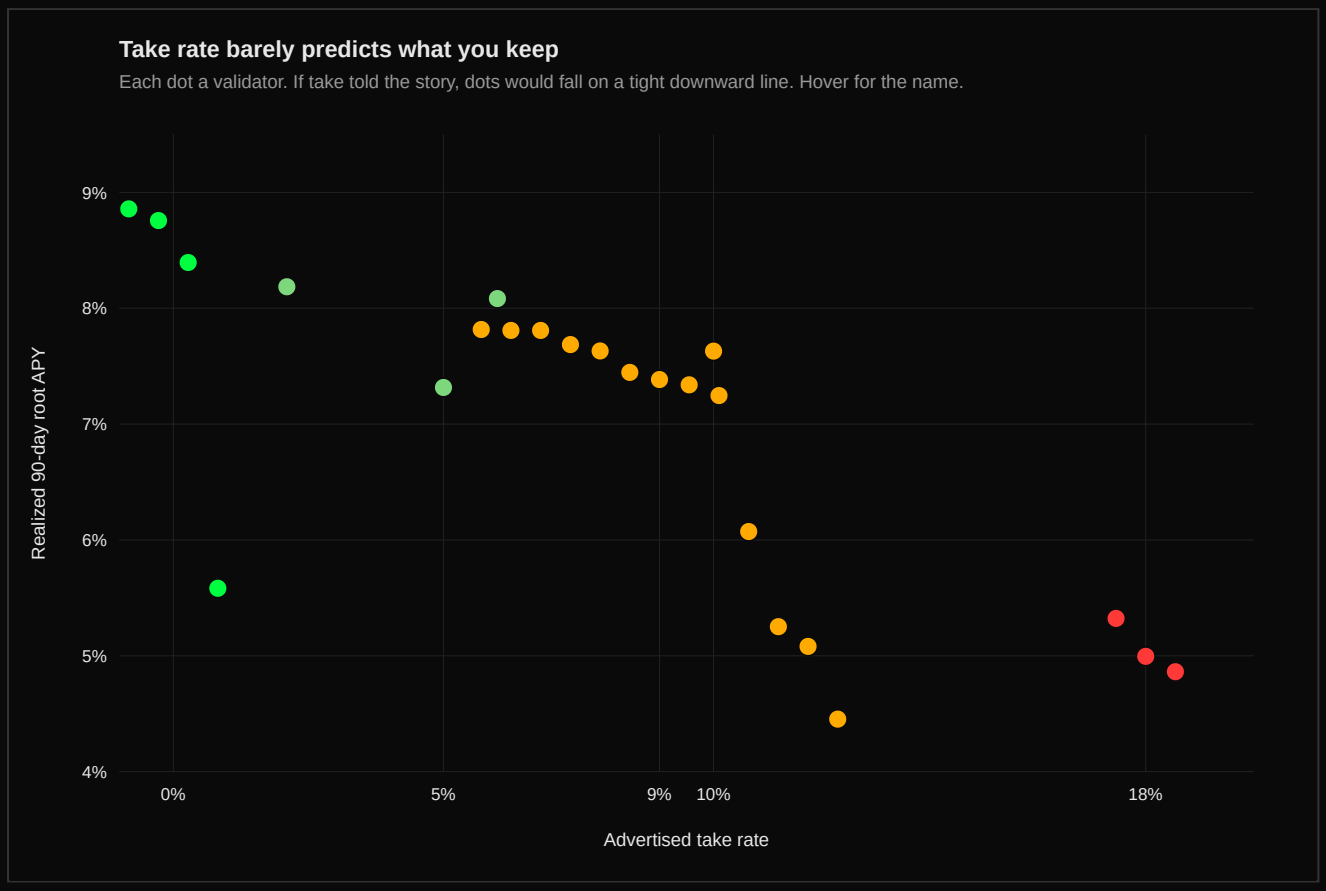

Figure 2. Advertised take rate against realized 90-day APY. If take predicted return, the points would fall along a tight downward line. They do not. The vertical spread within each take tier is the performance gap nobody advertises. Hover any point for the validator.

The most striking case is at 0% take. Four validators charge nothing at all. They range from 8.86% realized APY (1T1B.AI, the best in the entire study) down to 5.58% (Cortex Foundation, which finishes 18th of 24). Both are free. One earns you $328 a year more per $10k than the other.

> *Free is not the same as best. A 0% take validator that performs poorly costs you more than a 9% take validator that performs well.*

Figure 3. The range of realized returns within each take-rate tier. The width of each band is the performance spread among validators charging identical fees. The 0% and 9% tiers each span more than three points of yield. Hover for the validators at each end.

⌖ Where the gap actually comes from

To separate the two forces, we decomposed each validator's shortfall against the best into two components. The `take penalty` is how much of the gap comes from charging a higher fee than the leader. The `skill penalty` is how much comes from earning less gross yield in the first place, before any fee, a function of weight-setting, uptime, and how the network rewarded the validator's work.

To get there we back out each validator's gross APY from its net: if a validator nets 7.39% after a 9% take, its gross yield before the take was about 8.11%. Comparing gross yields strips the fee out entirely and shows pure performance.

#	VALIDATOR	TAKE	TOTAL GAP	TAKE PEN.	SKILL PEN.	GROSS APY
2	tao.bot	0%	+0.10	+0.00	+0.10	8.76%
3	Datura	0%	+0.46	+0.00	+0.46	8.39%
4	01Lamida, Inc.	2.1%	+0.67	+0.19	+0.49	8.36%
5	MUV	6%	+0.77	+0.53	+0.24	8.60%
6	Rizzo (Insured)	9%	+1.04	+0.80	+0.24	8.59%
7	5F2NCTUugqBHMP	9%	+1.05	+0.80	+0.25	8.58%
8	5E4hBXkG9uVc1y	9%	+1.05	+0.80	+0.25	8.58%
9	Taostats	9%	+1.17	+0.80	+0.37	8.45%
10	tao5	9%	+1.23	+0.80	+0.43	8.38%
11	K8STAKE.AI	10%	+1.23	+0.89	+0.34	8.48%
12	General Tensor (RoundTable21)	9%	+1.41	+0.80	+0.61	8.18%
13	TAO.com	9%	+1.47	+0.80	+0.68	8.11%
14	Yuma, a DCG Company	9%	+1.52	+0.80	+0.72	8.06%
15	Tatsu	5%	+1.54	+0.44	+1.10	7.70%
16	Neuralteq.com	9%	+1.61	+0.80	+0.81	7.96%
17	Crucible Labs	9%	+2.79	+0.80	+1.99	6.67%
18	Cortex Foundation	0%	+3.28	+0.00	+3.28	5.58%
19	TAO.app	18%	+3.53	+1.59	+1.94	6.49%
20	FirstTensor.com	9%	+3.61	+0.80	+2.81	5.77%
21	Tensorplex Labs	9%	+3.78	+0.80	+2.98	5.58%
22	Opentensor Foundation	18%	+3.86	+1.59	+2.27	6.09%

#	VALIDATOR	TAKE	TOTAL GAP	TAKE PEN.	SKILL PEN.	GROSS APY
23	Kraken	18%	+4.00	+1.59	+2.40	5.93%
24	5C9zwNEmLFjPyb	9%	+4.41	+0.80	+3.61	4.89%

The pattern is clear once you look. Near the top, the gaps are small and mostly fee-driven: Taostats, tao5, Yuma and the other established 9% validators trail the leaders by roughly a point, and most of that point is simply the 9% they charge. Their gross performance is excellent, within a few tenths of the best.

Near the bottom, the gaps are large and mostly skill-driven. Tensorplex Labs charges the same 9% as Taostats but trails the best by 3.78 points, and 2.98 of those points, nearly four fifths, come from earning less gross yield, not from the fee. FirstTensor.com tells the same story: 9% take, but a 2.81-point skill penalty. These validators are not expensive. They are underperforming, and the fee is a footnote next to the performance gap.

Cortex Foundαtion is the purest illustration. It charges 0%. Its entire 3.28-point gap is skill penalty, because there is no fee to blame. A nominator who chose Cortex specifically because it was free earned $328 a year less per $10k than they would have with 1T1B.AI, also free.

⌖ Is this just a bad month? The ranking is stable

The obvious objection to any single-window ranking is that we picked a window that happened to flatter some validators and punish others. So we ran the same analysis across six windows, from 7 days to a full year, and looked at how each validator's rank moved.

The extremes barely move. 1T1B.AI is first or second in every window measured. tao.bot is first or second in every window. At the bottom, Kraken sits at 23rd or 24th in all six windows, and 5C9zwNEmLFjPyb sits at 22nd to 25th in all six. These are not

artifacts of a lucky month. They are durable differences in performance that persist across the entire year.

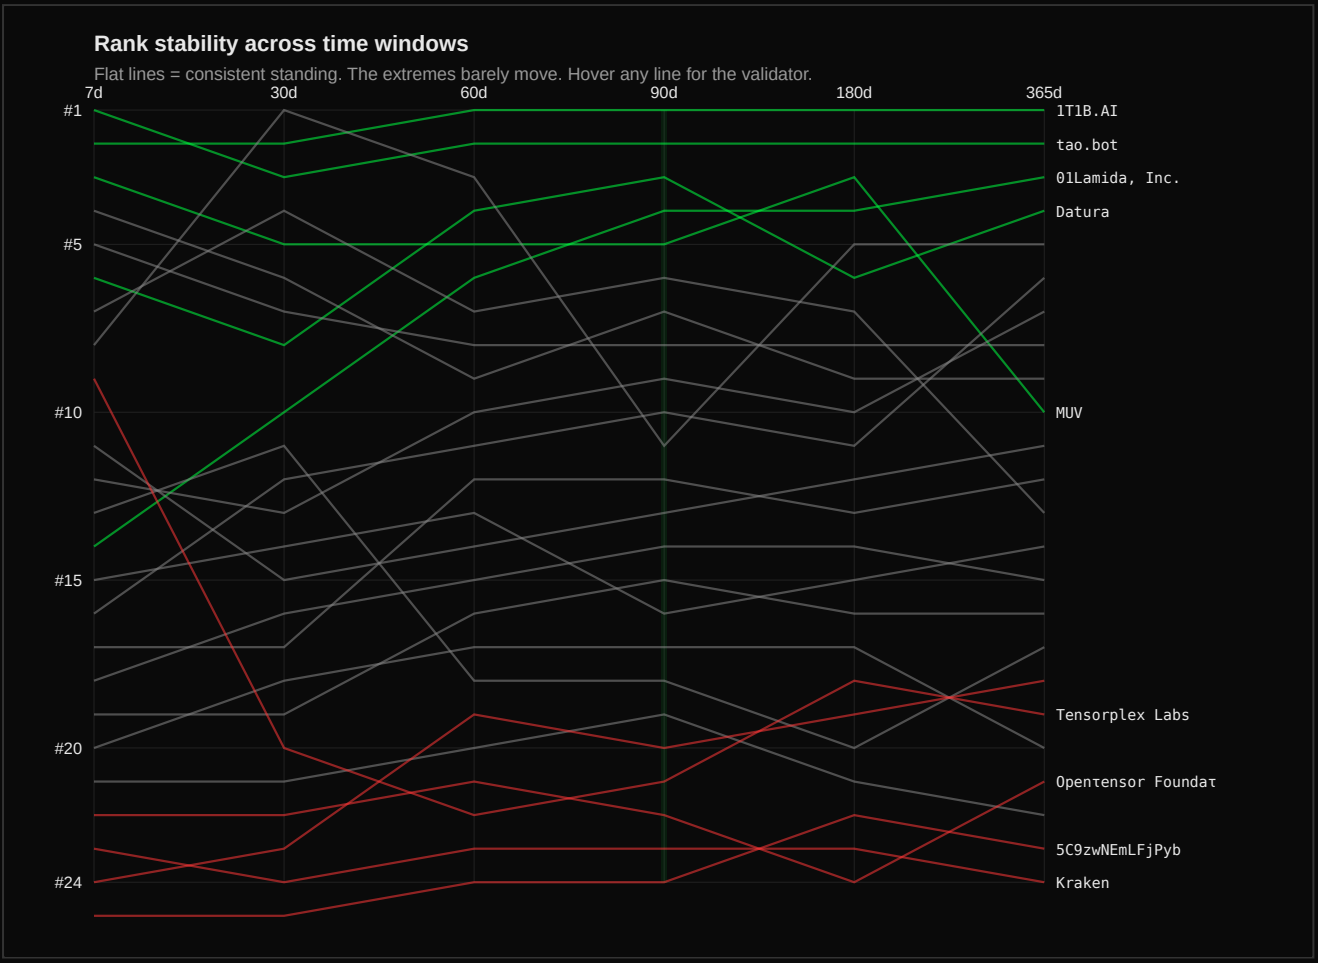

Figure 4. Each validator's rank at 7, 30, 60, 90, 180, and 365 days. Flat lines are validators whose standing is consistent. The 90-day column, used for the headline ranking, is highlighted. Hover any line for the validator's full rank history.

The middle of the pack does shuffle. K8STAKE.AI was first at 30 days and eleventh at 90. A handful of validators swing a dozen places between windows. This is real, and it is worth saying plainly: among the cluster of competent 9% validators, the month-to-month ordering is noisy and you should not over-read a few tenths of a point. But the top tier and the bottom tier are stable, and that is where the money is.

⌖ What this means if you stake TAO

The practical takeaways are simpler than the analysis behind them.

Stop choosing on take rate alone. The fee is the most visible number and one of the least informative. A 0% take tells you what the validator charges, not what it earns. Two validators at the same fee differed by more than three points of realized yield this year.

If you want the simplest safe default, the top of this list is durable. The validators that led across every window did so for a reason. You do not need to chase the exact number-one each month; the top handful are clustered within half a point and all of them are stable. Picking any of them protects you from the real risk, which is not landing in the top three, it is sitting in the bottom five without knowing it.

Check where your current validator actually ranks. If you are staked with one of the bottom-five names here, the cost is concrete: $350 to $440 a year per $10,000, every year, for as long as you stay. Moving your delegation costs nothing but a transaction fee.

The cost compounds and recurs. This is not a one-time $441. It is $441 every year you remain, on every $10k, and it compounds against the yield you would otherwise be reinvesting. Over a few years it is a meaningful fraction of your principal, quietly lost to a choice you made once and never revisited.

> THE BOTTOM LINE
>
> The validator you delegate to is one of the highest-leverage, lowest-effort decisions you make as a nominator, and the ecosystem gives you almost no information to make it well. The advertised take rate, the one number everyone uses, is a weak proxy for the only number that matters. This report is the number that matters, for the 25 validators holding the overwhelming majority of staked TAO.

⌖ Methodology

Every figure in this report is reproducible from public data. We pulled daily nominator yield history from the Taostats API endpoint `/api/dtao/validator/yield/history/v1` for 25 validators across all 129 subnets including root, for 365 days ending June 12, 2026. The result is 860,642 daily yield snapshots stored locally and available for verification.

Universe. The 25 validators are the largest commercial root validators by global nominator count, filtered to those active on at least 50 subnets. This deliberately excludes subnet-primary validators such as Chutes and Macrocosmos, which validate a single subnet and do not compete for general nominators. The excluded subnet-primary operators are listed separately below.

Realized APY. For each validator and window, we average the daily one-day nominator APY reported by the chain, net of the validator's take. This is the return a nominator actually experienced, already inclusive of take and performance. Where a validator has fewer than 70% of the window's days available, or fewer than 30 days, it is flagged as insufficient and excluded from that window's ranking.

Windows. The headline ranking uses 90 days as the primary horizon. We also compute 7, 30, 60, 180, and 365-day windows and publish the rank stability across all

six. Validators that joined recently have full coverage at short windows and partial or no coverage at long ones; this is noted rather than hidden.

Decomposition. To split each validator's gap-to-best into fee and performance, we back out gross APY from net using the advertised take: `gross = net / (1 - take)`. The take penalty is the portion of the gap attributable to charging more than the leader; the skill penalty is the portion attributable to earning less gross yield. The two sum to the total gap.

Cost per $10k. The annual dollar cost is the gap in percentage points applied to a $10,000 stake, holding for one year. Because the underlying APY figures are already annualized, this is a direct multiplication and does not assume compounding.

What we did not do. We did not assess validator value-adds beyond yield: some operators offer insurance, custom infrastructure, or integrations that may justify a return gap for specific nominators. We make no judgment on those. We did not pre-notify any validator. We make no predictions about future performance; every figure is historical and window-dependent.

Excluded subnet-primary validators

The following operators appear among the largest root entities but validate only their own subnet and are not part of the commercial comparison: Chutes (SN64 primary), Macrocosmos, and several single-subnet operators. They serve a different function and ranking them against general-purpose validators would be misleading.

⌖ Limitations

Realized yield is backward-looking. A validator's past performance reflects its weight-setting and uptime over the window, which can change. New validators with short histories are genuinely hard to rank and we flag rather than guess. The decomposition assumes the advertised take is the actual take applied to nominator rewards; if a validator applies take differently in practice, the split between fee and performance would shift, though the total realized gap would not. Cap-weighted cross-subnet returns

were not included in this edition because the subnet market-cap series needed cleaning; the root-only comparison, which every validator participates in identically, is the cleanest and is what we lead with.

This is a snapshot of a single year. We intend to refresh it quarterly, which will also reveal which validators improve in response to being measured.
