RESEARCH / REPORT 003

<< ON-CHAIN AUDIT · JUNE 2026 >>

The Subnet Hierarchy

A cross-subnet audit of validator quality on Bittensor. The common advice is to pick a specialized validator for each subnet you stake on. The data says the opposite: four validators win 91% of all 128 subnets. There is no best validator for a subnet. There are just the best validators, and they win almost everywhere.

FLOWSNIPER RESEARCH · JUNE 12, 2026 · 25 VALIDATORS · 128 SUBNETS · 365 DAYS

Our last report, The Validator Hierarchy, ranked the 25 largest Bittensor validators by what nominators actually earned on root. It found a 4.41-point spread between best and worst, and that the advertised take rate barely predicts realized return.

That report covered root staking only. This one asks the obvious follow-up: when you stake alpha on a specific subnet rather than root, does the same ranking hold? Or does each subnet have its own best validator, the way the common advice implies when it tells you to find specialists?

We took the same 25 validators and the same year of daily on-chain yield data, and measured how each validator performed on every one of the 128 subnets it validates. The answer is unusually clean.

> *Validator quality is not subnet-specific. It transfers. The validators that lead on root also lead almost everywhere else, and a handful of them win nearly the entire network.*

91%	52%	96%	1.12x
of 128 subnets won by just 4 validators	of all subnets won outright by a single validator	of subnets where the top validator places top-5	median capture of the best validator vs the median

The practical consequence for a staker is the opposite of the usual advice. You do not need to research a different validator for every subnet. The same small set of names that topped the root ranking also tops the subnet rankings, almost everywhere. One good validator choice covers your whole portfolio.

An important note on what these yields mean

Before the findings, a caveat that is central to reading this report honestly. The subnet yields here are **alpha-denominated**. A "58% APY" on a subnet means you earned 58% more of that subnet's alpha token, not 58% more value in TAO terms.

Our first report, The Yield Illusion, showed why that distinction matters: the median subnet alpha staker lost money in TAO last month despite headline APYs above 50%, because the alpha token's price fell faster than the yield accrued. Those high subnet APYs are mostly dilution, not profit.

So this report does not claim that staking alpha on a subnet at "58% APY" makes you money. It almost certainly does not, in TAO terms. What this report measures is something narrower and more reliable: **among validators on the same subnet, who captured the most of whatever yield existed.** Because every validator on a subnet faces the identical token price, the price decay cancels out of the comparison. The ranking of validators within a subnet is real even though the absolute yield number is an illusion.

> **HOW TO READ THE NUMBERS**
>
> Treat every subnet APY in this report as a relative measure, not a return you would earn. The honest metric is the `capture ratio`: how much more the best validator earned than the median validator on the same subnet. That number is price-neutral and is what the analysis leans on. If you are deciding whether to stake alpha on a subnet at all, read the Yield Illusion first. If you have already decided to stake on a subnet and only want to know which validator to pick, this report is your answer.

⌖ Quality transfers: the root leaders win everywhere

For each validator we took its rank on every subnet it validates, then summarized how often it lands near the top. The pattern is a clean two-tier structure. A small group of validators places top-5 on the overwhelming majority of subnets. After that group, performance falls off a cliff.

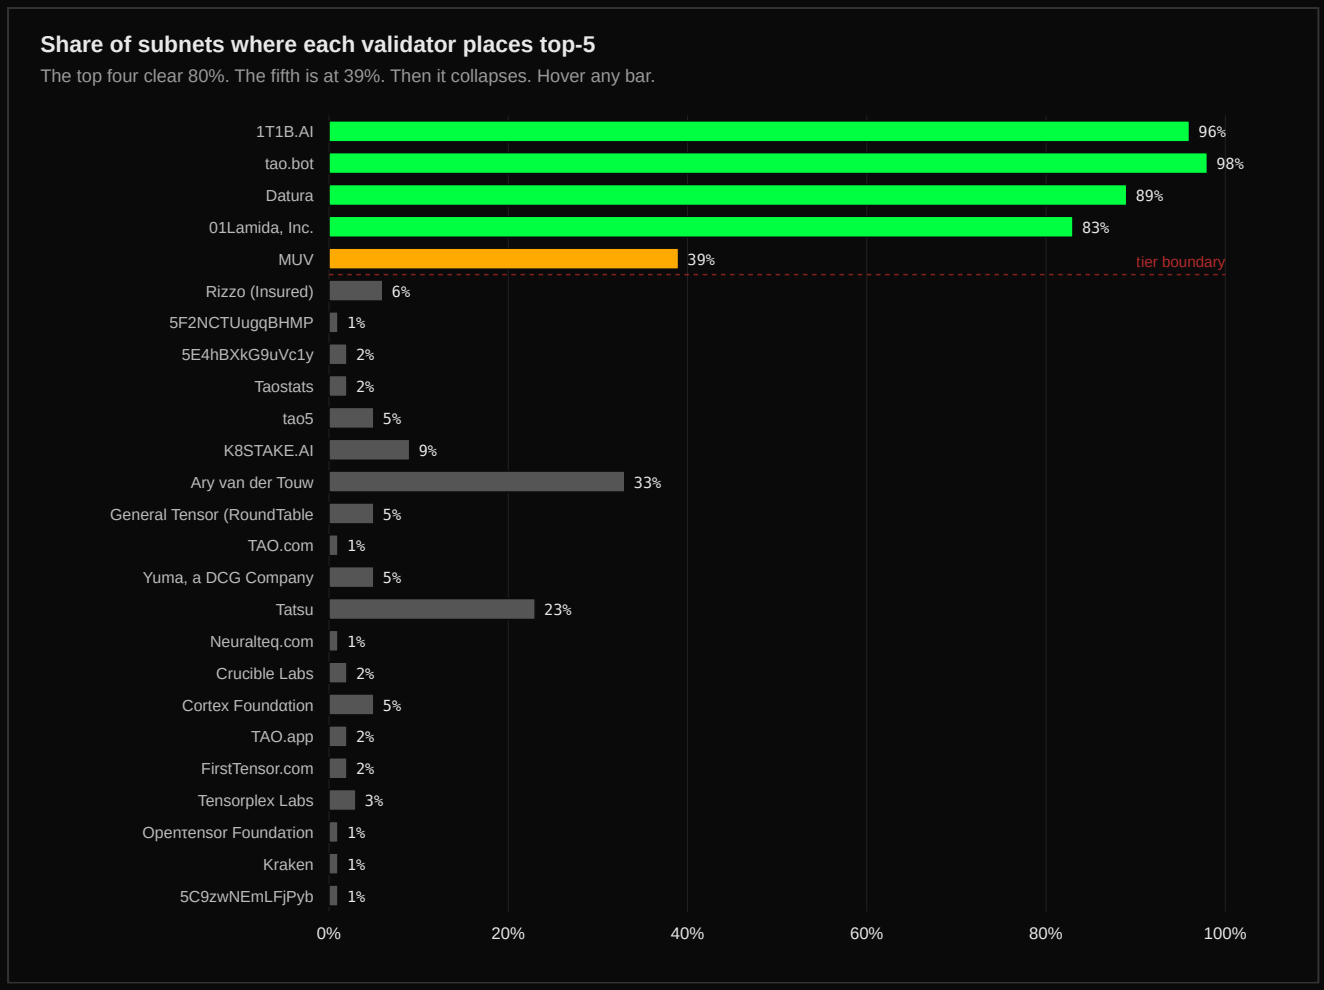

Figure 1. For each validator: the share of subnets where it places top-5. The top four validators land top-5 on 83 to 98 percent of subnets. The fifth is at 39 percent, and everyone below falls under 10 percent. Hover any bar for detail.

ROOT #	VALIDATOR	SUBNETS	MEDIAN RANK	% TOP-5	% TOP-3	% TOP-1
1	1T1B.AI	127	1	96%	90%	52%
2	tao.bot	126	2	98%	91%	24%
3	Datura	127	3	89%	52%	3%
4	01Lamida, Inc.	122	4	83%	28%	2%
5	MUV	126	6	39%	5%	2%
6	Rizzo (Insured)	127	12	6%	2%	1%
7	5F2NCTUugqBHMP	127	15	1%	0%	0%
8	5E4hBXkG9uVc1y	126	14	2%	0%	0%
9	Taostats	127	12	2%	1%	0%
10	tao5	127	11	5%	3%	0%
11	K8STAKE.AI	127	17	9%	9%	4%
12	Ary van der Touw	108	7	33%	18%	14%
13	General Tensor (RoundTable21)	127	11	5%	1%	0%
14	TAO.com	127	11	1%	0%	0%
15	Yuma, a DCG Company	127	11	5%	1%	1%
16	Tatsu	127	6	23%	2%	0%
17	Neuralteq.com	127	13	1%	0%	0%
18	Crucible Labs	127	17	2%	2%	0%
19	Cortex Foundαtion	127	18	5%	1%	0%
20	TAO.app	126	23	2%	1%	0%
21	FirstTensor.com	119	17	2%	0%	0%

ROOT #	VALIDATOR	SUBNETS	MEDIAN RANK	% TOP-5	% TOP-3	% TOP-1
22	Tensorplex Labs	117	15	3%	2%	1%
23	Opentensor Foundation	121	23	1%	0%	0%
24	Kraken	123	23	1%	1%	1%
25	5C9zwNEmLFjPyb	113	18	1%	0%	0%

Shaded rows are the five validators that place top-5 on more than a third of subnets. Root rank is from The Validator Hierarchy. Note how closely it tracks the subnet columns: the root-1 validator has a median subnet rank of 1; the root-2 validator a median of 2.

The correspondence with the root ranking is striking. 1T1B.AI, first on root, has a median subnet rank of 1 and places top-5 on 96 percent of subnets. tao.bot, second on root, has a median subnet rank of 2 and places top-5 on 98 percent. Datura, third on root, third on subnets. The validators that are good at validating are good at it everywhere. Quality is a property of the operator, not the subnet.

And the fall-off is sharp. After the top four (all above 80 percent top-5), the fifth validator drops to 39 percent and the sixth to 6 percent. There is no gentle gradient of slightly-worse validators. There is a small tier that is excellent across the board, and a large tier that is rarely near the top of anything.

⌖ Win concentration: two names, three-quarters of the network

The clearest way to see the concentration is to count outright wins: on how many subnets was each validator the single best performer?

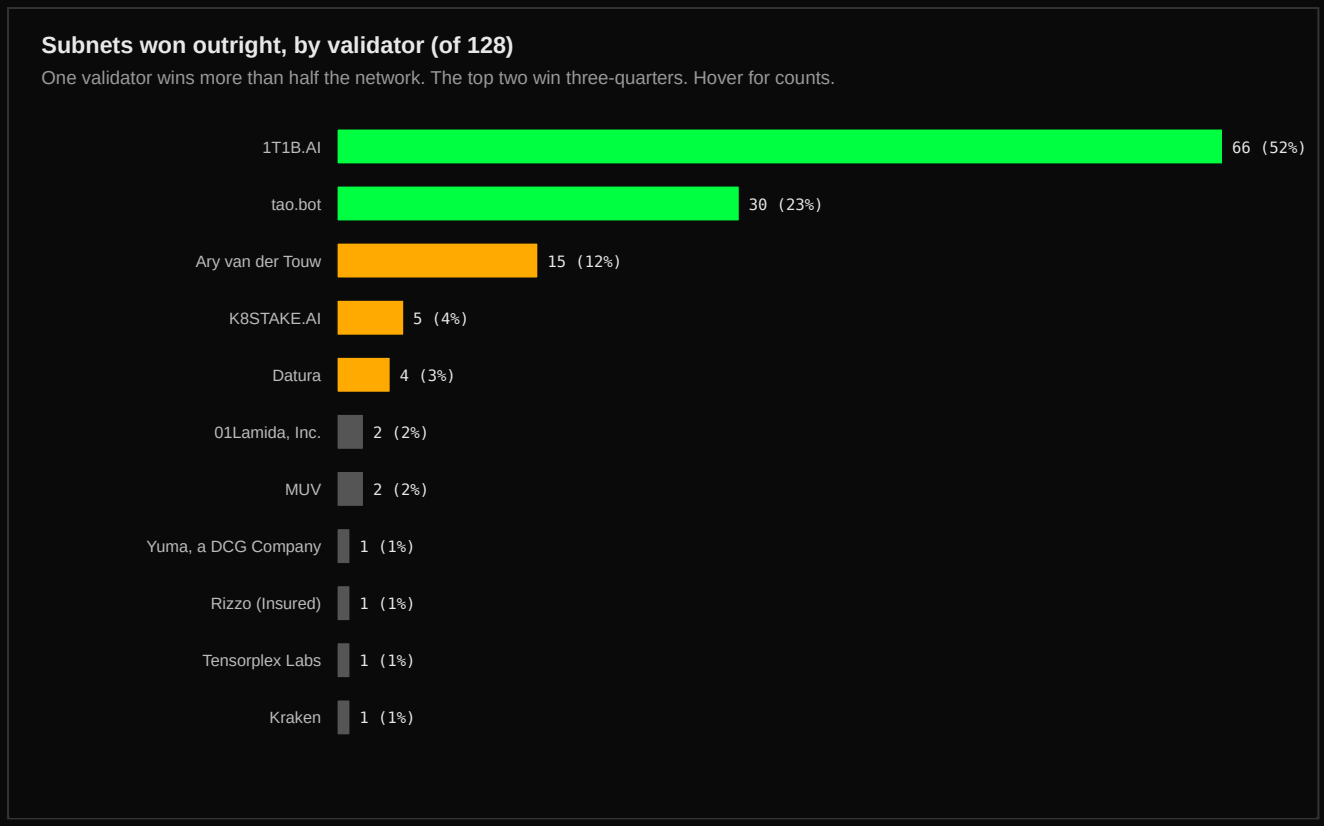

Figure 2. Number of subnets on which each validator was the single best performer, out of 128. One validator wins more than half the network. The top two win three-quarters. Hover for counts.

VALIDATOR	SUBNETS WON	SHARE
1T1B.AI	66	52%
tao.bot	30	23%
Ary van der Touw	15	12%
K8STAKE.AI	5	4%
Datura	4	3%
01Lamida, Inc.	2	2%
MUV	2	2%
Yuma, a DCG Company	1	1%
Rizzo (Insured)	1	1%
Tensorplex Labs	1	1%
Kraken	1	1%

1T1B.AI is the single best validator on 66 of 128 subnets, a clean 52 percent of the entire network. Add tao.bot and the top two cover 96 subnets, 75 percent. Add Ary van der Touw and K8STAKE.AI and the top four cover 91 percent. The remaining nine percent is split in ones and twos among everyone else.

This is the finding that overturns the common advice. If validator quality were subnet-specific, wins would be spread across many operators, each strong on the subnets they specialize in. Instead, a tiny number of validators win almost everywhere. The specialist-per-subnet strategy is chasing a pattern that is not in the data.

How much does the choice matter? The capture ratio

Knowing which validator is best on a subnet only matters if the gap to the others is meaningful. To measure it without the alpha-price distortion, we use the capture ratio:

the best validator's yield divided by the median validator's yield on the same subnet. Because both face the identical token price, this ratio is a clean, price-neutral measure of how much validator choice is worth.

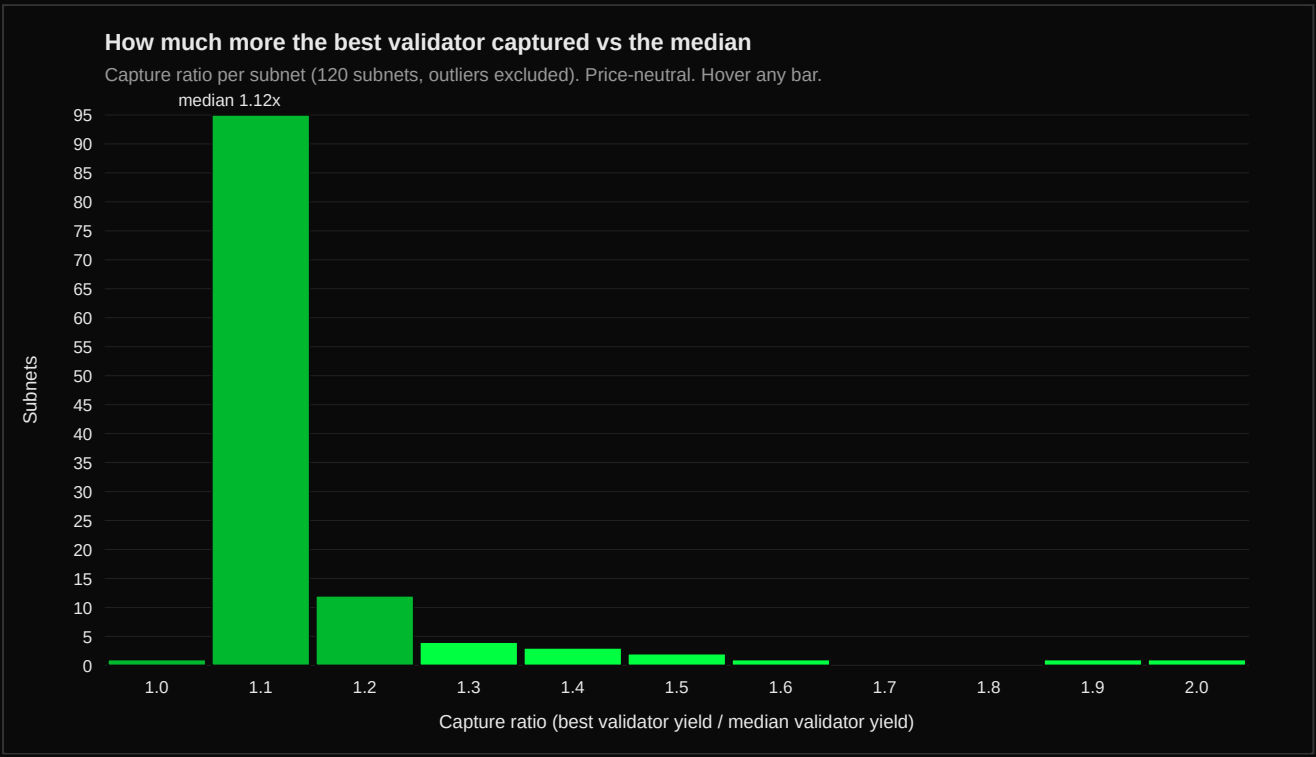

Figure 3. The capture ratio across the 120 non-outlier subnets. On the typical subnet the best validator captured 1.12x what the median did. On the most extreme, more than 2x. Hover any bar for the count.

On the median subnet, the best validator captured 1.12 times what the median validator did, a 12 percent edge in yield for the same staked alpha, the same subnet, the same token price. The difference is purely the operator. On some subnets the edge is far larger: the best validator on subnet 65 captured 2.01 times the median.

Twelve percent may sound modest next to the headline subnet APYs, but remember those APYs are largely illusory in TAO terms. The capture ratio is the part that is real. A consistent 12 percent edge in actual yield capture, compounded across a year and across every subnet position you hold, is the difference between a good validator choice and a mediocre one, and it costs nothing to choose well.

⌖ The per-subnet lookup

The complete table: for each of the 128 subnets, the validator that delivered the best realized nominator yield over the trailing 90 days, the capture ratio versus the median validator, and the spread between the best and worst validator on that subnet. Use the search box to jump to a subnet by number.

Capture ratio and spread are the reliable, price-neutral measures. Eight subnets are flagged `outlier` where a short-lived emission spike produced an implausible annualized figure (marked with *); their ranking is shown but the magnitude should be ignored. Spread is in percentage points of alpha-denominated APY and is comparable only within a subnet.

SUBNET	BEST VALIDATOR (90D)	CAPTURE VS MEDIAN	BEST-WORST SPREAD (PP)	VALIDATORS	FLAG
1	tao.bot	1.11x	15.2	25	
2	tao.bot	1.21x	15.8	25	
3	Ary van der Touw	1.22x	29.6	25	
4	1T1B.AI	1.10x	16.5	25	
5	Ary van der Touw	1.44x	70.9	25	
6	Ary van der Touw	1.67x	31.0	25	
7	01Lamida, Inc.	1.17x	10.9	25	
8	1T1B.AI	1.19x	19.4	25	
9	1T1B.AI	1.12x	12.0	25	
10	1T1B.AI	1.10x	20.1	25	
11	MUV	1.20x	27.6	25	
12	Yuma, a DCG Company	1.19x	12.2	25	
13	tao.bot	1.09x	27.2	25	
14	K8STAKE.AI	1.36x	33.3	25	
15	tao.bot	1.10x	38.7	23	
16	1T1B.AI	1.12x	9.8	25	
17	tao.bot	1.11x	15.9	24	
18	1T1B.AI	1.12x	24.7	25	
19	1T1B.AI	1.11x	25.1	25	
20	tao.bot	1.10x	10.2	25	
21	Ary van der Touw	1.32x	16.7	24	

SUBNET	BEST VALIDATOR (90D)	CAPTURE VS MEDIAN	BEST-WORST SPREAD (PP)	VALIDATORS	FLAG
22	1T1B.AI	1.12x	36.1	25	
23	MUV	1.15x	14.8	25	
24	1T1B.AI	1.17x	22.4	24	
25	Ary van der Touw	1.11x	12.0	25	
26	tao.bot	1.27x	46.5	24	
27	1T1B.AI	1.12x	8.5	25	
28	tao.bot	1.17x	8.4	25	
29	1T1B.AI	1.11x	8.9	23	
30	1T1B.AI	1.12x	10.9	25	
31	1T1B.AI	1.13x	30.7	25	
32	1T1B.AI	1.15x	8.0	25	
33	1T1B.AI	1.12x	15.8	25	
34	Datura	1.13x	12.7	25	
35	1T1B.AI	1.17x	8.9	24	
36	tao.bot	9.1x*	1397.0	21	outlier
37	1T1B.AI	1.13x	11.3	25	
38	1T1B.AI	1.12x	50.2	23	
39	Ary van der Touw	1.23x	33.6	25	
40	1T1B.AI	1.12x	8.8	25	
41	tao.bot	1.10x	12.6	25	
42	1T1B.AI	1.12x	11.9	24	
43	K8STAKE.AI	1.29x	32.2	25	

SUBNET	BEST VALIDATOR (90D)	CAPTURE VS MEDIAN	BEST-WORST SPREAD (PP)	VALIDATORS	FLAG
44	1T1B.AI	1.12x	20.2	25	
45	1T1B.AI	1.12x	11.6	25	
46	Datura	1.12x	9.7	25	
47	tao.bot	1.12x	43.6	23	
48	1T1B.AI	1.11x	16.9	25	
49	tao.bot	1.12x	31.2	25	
50	1T1B.AI	1.11x	17.8	25	
51	Ary van der Touw	1.14x	13.4	25	
52	1T1B.AI	1.13x	20.2	25	
53	tao.bot	1.12x	23.1	25	
54	1T1B.AI	1.10x	8.4	24	
55	1T1B.AI	1.27x	17.3	25	
56	1T1B.AI	1.10x	11.7	25	
57	Rizzo (Insured)	51.2x*	4869.1	25	outlier
58	Ary van der Touw	1.12x	11.0	25	
59	tao.bot	1.12x	9.9	25	
60	1T1B.AI	1.12x	12.0	25	
61	1T1B.AI	1.10x	8.7	25	
62	Ary van der Touw	1.18x	9.1	24	
63	1T1B.AI	1.19x	13.7	23	
64	Ary van der Touw	1.22x	17.1	25	
65	K8STAKE.AI	2.01x	46.2	25	

SUBNET	BEST VALIDATOR (90D)	CAPTURE VS MEDIAN	BEST-WORST SPREAD (PP)	VALIDATORS	FLAG
66	K8STAKE.AI	1.27x	13.2	25	
67	tao.bot	1.17x	70.3	22	
68	tao.bot	1.10x	10.1	25	
69	Ary van der Touw	1.58x	46.8	25	
70	tao.bot	1.8x*	179.9	22	outlier
71	tao.bot	1.10x	9.8	25	
72	1T1B.AI	1.11x	9.0	25	
73	K8STAKE.AI	1.93x	50.0	25	
74	1T1B.AI	1.14x	15.8	25	
75	1T1B.AI	1.10x	8.9	25	
76	1T1B.AI	1.12x	49.6	22	
77	1T1B.AI	1.13x	10.5	25	
78	tao.bot	6.0x*	673.6	25	outlier
79	1T1B.AI	1.10x	7.3	24	
80	tao.bot	1.26x	79.7	25	
81	Ary van der Touw	1.23x	44.3	25	
82	1T1B.AI	9.9x*	1243.1	25	outlier
83	1T1B.AI	1.15x	20.6	25	
84	tao.bot	1.50x	70.7	24	
85	1T1B.AI	1.10x	21.9	25	
86	1T1B.AI	1.11x	30.5	25	
87	1T1B.AI	1.11x	45.5	24	

SUBNET	BEST VALIDATOR (90D)	CAPTURE VS MEDIAN	BEST-WORST SPREAD (PP)	VALIDATORS	FLAG
88	1T1B.AI	1.11x	9.8	25	
89	1T1B.AI	1.12x	11.3	25	
90	1T1B.AI	1.11x	61.1	25	
91	1T1B.AI	1.1x*	94.4	21	outlier
92	1T1B.AI	1.16x	45.3	24	
93	1T1B.AI	1.15x	11.2	25	
94	tao.bot	1.11x	51.8	25	
95	1T1B.AI	1.13x	22.2	25	
96	tao.bot	1.1x*	66.8	20	outlier
97	Datura	1.32x	73.1	22	
98	1T1B.AI	1.12x	12.4	25	
99	tao.bot	1.13x	91.7	23	
100	Ary van der Touw	1.18x	56.8	23	
101	1T1B.AI	1.15x	41.2	24	
102	tao.bot	1.41x	121.8	21	
103	Ary van der Touw	1.12x	14.7	25	
104	tao.bot	1.2x*	55.0	3	outlier
105	tao.bot	1.12x	38.0	25	
106	tao.bot	1.10x	26.4	25	
107	1T1B.AI	1.11x	97.9	22	
108	tao.bot	1.12x	58.5	25	
109	1T1B.AI	1.13x	34.4	24	

SUBNET	BEST VALIDATOR (90D)	CAPTURE VS MEDIAN	BEST-WORST SPREAD (PP)	VALIDATORS	FLAG
110	1T1B.AI	1.16x	20.5	25	
111	tao.bot	1.10x	11.6	24	
112	1T1B.AI	1.11x	17.3	25	
113	1T1B.AI	1.12x	92.0	25	
114	1T1B.AI	1.14x	40.4	22	
115	1T1B.AI	1.13x	24.3	25	
116	Tensorplex Labs	1.34x	72.2	24	
117	Kraken	1.49x	78.6	24	
118	1T1B.AI	1.19x	27.7	25	
119	1T1B.AI	1.11x	40.8	24	
120	01Lamida, Inc.	1.11x	12.7	25	
121	1T1B.AI	1.16x	16.2	25	
122	Datura	1.17x	19.0	24	
123	1T1B.AI	1.10x	24.6	25	
124	1T1B.AI	1.12x	13.6	25	
125	1T1B.AI	1.11x	22.1	25	
126	1T1B.AI	1.13x	64.9	21	
127	Ary van der Touw	1.25x	26.5	25	
128	1T1B.AI	1.11x	9.3	24	

⌖ What this means if you stake alpha on subnets

You do not need a different validator for each subnet. The single most useful takeaway: the validators that are best on root are best almost everywhere. Picking one of the top handful covers your entire subnet portfolio. The effort of researching a specialist for each position buys you nothing, because the specialists mostly do not exist.

The same names keep winning. 1T1B.AI and tao.bot together are the best validator on three-quarters of all subnets. If you stake alpha and have not thought about your validator, defaulting to one of the consistent top performers is the safe, simple choice, the same conclusion the root report reached, now confirmed across the whole network.

Choosing well is worth a real, if modest, edge. The best validator captures about 12 percent more yield than the median on the typical subnet, more on some. That edge is price-neutral and recurring. It will not rescue a losing alpha position, but among validators it is free money left on the table if you pick poorly.

But the validator is the smallest of your subnet decisions. The honest hierarchy of what determines your outcome when you stake alpha: first, whether the subnet's token holds its value (mostly it does not, see the Yield Illusion); second, your timing; a distant third, which validator you pick. This report optimizes the third. Do not let it distract from the first.

> THE BOTTOM LINE
>
> There is no best validator for a given subnet. There are just the best validators, and they win almost everywhere. The same four operators are the best choice on 91 percent of the network. Pick one of them and you have made the validator decision correctly for your entire portfolio, on root and on every subnet. Then spend your real attention on the decision that actually drives your return: whether to be in that subnet's alpha at all.

⌖ Methodology

This report uses the same dataset as The Validator Hierarchy: 860,642 daily nominator-yield snapshots pulled from the Taostats API endpoint `/api/dtao/validator/yield/history/v1`, covering 25 validators across all 129 subnets including root, for 365 days ending June 12, 2026.

Universe. The 25 validators are the largest commercial root validators by global nominator count, active on at least 50 subnets. Subnet-primary operators that validate only their own subnet are excluded, as in the prior report.

Per-subnet ranking. For each subnet we average each validator's daily one-day nominator APY over the trailing 90 days, then rank validators within that subnet. The validator with the highest average realized yield is the subnet's best. Validators with fewer than 70 percent of the window's days, or fewer than 25 days, are excluded from that subnet's ranking.

Capture ratio. The price-neutral quality measure: the best validator's average APY divided by the median validator's average APY on the same subnet. Because all validators on a subnet earn the same alpha token and face the same token price, dividing one by another cancels the price effect, leaving a clean comparison of validator performance.

Outlier handling. Eight subnets where the best validator's annualized APY exceeded 200 percent are flagged as outliers. These reflect short-lived emission spikes or low-liquidity dislocations that produce implausible annualized figures, the most extreme being a single subnet reading above 1,400 percent. We show their validator ranking but the magnitudes should be disregarded; they are excluded from the summary statistics.

Transfer analysis. For each validator we compute its rank on every subnet it validates, then report the median subnet rank and the share of subnets where it places top-5, top-3, and first. Comparing these to the root ranking shows whether quality transfers across subnets. It does.

What this is not. The subnet APY figures are alpha-denominated and are not TAO returns. This report ranks validators within subnets; it does not endorse subnet staking as profitable. See The Yield Illusion for the TAO-denominated picture.

> DISCLOSURE
>
> FlowSniper delegates by default to tao.bot (`5E2LP6E...d6ddd748`), which appears prominently in this report as the best validator on 30 subnets. We have no commercial relationship with tao.bot and no contact with their team. We selected them as our default before any of this analysis, on publicly observable characteristics, and the data has repeatedly confirmed that choice rather than been shaped by it. As with our prior reports, had the data ranked tao.bot poorly we would have published it and switched.

⌖ Limitations

Realized yield is backward-looking and reflects the trailing 90 days; validator performance can change. The subnet APY figures are alpha-denominated and should not be read as TAO returns. The capture ratio and within-subnet rankings are reliable; the absolute yield magnitudes are not, and on the eight flagged outlier subnets even

the magnitude of the spread should be disregarded. Ary van der Touw appears strong where present but validates fewer subnets (108 of 128) with a shorter history, so its standing carries more uncertainty than the fully-covered validators. This is a single 90-day snapshot within the broader year of data; we intend to refresh it alongside the root ranking quarterly.
