Obol

x

LIDO

Performance Report

March 23rd - May 20th, 2023

Introduction

Obol Labs is excited to share the performance report of our testing effort with Lido. Distributed Validator Technology (DVT) is one of the key technology primitives that will allow Lido to safely scale their node operator set as it becomes permissionless.

DVT allows an Ethereum PoS Validator to operate across multiple nodes, creating byzantine fault-tolerant node clusters with active/active redundancy. By enabling high-availability deployments for validator operations, DVT can improve the safety and liveness of Ethereum's staking ecosystem while mitigating single-operator risks. For more benefits on DVT, check out this post.

At Obol, we have been actively testing with the Lido team to further DVT's integration into their validator stack. Our testing period spanned 59 days, from March 23rd to May 20th, 2023. During this time, we focused on key metrics for our Lido clusters, also drawing comparisons with industry peers. We're excited to share that the data displays strong performance of our DVT clusters as we continue to improve and enhance our middleware client, Charon.

Lido conducted their testing with 10 clusters of different sizes. Each of these clusters was running 5 validators per cluster. The specific sizes are as follows:

- 4-node cluster: Allows for 1 node to be offline.
- 7-node cluster: Allows for 2 nodes to be offline.
- 10-node cluster: Allows for 3 nodes to be offline.

At Obol, we have successfully completed our Bia Testnet with exceptional participation and results. Simultaneously, we initiated our Mainnet Alpha, collaborating with our operator community to introduce the first distributed validators on Ethereum mainnet. We've also established a partnership with Nethermind for researching and developing a distributed validator protocol specification for Obol V2.

This performance report provides a comprehensive selection of key metrics and data points, illustrating our validators' performance on Ethereum's Göerli testnet. As we refine our software for Mainnet Beta and our forthcoming V1 launch, these results lay out a very optimistic outlook.

Obol

Contents

Obol

Lido DV Clusters Overview

Obol	Cluster 1: "Quadforce"	Cluster 2: "Group2 Guardians"	Cluster 3: "dvt2-g3"	Cluster 4: "Group 4"	Cluster 5: "Group 5 Rangers"
Network	Göerli	Göerli	Göerli	Göerli	Göerli
# of Validators	5	5	5	5	5
# of Operators	4	4	4	7	7
Mean Threshold Latency	242 ms	89 ms	57 ms	67 ms	202 ms
Operators	• Hashquark • DSRV • Blockseal • Stakely	• Simply VC • Kukis Global • Mav3rick • SenseiNode	• Mahof • Cosmostation • GraphOps • Amonxx	• CryptoManufaktur • Chorus One • Nethermind • Everstake • Uniqlabs • Luganodes • Blockscape	• Kukis Global • Bellatora • a41 • D-Stake • Stakely • Staking Facilities • Everstake
Dashboard	Beaconcha.in				

Obol

Lido DV Clusters Overview

Obol	Cluster 6 "SevenNodes"	Cluster 7: "Group 7"	Cluster 8: "Clustery McClusterfaces"	Cluster 9: "Group 9 Hunter"	Cluster 10: "Mostly Harmless"
Network	Göerli	Göerli	Göerli	Göerli	Göerli
# of Validators	5	5	5	5	5
# of Operators	7	7	10	10	10
Mean Threshold Latency	125 ms	21 ms	439 ms	95 ms	63 ms
Operators	• SenseiNode • NodesKuge • Secard Node • Wallclimbr • Spacesider • Oisin (Obol) • Imperator	• P2P • P-OPS • Minivipers • Alkadeta • DSRV • Blockscape • Chorus One	• Archimedes • Coinstamp • Cryptomanufaktur • H20 Nodes • Hashquark • Kukis • Nethermind • Simply Staking • Staking Facilities • Talha	• Chorus One • DSRV • P2P • Hashquark • Everstake • TheNop.io • Polkachu.com • Wabut.club • Farukyasar • Staking4all	• Swiss Staking • Piconbello • Yesaynow • Eridian • Luke (Obol) • Smart Node Capital • Furkan • Obol Ar Line • Sodiumstar • Nakoturk
Dashboard	Beaconcha.in				

Obol

Nodes Geographical Distribution

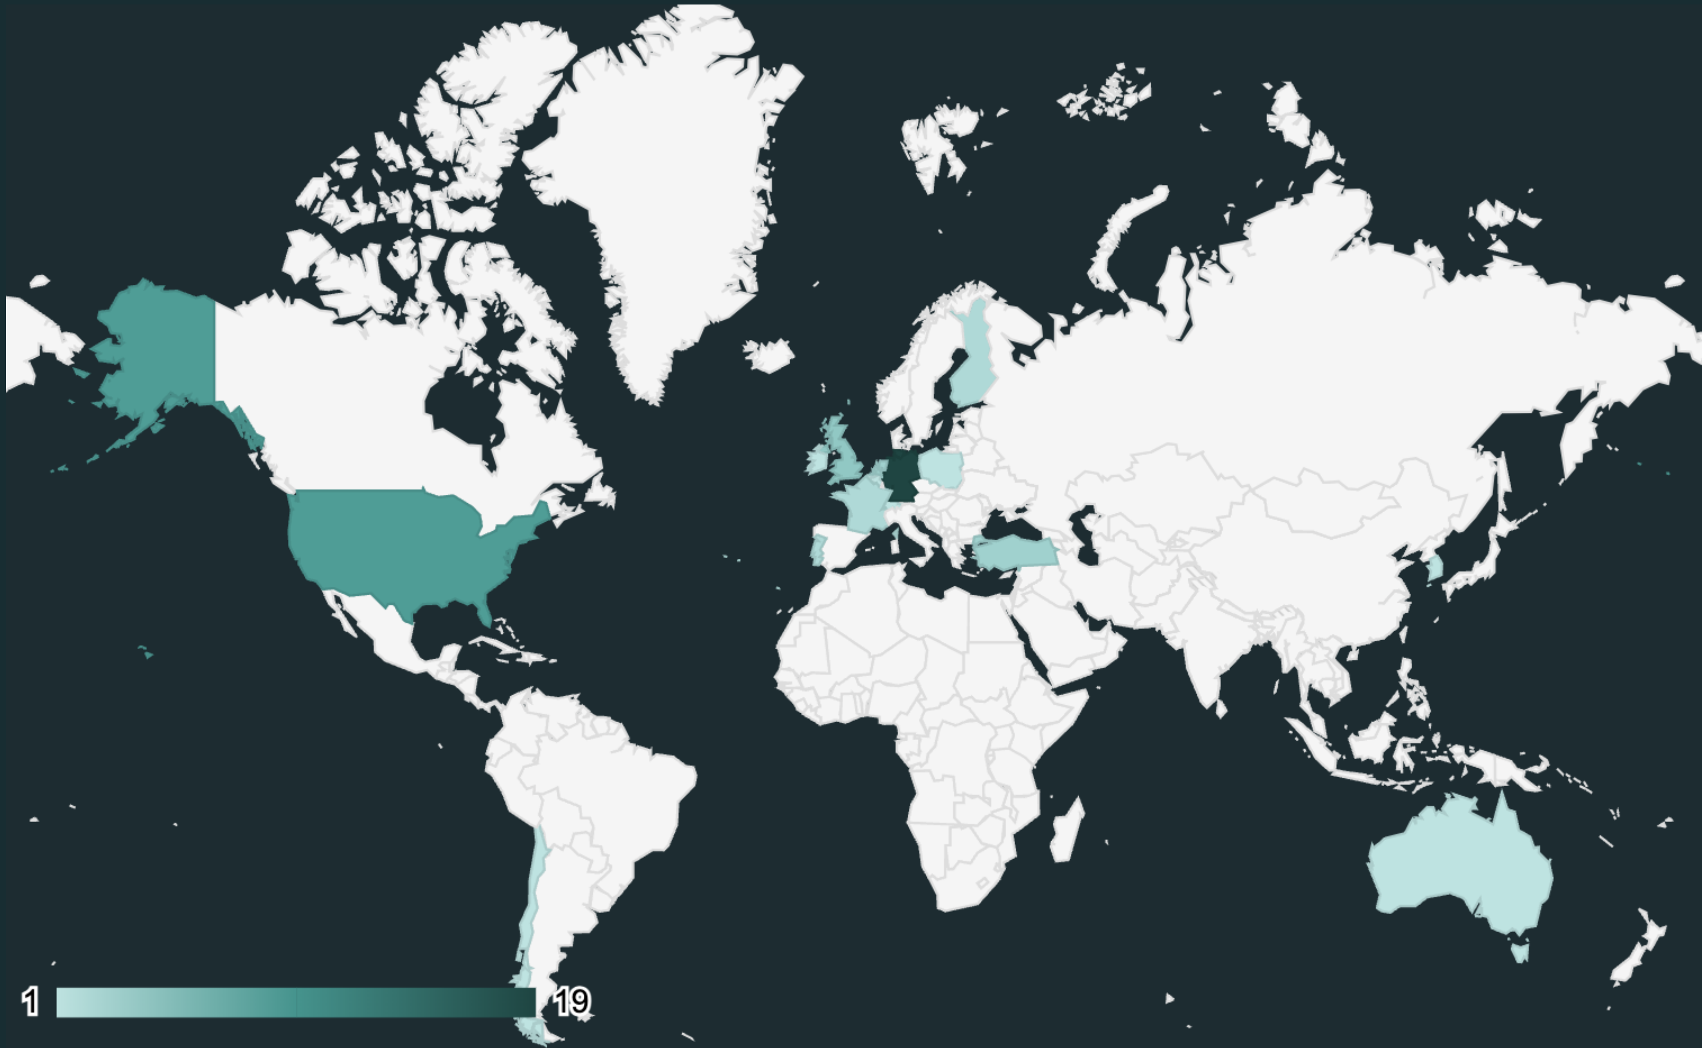

CPU Types: Virtualized or Physical Cores

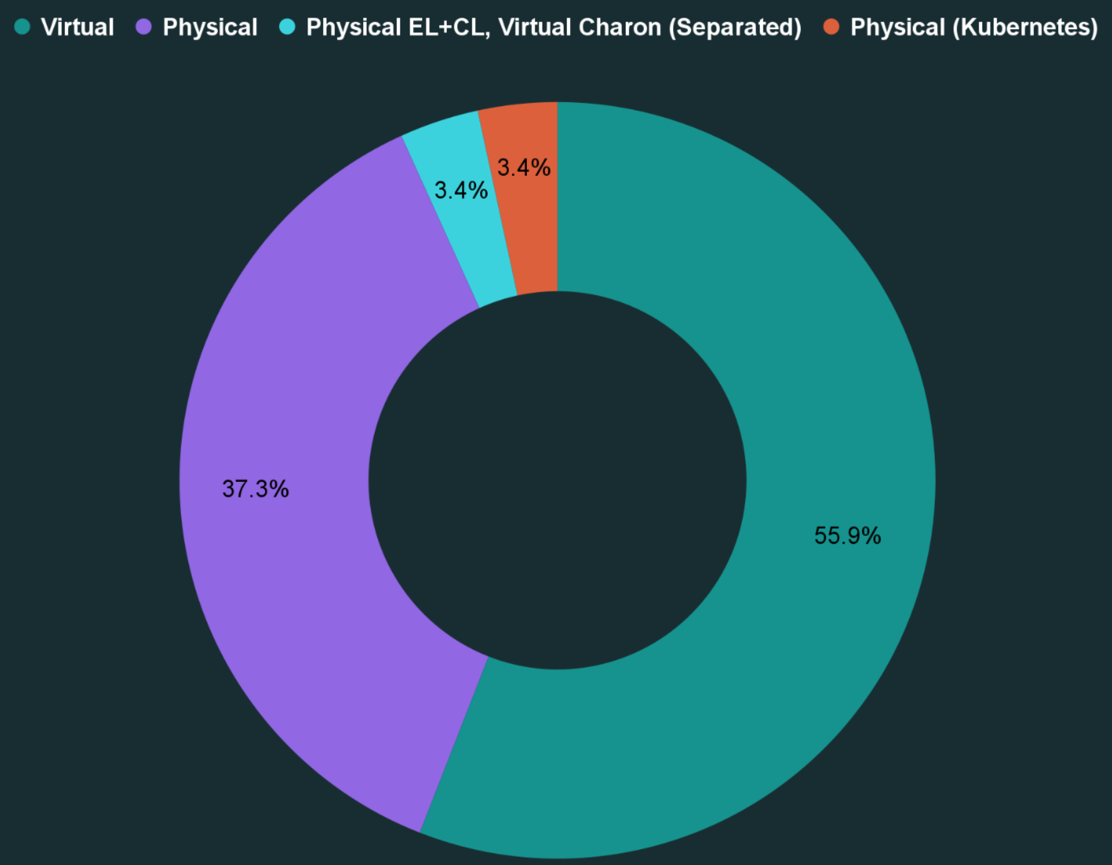

Latency among peers in a cluster

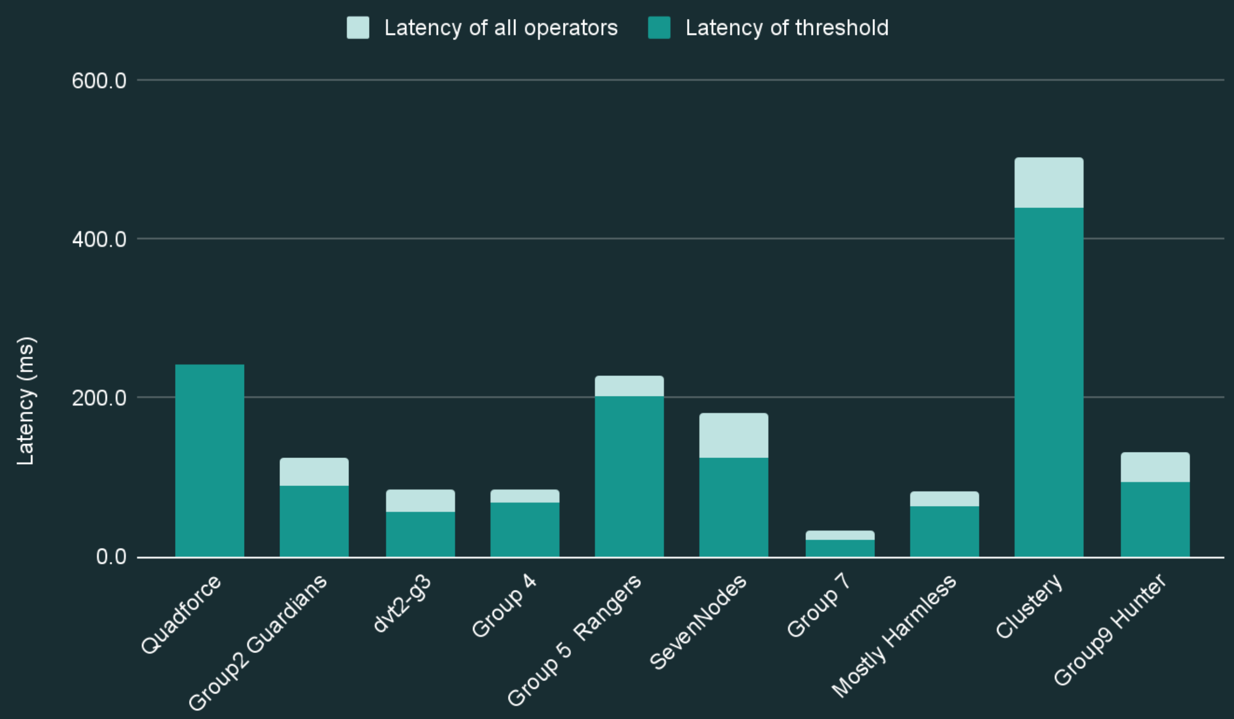

The chart displayed above portrays the average latency between peers across the different clusters. Performance can degrade when values exceed 250 ms, leading to an increased likelihood of missed validator duties by a cluster. The average latency across all clusters stood at 140 ms. Cluster 8, being a heavily geo-distributed cluster, recorded the highest average latency of 439 ms.

Obol

Lido DV Clusters Performance Metrics (Göerli)

Metric	Obol Lido Pilot	Attestant Benchmark Operator
Number of Validators	50	2010
Slashings Received	0	0
Avg. Uptime	99.2%	99.5%
Avg. Inclusion Distance	1.72	1.71
Avg. Attester Effectiveness	53.3%	53.7%
Avg. Proposer Effectiveness	57.8%	96.9%
Avg. Validator Effectiveness	53.3%	53.9%

Time-Series Charts for Selected Metrics

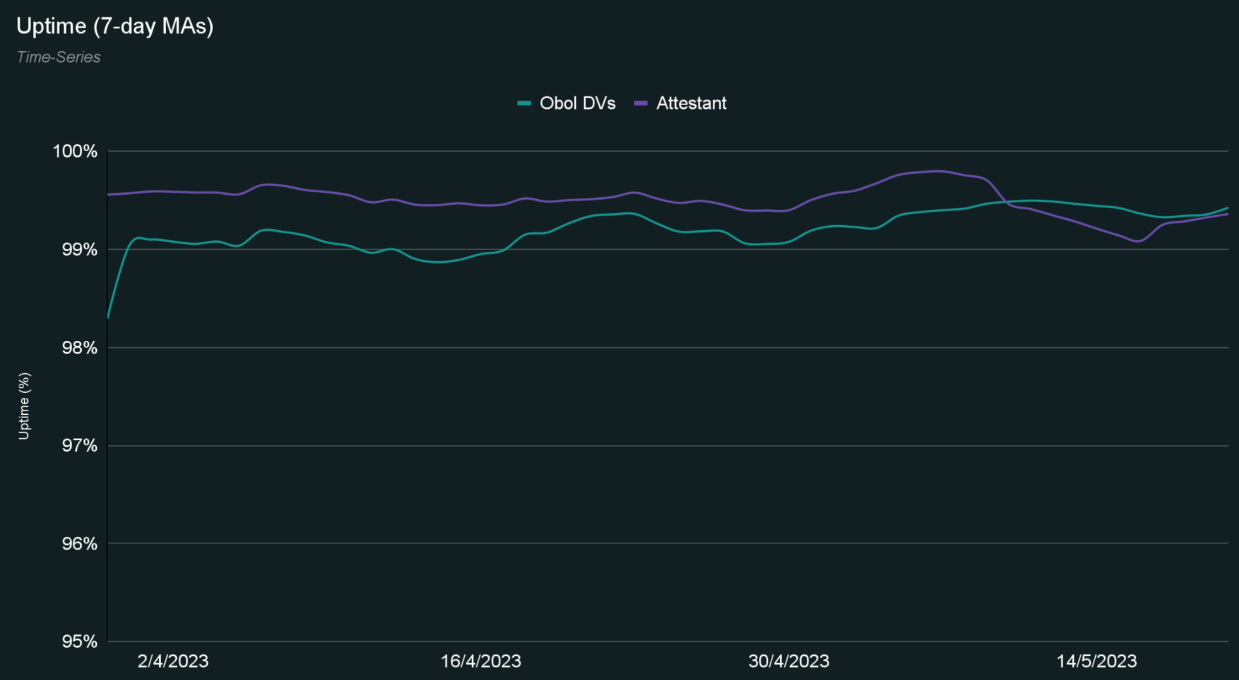

Figure 1

Figure 1 displays a time-series for uptime of the 10 DV clusters. Uptime is the percentage of time a validator's attestation is successfully included on-chain, relative to a validator's total active time in the network. Throughout our testing phase, we observed an exceptional average uptime of 99.2%. In comparison, our benchmark operator, Attestant, had an average network uptime of 99.5%, a marginal difference of 0.3%. Towards the end of the testing phase, DV clusters surpassed Attestant's average uptime, a very encouraging result showcasing how performant DV clusters can be.

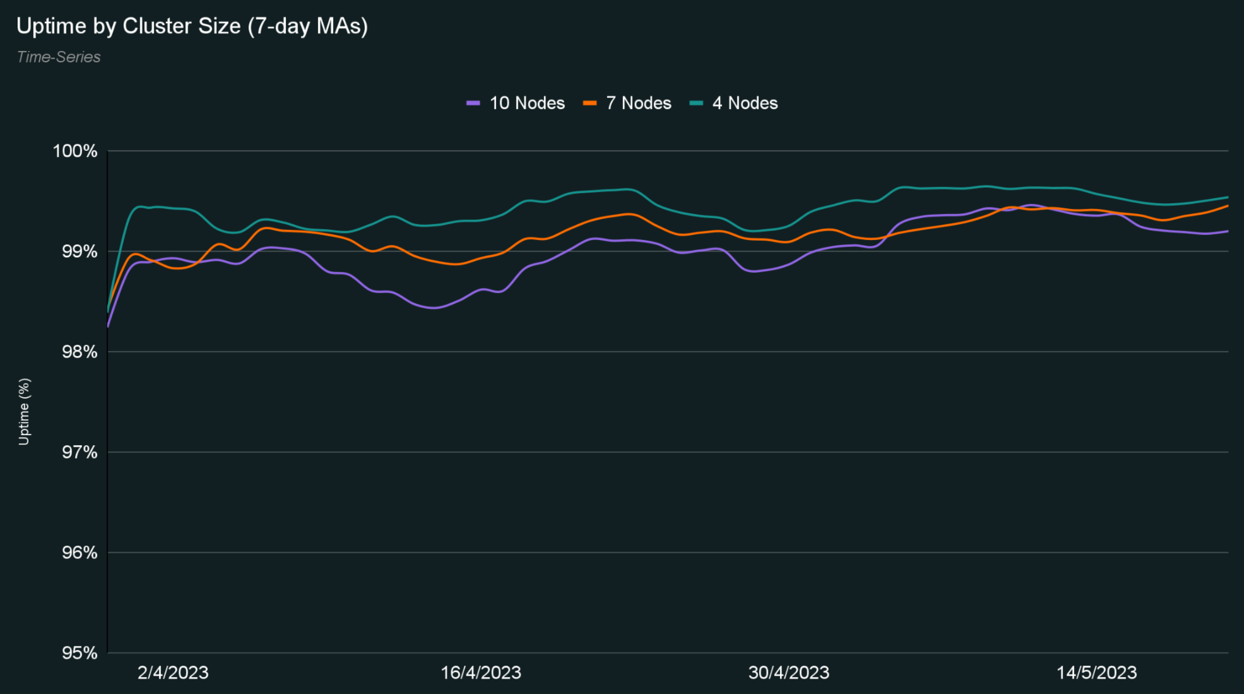

Figure 2

Figure 2 displays a time-series graph of uptime, segmented by cluster size. During testing, decreased uptime was observed in clusters with larger node counts. This situation improved towards the end of the testing effort. Still, large clusters experienced slightly degraded performance than 4-node clusters. Specifically, one of the 10-node clusters, heavily geo-distributed, can help explain the degraded performance (mean latency 439ms, roughly triple the average latency).

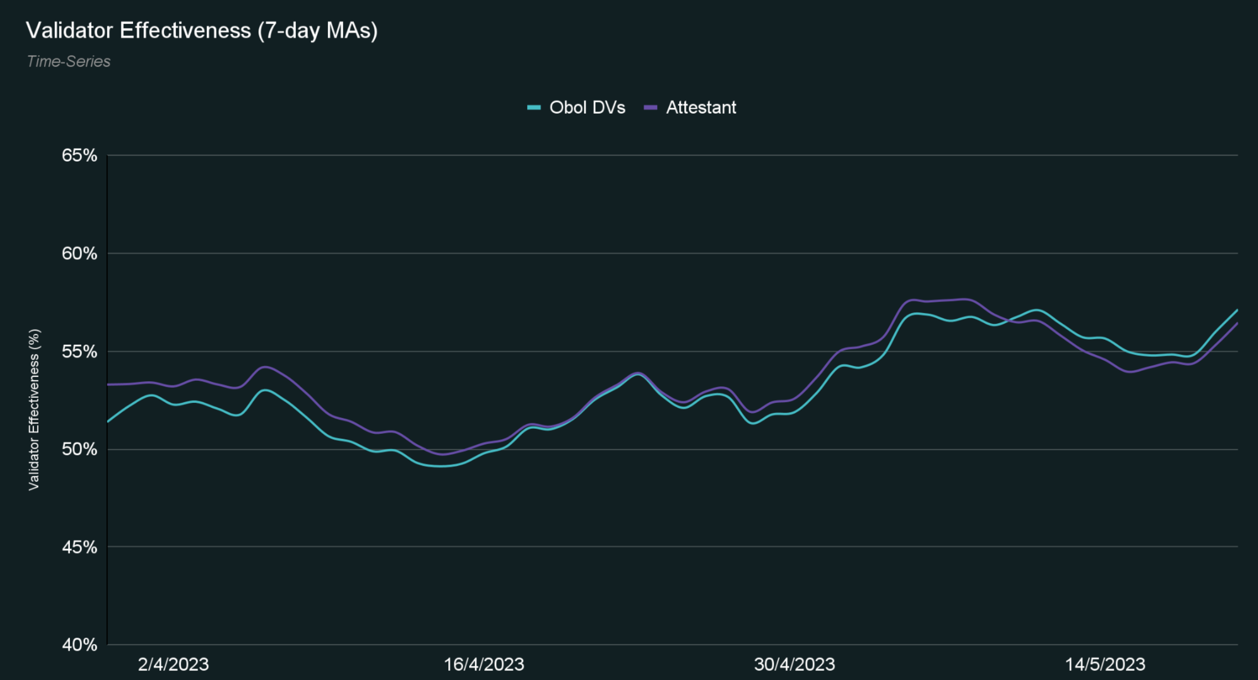

Figure 3

Figure 3 presents a time-series graph charting validator effectiveness. This metric, curated by Rated Network, measures a validator's reliability related to its duties by examining proposer and attester effectiveness along with the validator's slashing record. Our DV clusters exhibited an average effectiveness of 53.3% throughout the testing phase, displaying a slight difference of 0.6% compared to our benchmark operator, Attestant.

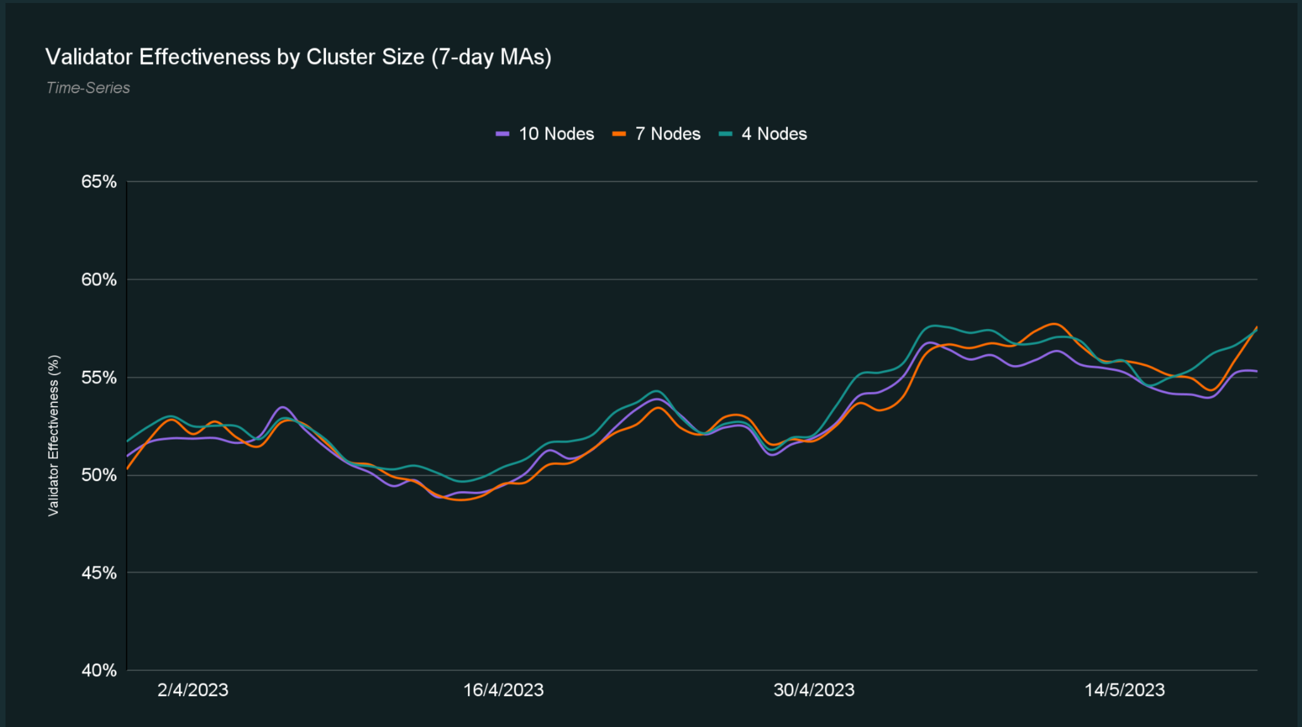

Figure 4

Figure 4 presents a time-series graph depicting validator effectiveness, segmented by cluster size. Across various cluster sizes, the effectiveness remains remarkably consistent. Notably, the 4-node cluster exhibits a slightly superior performance compared to other configurations for the majority of the testing phase.

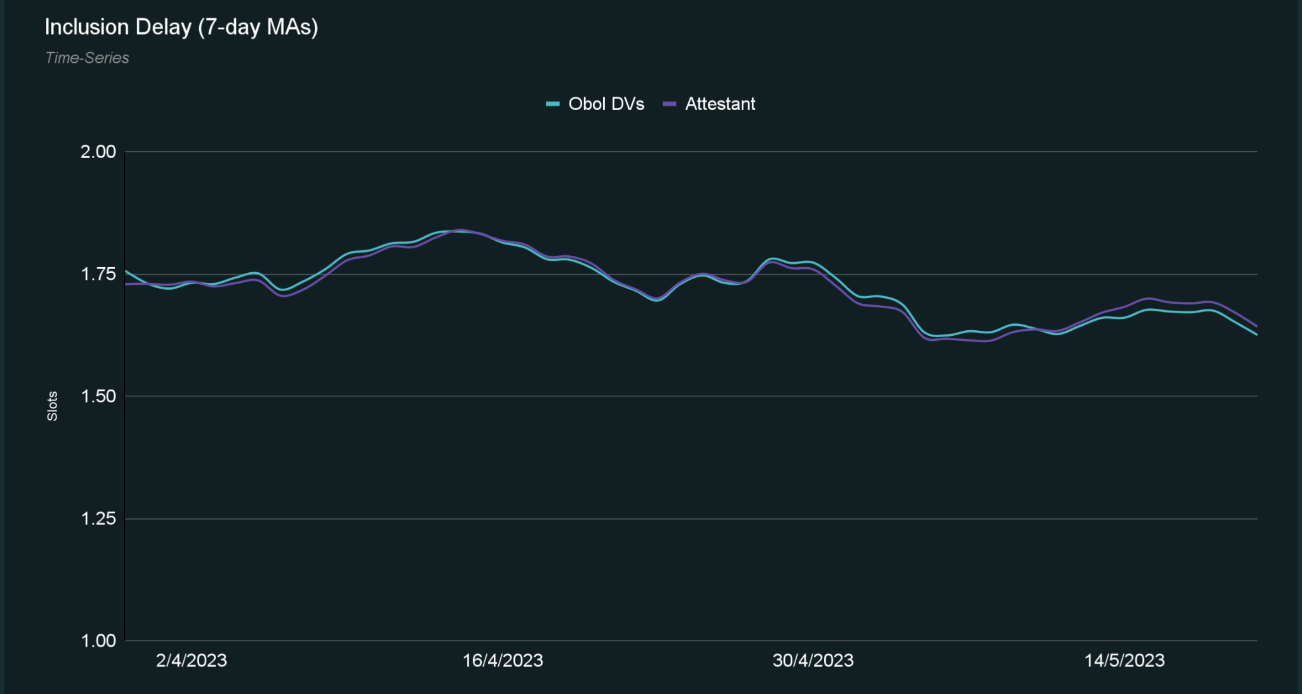

Figure 5

Figure 5 displays a time-series graph of inclusion delay. Inclusion delay is the aggregate slot distance between the attestation slots attributed and the actual slots the votes were included in. The inclusion delay closely mirrors our benchmark operator, Attestant, displaying a negligible disparity—an exceptionally encouraging result.

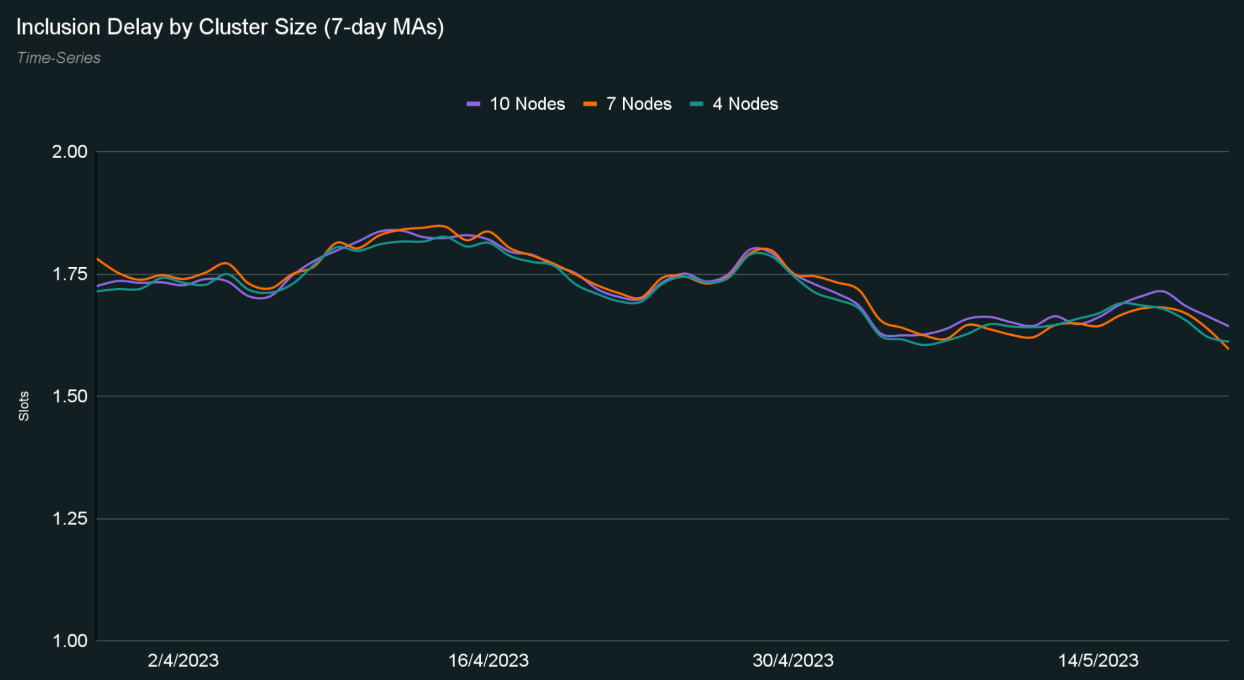

Figure 6

Figure 6 showcases a time-series graph of inclusion delay, segmented by cluster size. This graph exhibits a pattern akin to the validator effectiveness chart, revealing no significant discrepancies across clusters of varying size - indicative of robust results across all cluster sizes.

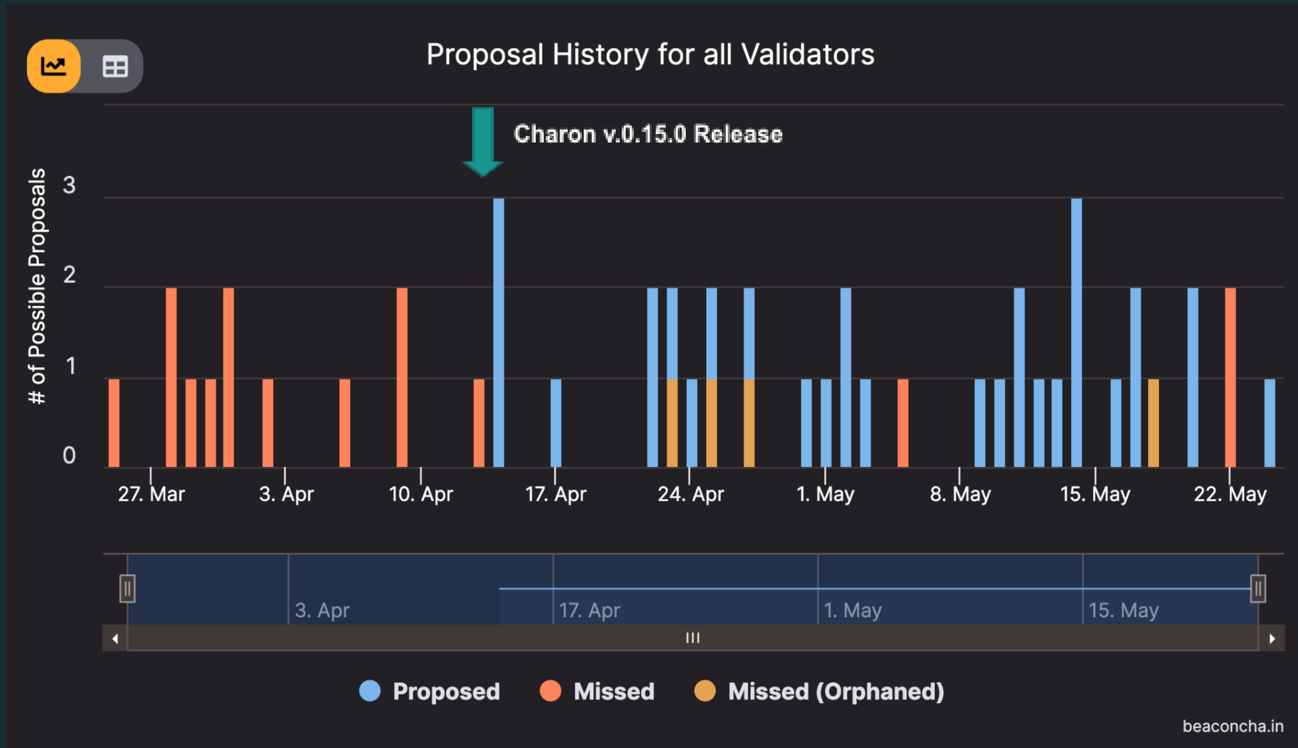

Figure 7

Figure 7 displays a time-series graph of the proposal history of our validators. On April 12, Charon v.0.15.0 was released, completing the QBFT wire protocol migration resulting in a substantial decrease in bandwidth requirements for block proposals. This addressed sporadic missed block proposals in large geo-distributed clusters.

Data Considerations

- *Obol Data was retrieved using Rated Network's API (/v0/eth/validators/{validator_index_or_pubkey}/effectiveness) for days 731 to 795 since genesis on Prater/Görli.*
- *Data from Attestant, our benchmark node operator, was retrieved using Rated Network's API (/v0/eth/operators/{operator_id}/effectiveness) for days 731 to 795 since genesis on Prater/Goerli.*
- *All metrics are as is from Rated Network or deduced from simple maths definition. See Rated docs and definitions here.*
- *Pubkeys and Raw data for Obol's Lido Pilot and Attestant can be accessed freely here.*

Supported Duties

Obol's middleware, Charon, is a distributed validator middleware – it works to coordinate validators to work in a distributed manner with fault tolerance. The following table outlines which validator clients have produced which duties on a Public Testnest, and which are still under construction (🚧)

Stay up to date with Obol's support for all validator clients at https://dvt.obol.tech

Duty \ Client	Teku	Lighthous e	Nimbus	Prysm	Lodestar	Vouch
Attestation	✅	✅	✅	✅	✅	✅
Attestation Aggregation	🚧	🚧	🚧	🚧	✅	🚧
Block Proposal	✅	✅	✅	🚧	✅	✅
Blinded Block Proposal (MEV-boost)	✅	✅	🚧	🚧	🚧	🚧
Sync Committee Attestation	✅	✅	✅	🚧	✅	✅
Sync Committee Aggregation	🚧	🚧	🚧	🚧	✅	🚧