

Deep Dive:

# Solana Onchain Activity



# Part I

# Network

# Validators and Block Production

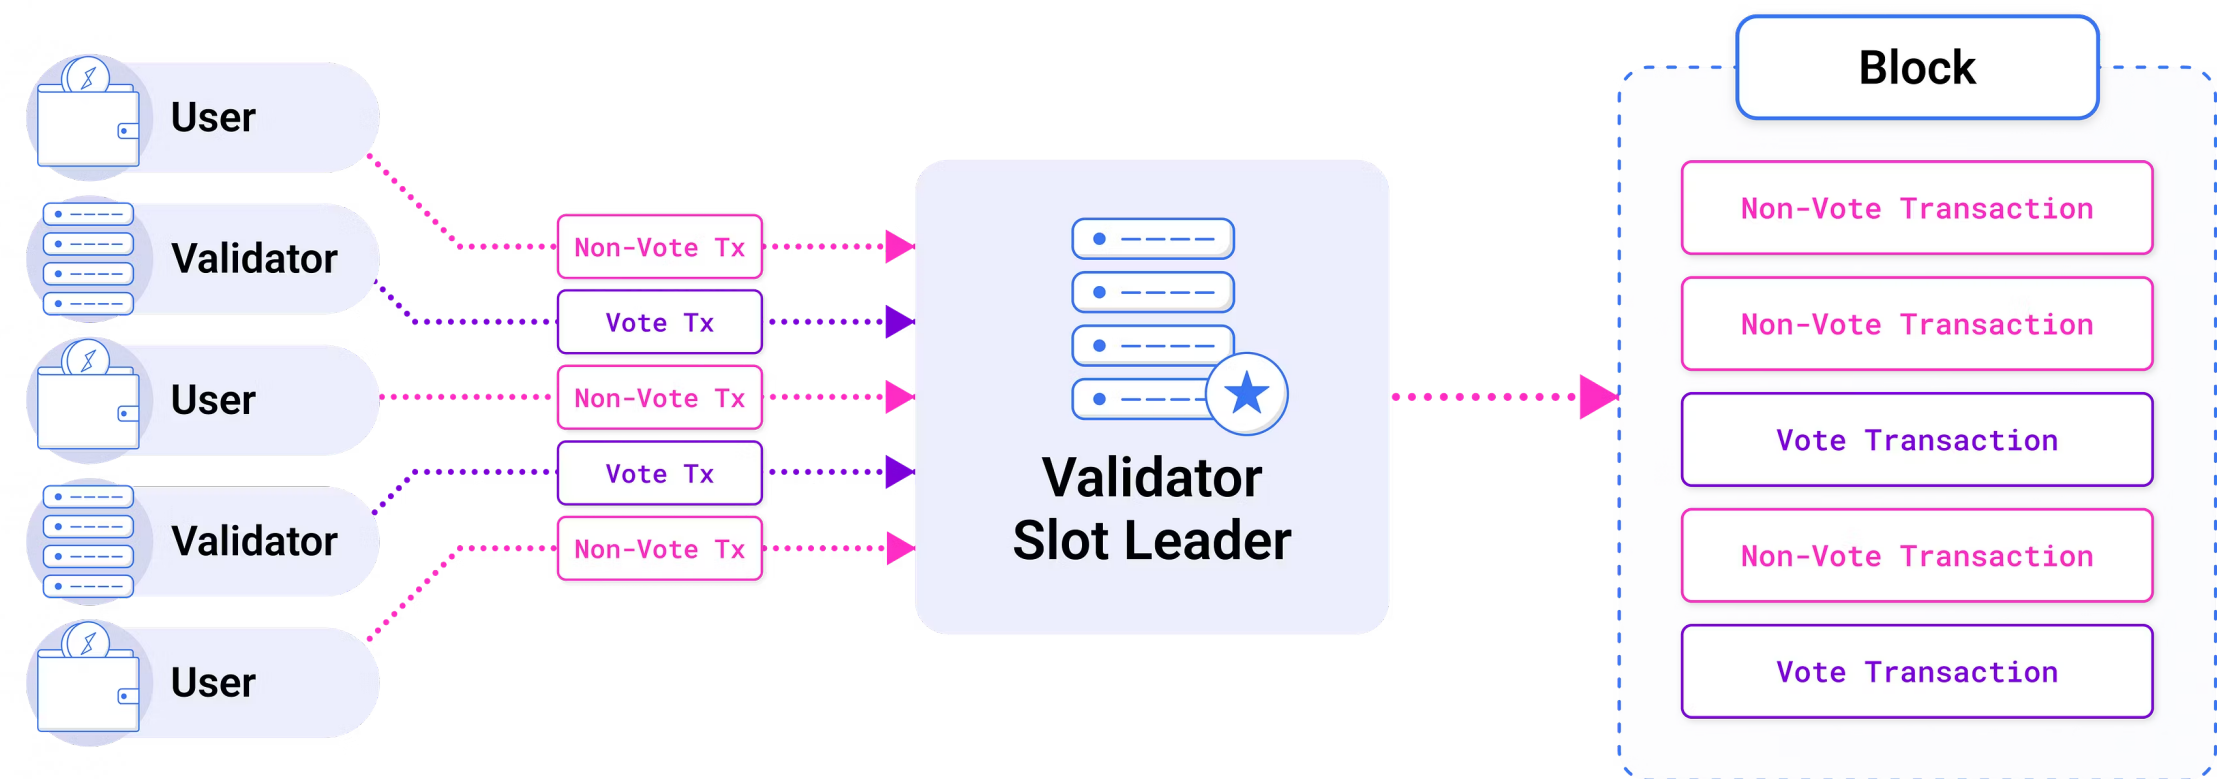
A **validator** is a computer or node that verifies the accuracy of transactions, creates new blocks, and participates in consensus.

A **slot** represents an allotment of time for which a block may be produced. Validators are assigned to act as **slot leaders** for a set of 4 consecutive slots.

During these 4 slots, the designated slot leader is responsible for creating blocks by bundling hundreds to thousands of transactions, which include both vote and non-vote transactions.

Validators submit **vote transactions** to reach consensus. One can understand it as validators casting votes to agree on what is happening within the Solana blockchain and what the correct order of events is.

**Non-vote transactions** are any other kind of transaction, like a swap, a token transfer, etc.

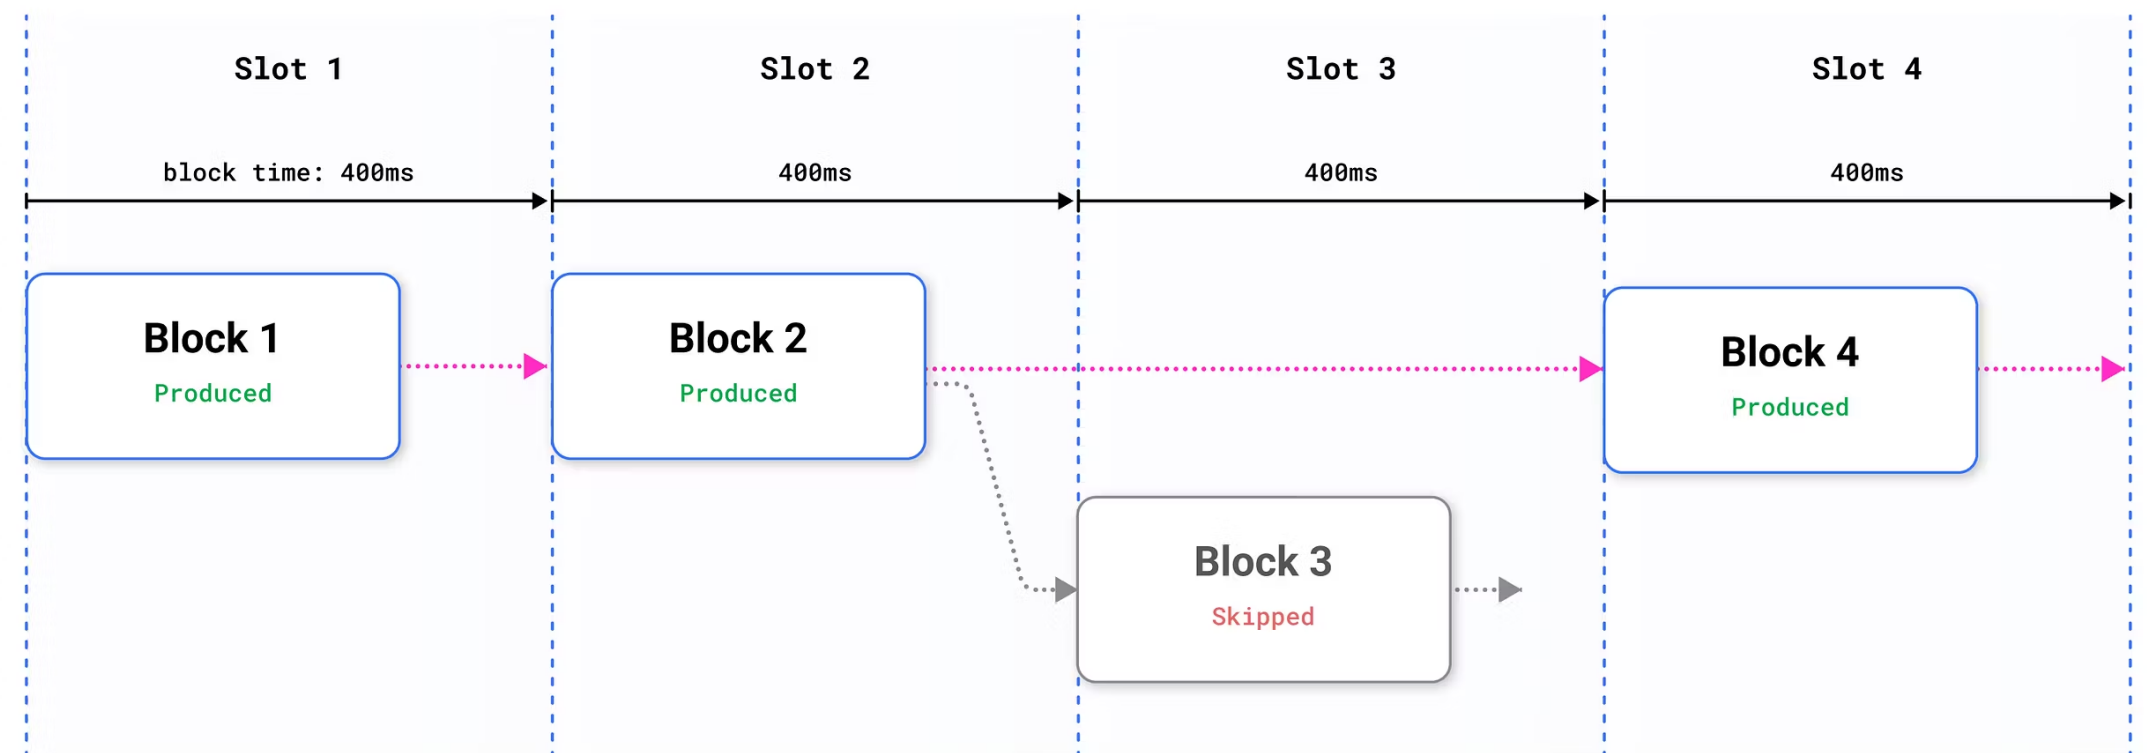




# Measuring Network Performance

We will use four main metrics to measure Solana's network performance:

1. **Block time:** Measures how quickly blocks are created, with Solana aiming to produce one block every 400ms.
2. **Blocks produced:** Measures how many successful blocks Solana outputs, targeting 216,000 blocks per day.
3. **Skip rate:** When a slot leader fails to produce a block in their assigned slot, the block gets skipped. Skip rate tracks the ratio of skipped versus total slots.
4. **Vote latency:** Measures the time between when a block is produced and when a validator's vote for that block is included in a future slot. It is measured in slots, and the minimum possible latency is one slot.






# Solana achieves 80% reduction in skip rate since January.

Starting at 0.6% in early January, the 7-day rolling average has fallen to 0.125% by August.

The steepest improvement occurred in February, with skip rate stabilizing at lower levels since late May.

## Skip Rate By 7 Day Rolling Average






# Solana now consistently operates faster than its 400ms target.


From January to August, Solana's average block time improved from 407.2ms to a target-beating 396.1ms.

This 11ms reduction in block time resulted in over 5,900 additional blocks being produced each day, pushing the daily average from 212,200 to 218,100.

Block Time  
By 7 Day Rolling Average



Block Production  
By 7 Day Rolling Average



99% of votes  
now land  
within a  
single slot.

Share of votes landing within 1 slot  
7 Day Rolling Average







# p99 vote latency cut in half since January.


The average vote latency has fallen from 1.35 slots in January to 1.03 in August.

p99 latency has dropped and become far less volatile: from an average of 3.8 slots with peaks of 6 slots in January to 1.75 slots with peaks of 2 slots in August.

Average Vote Latency Evolution  
7 Day Rolling Average



P99 Vote Latency  
By Day



# Solana achieves 100,000 non-vote TPS on mainnet

On August 17, Cavey from Temporal conducted a live throughput test, reaching 102,228 non-vote TPS in a single slot with blocks containing 40,000 non-vote transactions.

The test averaged over 90,000 non-vote TPS across four consecutive slots using mainly "no-op" transactions, effectively empty operations that still require significant network overhead for signature validation and processing.

This burst test demonstrates what is possible on Solana today, even before further Agave improvements or Firedancer being rolled out. The test also highlights the importance of efficient onchain programs. Cavey is a leading advocate for "p-token," a more compute-efficient replacement for the current SPL token implementation. This proposal, formalized under SIMD-0266 and heading to governance vote in the coming weeks to months, could enable Solana to sustainably process 10,000-20,000 swaps and 80,000 token transfers per second, according to Cavey.



Cavey Cool

CaveyttUBTKttncu1e4RF814XjuoGfYv8cEsiKGDNCPX

Frankendancer • 716,774 SOL • 208.91.110.44

v0.703.20300 (0.1769%)

Aug 17, 2025, 7:39:21 AM (1m 49s ago)


Slot		Votes	Non-votes	Fees	Tips	Duration	Compute Units
360650523		948	42,118	0.112	0.003	412 ms	59,947,125 (100%)
360650522		222	39,599	0.104	0.016	408 ms	59,868,323 (100%)
360650521		861	39,866	0.125	0.006	409 ms	59,997,547 (100%)
360650520		1,064	28,017	0.080	0.004	420 ms	59,945,006 (100%)

# Successful transactions rebound in both number and rate.


After a dip in early 2025, the percentage of successful non-vote transactions climbed to a yearly high of 75% in July.

Meanwhile, the number of successful non-vote transactions per block increased from a low of 230 in March to a high of 390, a 70% increase.

Successful Non-Vote Transactions per Block  
7 Day Rolling Average



Successful VS Failed Rate of Non-Vote Transactions  
7 Day Rolling Average





# Non-vote TPS breaks yearly records in July before dipping in August.


Non-vote TPS reached a yearly high in July, with a daily average of 1,320. Peak activity was even stronger, with the top 1% of blocks (p99) at 2,460 TPS.

July activity was higher than in the previous January peak, which was driven by the TRUMP memecoin launch, and represented a total recovery from the dip in late Q1.

Activity declined in August, however, to the lowest values of the year, with the average falling to 800 and p99 to 1100 non-vote TPS by month's end.

## Non-Vote TPS 7 Day Rolling Average

● P99 ● Average



# What are compute units (CU)?

When you execute an on-chain transaction, it uses up computational resources or compute. Compute is measured in compute units (CU).



## Block Compute Budget

- Each block: limit of 48M CU (raised to 50M CU on April 10 and again to 60M CU on July 22)
- Blocks consist of multiple transactions
- Each transaction deducts CU from the block budget



## CU Limits

- Maximum: 1.4M CU per transaction
- Transactions are split into instructions
- Instruction Limits:
  - Default Limit: 200k CU
  - Custom Limit: users can request a specific limit on compute units


# Block CU demand falls to yearly lows in August.

August marked the lowest network utilization of 2025, with median blocks at 35.6M CU and averages of 37.9M CU. Both metrics hit yearly lows on August 31 at 25.5M and 27.5M respectively.

This is in contrast to typical values between 40-46M CUs on average and p50 often pushing against the block limit. The decrease in activity is further highlighted by the block limit raise to 60M CU on July 22nd (epoch 823) from 50M CU.


## Average Compute Units per Block By Day

● Average Comput Units — Block Limit



## p50 Compute Units per Block By Day

● P99 Comput Units — Block Limit






# Solana's busiest blocks were consistently full until August.

Throughout most of 2025, the top 10% of blocks (p90) ran against the compute limit. This pattern continued into mid-August despite the July 22 increase to 60M CU.

Usage declined in the second half of August, with p90 averaging 54.6M CU, dropping well below previous limits and ending the month at a yearly low of 40.6M CUs.

## p90 Compute Units per Block By Day

● P90 Comput Units    — Block Limit



## Part II

# Validator Clients

# Clients live on mainnet

*Note: In the following section, we analyze the performance of these clients and their flavours across metrics like vote latency and skip rate since January 2025. All metrics are weighted by the number of blocks produced by each client during the analysis period.*

*These metrics are influenced by factors beyond client choice, including geographical location, or how the validators prioritize transactions and build blocks. Though ideal comparisons would control for validator specifications, stake and locations, this historical analysis still offers useful insights.*

*Most of the data used in this analysis comes from the dashboard [reports.firedancer.io](https://reports.firedancer.io), maintained by the Firedancer team.*

## Base Clients

- **Agave:** main Solana validator Client maintained by Anza, continuing the work of the Solana Labs client. Serves as the reference implementation for the network.
- **Frankendancer:** hybrid client as a transitional step towards the full Firedancer client being developed by Jump Crypto. It combines the original consensus and runtime components from Agave with Firedancer's new, optimized networking components, signature verification, and block packing.

## Flavours of Base Client

- **Vanilla:** default Agave/Frankendancer clients.
- **Jito:** modified version of the core client with built-in Maximum Extractable Value (MEV) auction and block engine to maximize revenue and potentially increase reliability and performance.
- **Paladin:** modified Jito client that aims to protect validators from sandwiching and distribute MEV rewards across Paladin validators.



# Part II.A


# Stake

# Three-quarters of Solana stake runs Agave Jito.

Agave Jito ends August with over 300 million SOL in stake or 76% the total, almost 6 times as much as Frankendancer Jito.

## Stake Distribution by Client August 2025

Agave Jito Agave Paladin Agave Vanilla Frankendancer Jito




# The validator ecosystem diversifies as Agave Jito loses dominance.

Starting the year with 93% of network stake, the stake of Agave Jito validators dropped to 76% by August as competitors gained ground.

Paladin captured 9% of stake, Frankendancer grew from under 1% to 13%, while Agave vanilla declined from 6% to 2%.


## Stake Distribution by Client By Day

Agave Jito Agave Paladin Agave Vanilla Frankendancer Jito Frankendancer Vanilla



Frankendancer  
crosses 13%  
of total  
staked SOL.

Frankendancer's Share of Stake  
By Day





# Jito accounts for virtually all of Frankendancer stake.


While vanilla validators led early Frankendancer adoption, March saw Jito become the dominant choice.

As of late August 2025, Frankendancer Jito holds 13.1% of the total Solana stake.

Frankendancer vanilla validators hold 36,000 SOL, virtually 0% of total stake.

## Frankendancer Stake by Client Flavour By Day

Frankendancer Jito Frankendancer Vanilla



## Part II.B

# Skip Rate


# Agave v2.3 achieves skip rates 35 times lower than 1.18.

Since v1.18 through v2.2, each new version of the Agave client has lowered validator skip rates.


The skip rate fell the sharpest when upgrading to v2.0, dropping from 4.6% down to 0.6%. Subsequent releases continued the trend, with v2.2 achieving the lowest skip rate at 0.12%.

However, v2.3 shows a 12% increase in skip rate over v2.2.

**Skip Rate by Agave Version**  
v1.18-v2.3



**Skip Rate by Agave Version**  
v2.2 vs v2.3




# Agave and Jito achieve skip rates less than half of Frankendancer and vanilla clients.

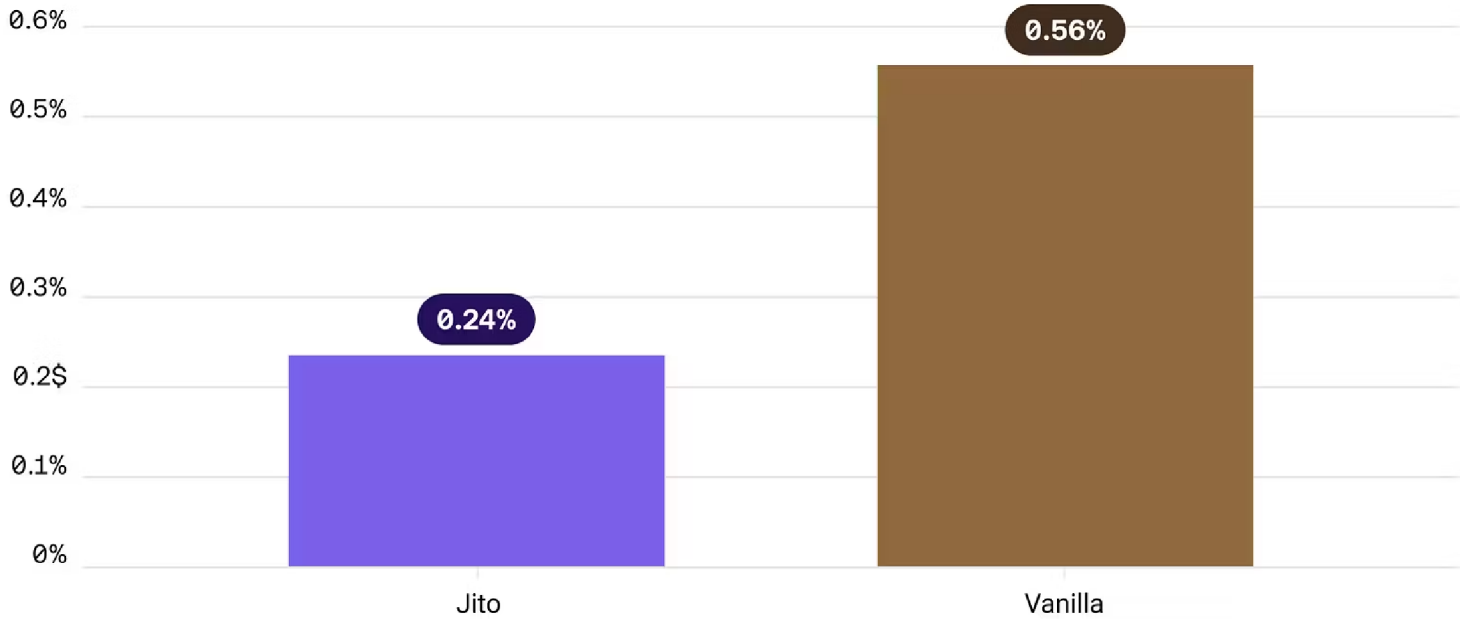
From January to August 2025, Agave validators achieved a 0.23% skip rate compared to Frankendancer's 0.50%.

Jito-enabled validators achieved 0.24% while vanilla clients hit 0.56%.

Average Skip Rate: Agave vs Frankendancer  
January-August 2025



Average Skip Rate: Jito vs Vanilla  
January-August 2025






# Jito integration cuts skip rates by half.

Validators running MEV-enabled clients were twice as reliable from January to August 2025.


Agave Jito (0.22%) and Frankendancer Jito (0.43%) both achieved skip rates less than half of their vanilla counterparts.

In August, Agave Paladin was the top performer with a 0.07% skip rate.

Average Skip Rate by Client  
January-August 2025



Average Skip Rate by Client  
August 2025



# Agave consistently outperforms Frankendancer on skip rate.


Agave has maintained a significantly lower and more stable skip rate throughout 2025, improving from 0.7% to below 0.15% and never surpassing 0.4% since March.

While Frankendancer has been volatile with spikes over 1%, the performance gap narrowed in August, with Frankendancer's worst day (0.22%) being just over twice that of Agave.

## Skip Rate: Agave vs Frankendancer

7 day rolling average

Agave Frankendancer




**Jito delivers a more stable and predictable skip rate compared to vanilla clients.**

## Skip Rate: Jito vs Vanilla

7 day rolling average

Jito Vanilla



## Skip rate outliers stand out in otherwise reliable validator set

We focus now on the validators that produced at least 10,000 blocks during the month of August, which resulted in 126 validators and around 75% of total stake.


Among these validators, median skip rate was 0.0478% and five validators did not miss a single slot.

However, 9 validators exceeded 0.5% skip rates, with the worst three being rubyWZkf (1.11%), HLnodbYk (1.54%), and 2mMGsb5u (1.98%).

### Validator Average Skip Rate

August 2025

● Top 3 (Best Skip Rate) ● Normal ● Bottom 3 (Worst Skip Rate)





## Part II.C

# Vote Latency


## Each Agave version votes faster than the last.

The average latency dropped from 1.58 slots on v1.18 to just 1.03 slots on v2.3, with each intermediate version showing progressive improvement.

The likely main driver is Timely Vote Credits (TVC) upgrade implemented on November 26 2024 (Epoch 703).

With TVC, votes within two slots of the voted-on slot earn bonus rewards, and votes cast later than four slots afterwards lose rewards. This incentivizes validators to cast votes quickly, reducing forks and reaching consensus faster, which results in a lower skip rate and improved latency.

### Vote Latency by Agave Version January-August 2025

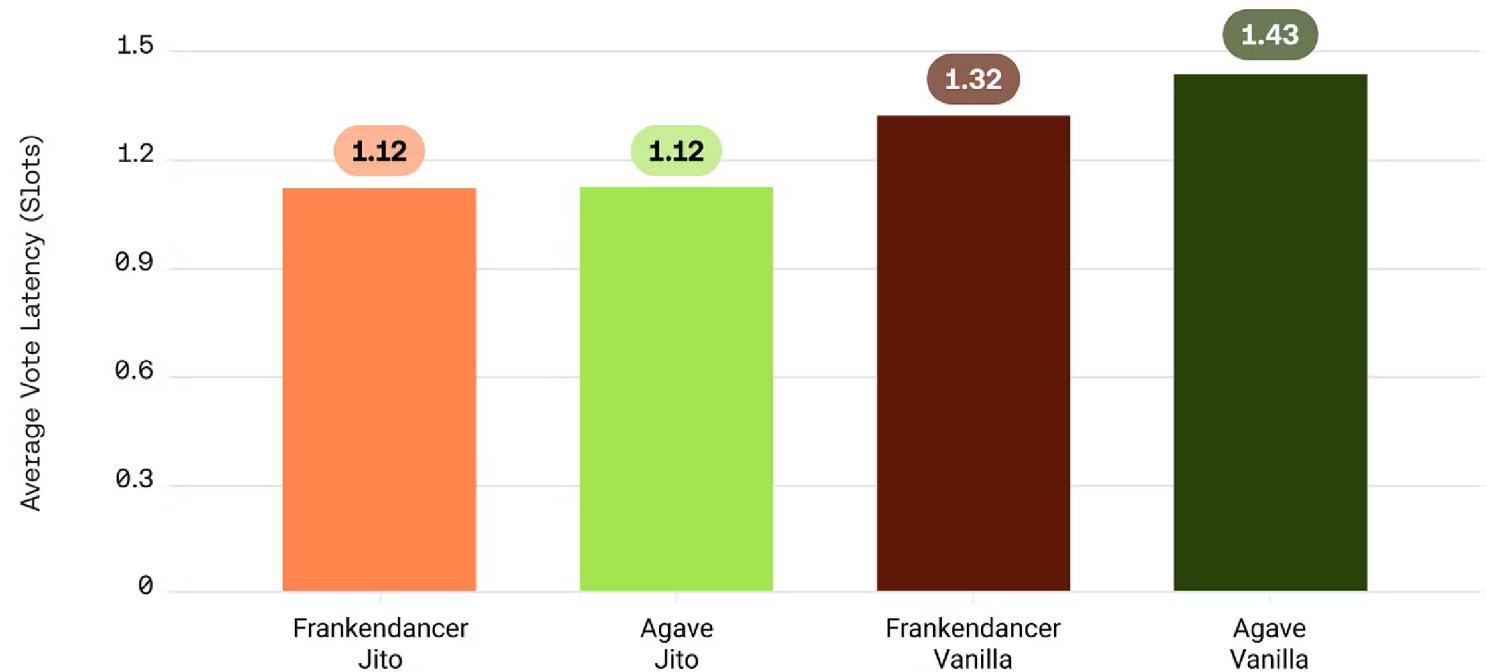


# MEV-enabled clients have the lowest vote latency


So far in 2025, Jito clients have shown lower vote latency, with Agave Jito and Frankendancer Jito being 20-25% faster than their vanilla counterparts.

In July Agave Paladin recorded the lowest average vote latency at 1.015 slots, just ahead of Agave Jito (1.025 slots) and around 4% faster than Frankendancer Jito.

Average Vote Latency by Client  
January-August 2025



Average Vote Latency by Client  
August 2025



## Vote latencies converge, but Agave remains more stable


While vote latency has improved dramatically across all clients in 2025, Agave clients have proven to be more consistent.

Looking at August data, Agave Jito delivered the most predictable latency, with daily averages fluctuating just  $\pm 2.5\%$  around its monthly mean.

Agave Paladin was twice as volatile as Agave Jito, and Frankendancer Jito five times as much.

### Vote Latency by Client 7 day rolling average

● Agave Jito
 ● Agave Paladin
 ● Agave Vanilla
 ● Frankendancer Jito
 ● Frankendancer Vanilla







# Agave Paladin and Jito achieve p99 vote latency below 2 slots.

Agave Jito improved from 2.80 to 1.71 slots since January, while newcomer Agave Paladin leads at 1.35 slots.


Frankendancer Jito has plateaued at 2-slot latency with periodic spikes, the latest one in August reaching as high as 9 slots.

## p99 Vote Latency by Client



7 day rolling average

Agave Jito Agave Paladin Frankendancer Jito




## p99 Vote Latency by Client



7 day rolling average

Agave Vanilla Frankendancer Vanilla



## Vote latency clusters near optimal value of 1 slot with notable outliers

Among the same 126 major validators producing over 10,000 blocks in August, median vote latency was 1.0139 slots.


The best performers (GoLd56YQ, oPaLtitM, GUdrZwgT) achieved ~1.004 slots. The worst performer was validator 5EhGYUyQ (#21 by stake, 3.6M SOL) at 1.29 slots.

All top and bottom performers run Agave Jito.

### Vote Latency by Validator

August 2025

● Top 3 (Best Latency) ● Normal ● Bottom 3 (Worst Latency)



Part II.D


# Compute Units (CUs)

# Frankendancer and Jito clients build fuller blocks.


Frankendancer clients lead in compute usage, averaging 44M compute units (CUs) per block compared to Agave's 42.5M.

Similarly, Jito-enabled clients outperform their vanilla counterparts, packing 42.8M CUs per block versus 40.6M for vanilla.

Compute Units per Block: Agave vs Frankendancer  
January-August 2025



Compute Units per Block: Jito vs Vanilla  
January-August 2025







# Frankendancer Jito builds the fullest blocks, with Paladin emerging as a top competitor.

Frankendancer Jito consistently packs blocks the fullest, averaging 44.4M CUs in 2025, followed by Agave Jito at 42.7M.

In August, Frankendancer remained the top performer but Agave Paladin edged out the second spot from Agave Jito with 0.2M more CU per block on average.






# Frankendancer's CU performance lead over Agave widens.

While performance was comparable early in the year, with Agave even leading in February, Frankendancer has consistently built 5-10% fuller blocks since May.

## Compute Units per Block: Agave vs Frankendancer

7 day rolling average

Agave Frankendancer Block Limit



## Top 3 validators pack 30% more compute units per block


Among the same 126 validators, the top three validators (84gC25fb, 4mtXJ5pU, GUdrZwgT) averaged 49.5M CU per block in August. That is 11M CU, or 30% more than the median, and 6M CUs more than the fourth-ranked validator.

The top 10 includes 7 Agave Jito and 3 Frankendancer Jito clients, while all 3 Agave Vanilla validators rank among the worst 10.

### Compute Units per Block by Validator

August 2025

● Agave Jito
 ● Agave Paladin
 ● Agave Vanilla
 ● Frankendancer Jito
 ● Frankendancer Vanilla




Part II.E

# Transactions per Block


# Frankendancer and Jito pack more user activity into blocks.

Since January 2025, Frankendancer validators have averaged 505 non-vote or user transactions in each block. That's 12% more than Agave validators, approximately the same difference between Jito versus vanilla validators.

Non-Vote Transactions per Block: Agave vs Frankendancer  
January-August 2025



Non-Vote Transactions per Block: Jito vs Vanilla  
January-August 2025



# Frankendancer's efficiency packing blocks is clear across all versions.


Frankendancer Jito led all clients in August with 502 user transactions per block. That's 16% and 18% more than Agave Paladin and Jito, respectively.

More impressively, Frankendancer Vanilla performed just 2% worse than Agave Jito and 20% better than its direct counterpart, Agave Vanilla.

Non-Vote Transactions per Block: Agave vs Frankendancer

Syndica 


January-August 2025



Non-Vote Transactions per Block: Jito vs Vanilla

Syndica 

August 2025





## Frankendancer packs over 10% more user transactions in each block since Jito becomes default.


Frankendancer began to consistently outperform Agave in March once its Jito version surpassed vanilla in popularity.

Since that turning point, Frankendancer has maintained a lead of at least 10% more user transactions per block, and averaged 19% more in August.

### Non-Vote Transactions per Block: Agave vs Frankendancer

7 day rolling average

Agave Frankendancer



## Top compute performers also dominate user transaction throughput.

The top three validators for user transactions are identical to the compute unit leaders: 84gC25fb, 4mtXJ5pU, and GUdrZwgT.

At 636 non-vote transactions per block, they pack 50% more than the median and 20% more than the fourth-ranked validator (Temporal's Ruby validator running Frankendancer Jito).

The top 10 splits evenly between Agave Jito and Frankendancer Jito (5 each), while the bottom three performers are all Agave Vanilla clients.

### Non-Vote Transactions per Block by Validator

August 2025

Agave Jito Agave Paladin Agave Vanilla Frankendancer Jito Frankendancer Vanilla

