Obol

Simple DVT on Holesky Testnet

Performance Report

LIDO

28th Nov 2023 - 11th Jan 2024

Contents

1. Introduction 3
   - Why DVT matters to Lido 3
   - Summary of Testing Results 4
   - Summary of Key Learnings & Recommendations 5
2. In-depth Performance Review 6
3. Latency causes and implications 12
4. MEV-Boost and Relays 17
5. Hardware 19
6. Notes and References 21

1. Introduction

Obol Labs is excited to share this performance report detailing our testing effort with Lido on Holesky testnet. **Overall, Obol's DV middleware satisfied the performance requirements to move forward with a mainnet deployment of Lido's Simple DVT module.**

Why DVT matters to Lido

Distributed Validator Technology (DVT) is a standard security primitive that allows a single validator to be run across multiple machines. Lido will use DVT as a technology primitive to permissionlessly scale their node operator set - all while significantly reducing risks and improving Lido's resilience, scalability, and decentralization. These features boost participation and collaboration by allowing smaller operators to join forces with larger ones, fostering a more diverse and resilient network. This inclusivity paves the way for a trustless future where even at-home validators can join large well-established staking protocols like Lido.

More specifically, the use of DVT reduces correlation risk by enabling client diversity and geographical dispersion within a distributed validator cluster. DVT also reduces slashing risks by making operators only responsible for a key share, instead of an entire validator key. Validator downtime is also minimized, because only a threshold of keys are required to be online for the validator to stay active. (i.e. 3 of 4, 5 of 7, 7 of 10, etc.)

Ultimately, DVT democratizes stake in Lido, reducing entry barriers for smaller validators and moving towards a more decentralized, trustless staking environment, as seen in the adoption plans of major LSPs like Lido and Stakewise.

To learn more about the benefits of DVT for liquid staking protocols, read our blog post.

Summary of Testing Results

The Simple DVT testing program, conducted by the Obol and Lido teams on Ethereum's Holesky testnet, spanned 69 days from November 3rd, 2023, to January 11th, 2024. The metrics presented here were derived from a 45-day monitoring period starting on November 18th, 2023. Throughout this monitoring phase, the aggregate metrics of Obol's DVT clusters surpassed both Lido's minimum requirements and the Holesky network averages, as shown in the table below.

Metric	Simple DVT Clusters	Satisfies Mainnet Proposal?	Holesky Network Average
Number of Validators	3100		1.5M
Slashings Received	0	✅	2
Avg. Uptime	97.9%	✅	93.9%
Avg. Attester Effectiveness	84.5%	✅	78.8%
Avg. Proposer Effectiveness	71.2%	✅	87.9%
Avg. Validator Effectiveness	84.5%	✅	78.9%

Testing included 32 DV clusters utilizing a "5-of-7" threshold configuration. Initially, this included 214 participants, which later fell to 196. Each cluster initially ran 5 validators, a number that was subsequently scaled up to 50 and then 100 for most clusters, except for "Glacial Gull" and "Crimson Coyote", which remained with 50 validators due to performance constraints. The testing concluded with all DV clusters exiting on January 22nd, with node operators subsequently claiming their rewards with the Obol Splits smart contracts. These results show that the mainnet deployment of the Simple DVT module can safely and efficiently onboard a significant number of new Node Operators to the Lido operator set.

Throughout the testing, various challenges were encountered, primarily stemming from infrastructure limitations attributable to the nascent stage of the Holesky testnet. The absence of established community infrastructure components such as SAFE, 0xSplits, The Graph, MEV-Boost, and MEV relay infrastructure further compounded these challenges. However, despite these hurdles, the testing managed to achieve a Block Proposal Success Rate of 71.2% during the monitoring period from November 28th to January 11th, surpassing the targeted 70% and eventually **exceeding 85% during the final week of observation**.

Summary of Key Learnings & Recommendations

Ensure Beacon Node Health: The Charon middleware exhibited no inherent flaws but underscored the importance for operators to tailor their node settings adeptly. The beacon node plays a critical role in the network's operation. It is imperative to ensure the continuous health and stability of the beacon node to maintain overall network integrity and functionality.

Maintain Low Intra-cluster Latency: Challenges with latency were particularly pronounced for operators located in Asia and Australia, as well as for multi-continent clusters, impacting their consensus times and overall performance. It is recommended to keep intra-cluster latency below 235 milliseconds for all nodes. In cases where latency exceeds this threshold, efforts should be made to reduce physical distance between nodes or optimize Internet Service Provider (ISP) settings accordingly.

Provision Sufficient Hardware, ideally Bare Metal or on Prem: Some participants' inadequate hardware (low-performance virtualized servers and/or slow HDD storage) hindered performance, indicating the necessity of provisioning adequate resources. CPU clock speed and Disk throughput/latency are the two most important factors for running a validator. Note that the Charon client itself takes less than 500MB of RAM and minimal CPU load. In order to optimize both performance and cost-effectiveness, it is recommended to prioritize physical over virtualised setups. Such configurations typically offer greater performance and minimize overhead associated with virtualisation, contributing to improved efficiency and reliability.

Carefully Select MEV-Boost Relays: Integration of MEV-Boost and relay functionalities, especially with newly introduced support on Holesky, highlighted potential optimizations to mitigate latency and relay performance disparities. More research is needed in this area but it is recommended that a thorough analysis of performant, low latency relays be done before selecting which ones to include as part of your DV node configuration.

Lido Exit Sidecar Optimization: Testing also facilitated optimization of the Lido Exit Sidecar feature, presenting potential applications beyond its initial scope.

Obol

2. In-depth Performance Review

Approximately 40% of the performance variability among clusters can be attributed to a few specific factors, which are covered in-depth in the following section. Most of these factors are also relevant to non-DV validators, but their implications can be exacerbated in a DV setup:

- **Latency:** The latency within a DV cluster serves as an important predictor of performance. High latency across cluster nodes prolonged the consensus process, heightening the risk of missing participation in duties.

- **Hardware:** The use of performant hardware, — preferably 8 physical cores, 32GB of RAM — and avoiding HDD storage entirely, can markedly enhance cluster performance. Considering CPU clock speeds such as 2.5 GHz and 3.5 GHz would significantly enhance performance as well.

- **Beacon Node:** Although the Obol DVT solution is confined to a middleware, Charon, it remains crucial for operators to maintain a robust beacon node. The efficacy of Charon in propagating duties across the network hinges on the health and performance of its connected beacon node(s).

- **MEV-Boost and Relays:** During the initial phase of the testnet, Holesky provided limited MEV-boost support. Furthermore, this support was contingent on the availability of only a few relays, some of which exhibited artificial behaviours such as overbidding and suboptimal geographical placement. This limitation underscores the need for a comprehensive evaluation of MEV-boost capabilities and relay performance within the network.

The below figure combines the above factors into a cluster score, which explains 40.6% of the variability in cluster performance:

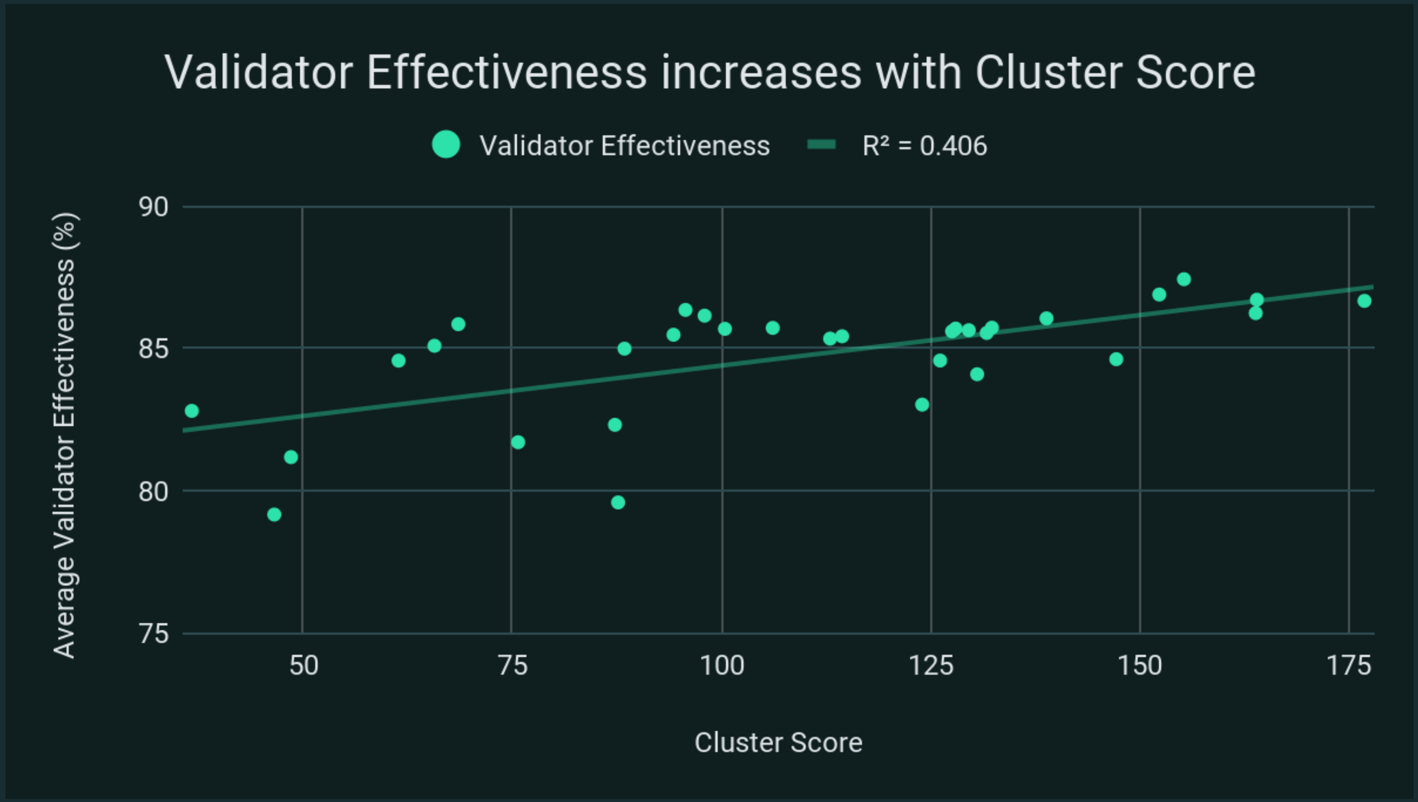

The cluster score is based on each cluster's rankings in the following categories: average latency to other peers in the cluster, latency to MEV-relays, quality of hardware, performance of the beacon node, and number of cluster peer impairments suffered (impairments due to offline nodes, or disconnected Validator Clients or Beacon Nodes).

Considering the above factors, it is easier to understand why some clusters performed much better than others, especially on the proposal rate. As shown below, multiple clusters did achieve a 100% block proposal rate for the majority of the observation period, demonstrating that a 100% proposal rate is possible with the right setup. We can also notice the drop in performance around the holiday period which may indicate the need to keep investing in alerting systems. We are pleased to see that Lido is investing in those thanks to their [infrastructure grants for Simple DVT](#).

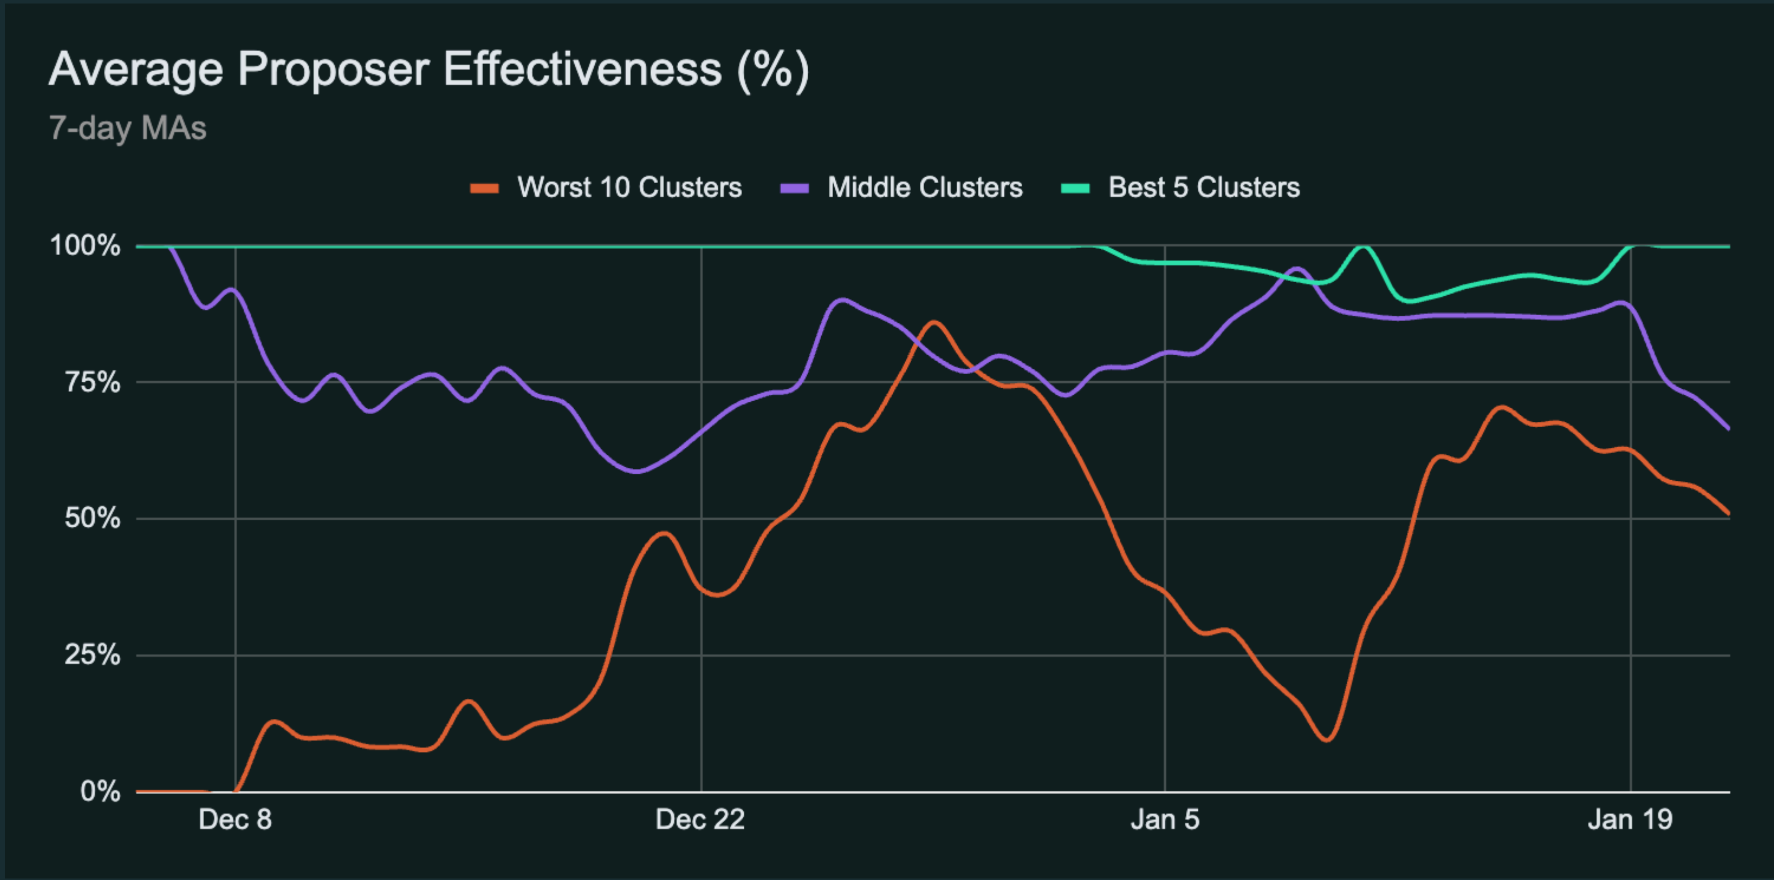

The main reasons for missed proposals included the following:

- **Chain Inclusion Failure:** As shown in the pie chart below, 76.1% of missed block proposals were because latency exceeded the maximum value of 4 seconds required for timely inclusion of proposals. Latency when communicating with MEV-Boost relays led to long wait times before starting consensus, and latency between peers led to long consensus times.
- **Threshold Not Reached:** 23.9% of missed block proposals were due to instances of multiple offline nodes causing the cluster to fail to reach the threshold required to perform duties. Some of this downtime was due to validator client misconfigurations, which prohibited these nodes from taking part in consensus, highlighting the significance of properly configuring nodes during the cluster setup process.

The chart below shows the proportion of missed proposals attributed to these two reasons:

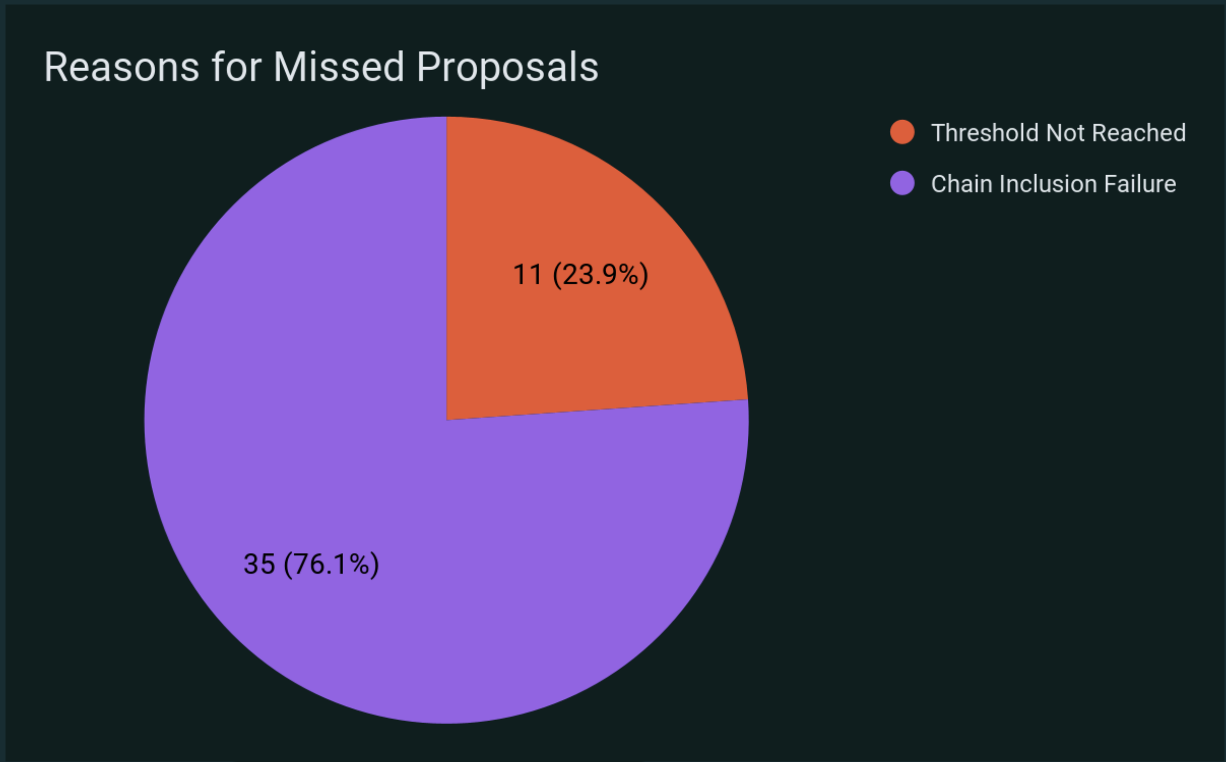

The chart below displays a time-series of average uptime for all clusters that participated in the testing effort. Uptime is defined as the percentage of time a validator's attestation is successfully being included on-chain. Throughout the testing period, we observed an exceptional average uptime of 97.9%. This shows a significant improvement of distributed validators over Holesky's network average of 93.9%.

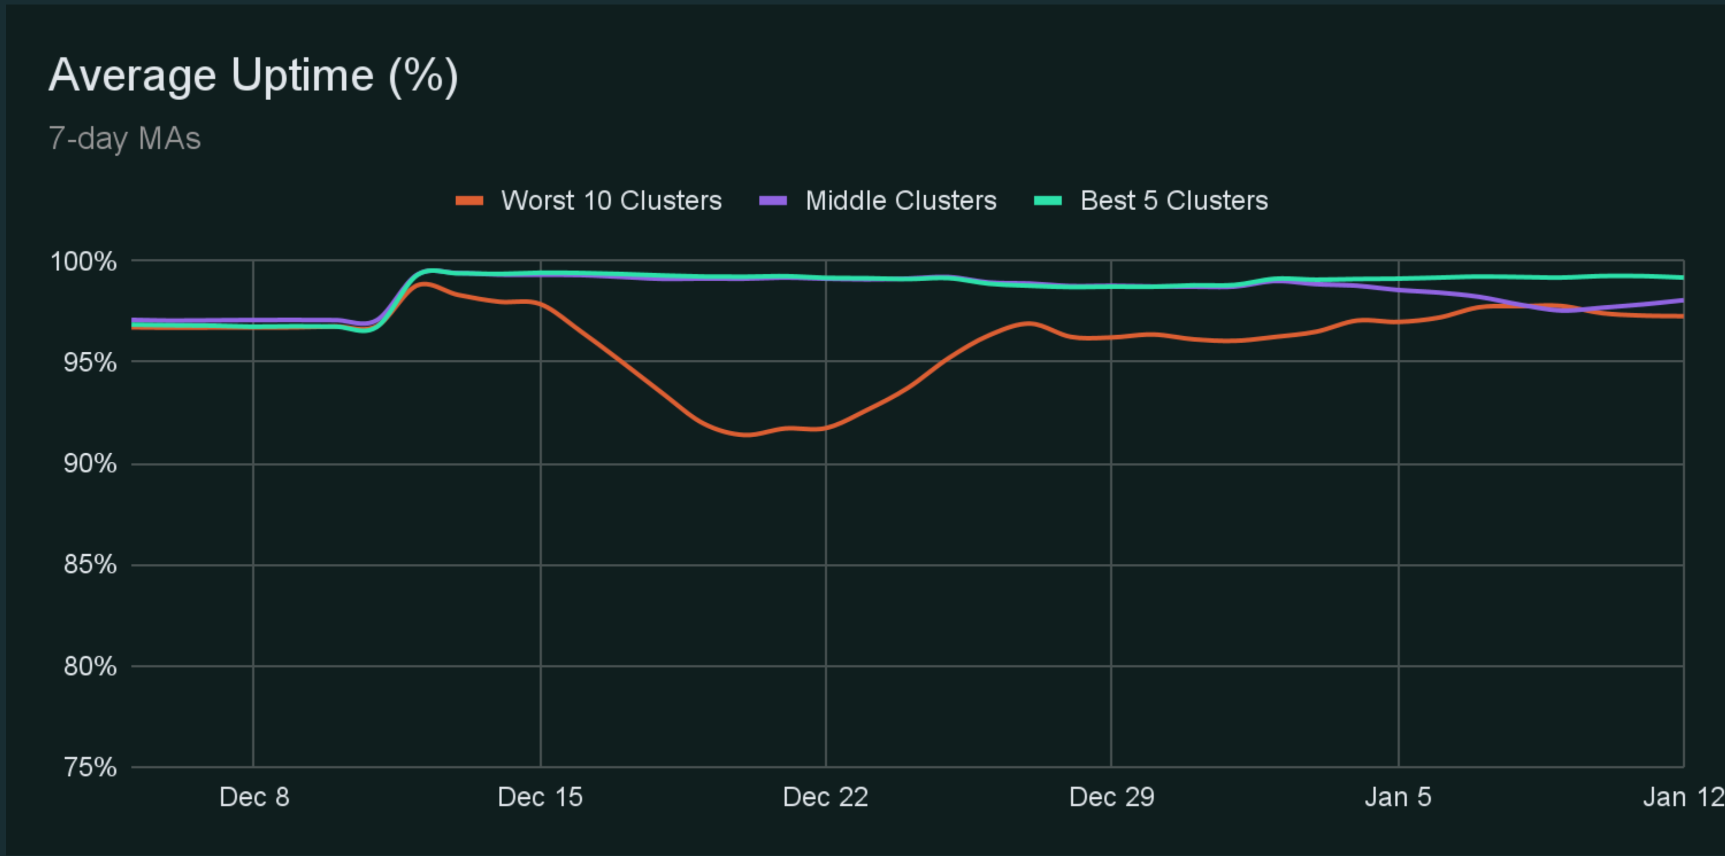

The chart below shows that all DV clusters managed to achieve a high level of average attester effectiveness across the entire testing period, averaging 84.5%, significantly higher than the Holesky network average of 78.8%. Attester Effectiveness is defined in the Rated documentation, as a combination of participation rate, correctness, and inclusion delay. (The declining effectiveness shown in the chart is also reflected in a declining Holesky network average effectiveness and was not attributed to decaying performance on the part of DV clusters in the test group.)

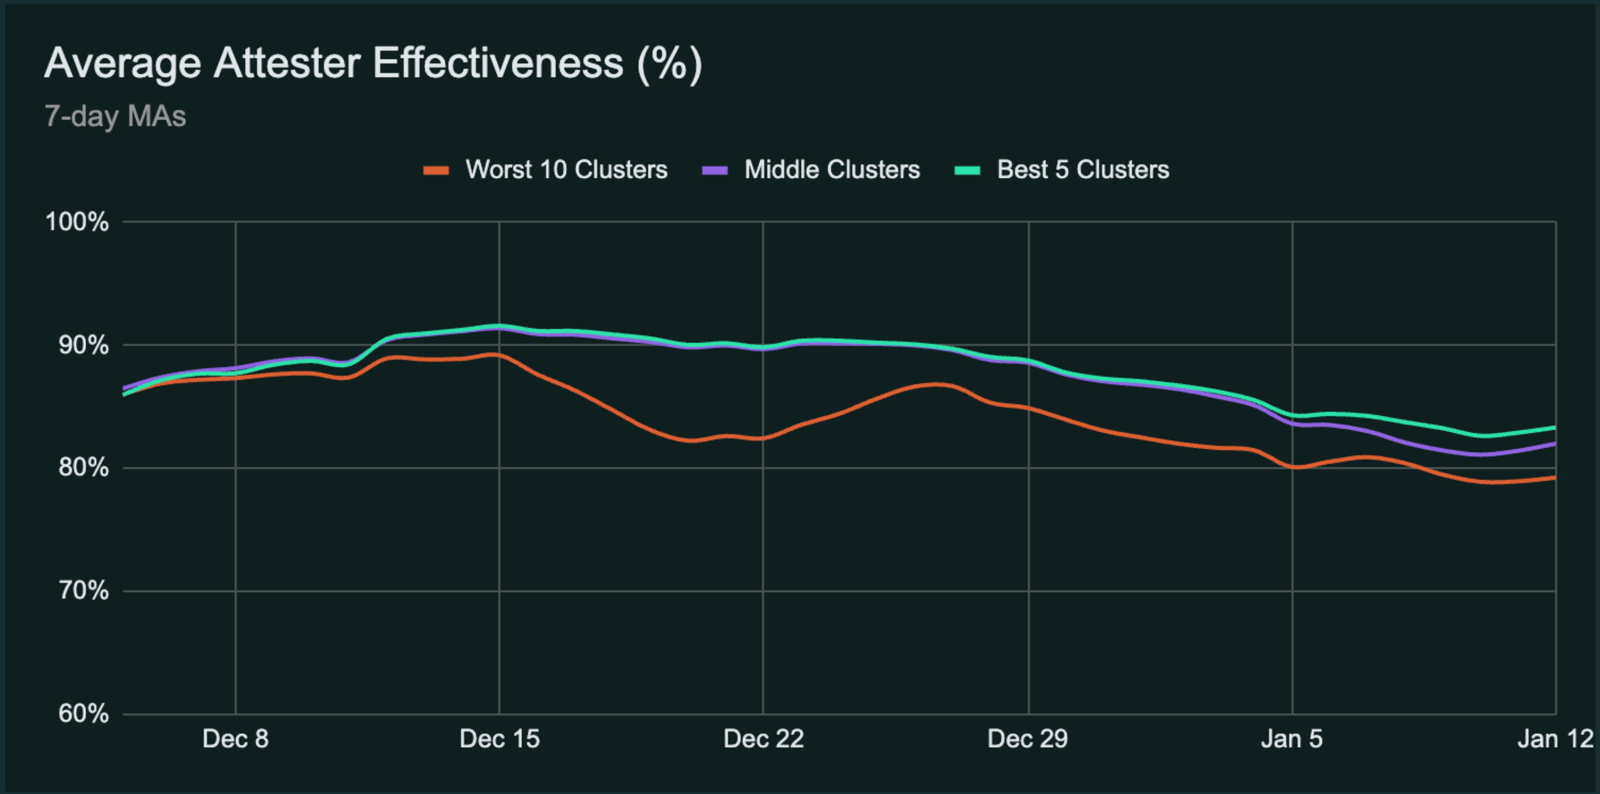

3. Latency causes and implications

Cluster latency serves as a key indicator of cluster performance. DV clusters exhibiting high latency across nodes increase the time needed to reach consensus, elevating the likelihood of missed participation in assigned duties. For instance, the action of submitting a (blinded) block for inclusion must be accomplished within a time frame of ***4 seconds.***

Illustrated in the left figure below, threshold latency (i.e. minimum latency to reach consensus within the DV cluster) exhibits significant variability across clusters. Higher cluster latency poses an increased risk of delayed duty execution and compromised performance. As long as a threshold of nodes in the cluster have low latency, the duties will be included on-chain, however, as per the left figure below, even threshold latency can vary widely based on geo and setup. The negative effect of higher latency on consensus times can be seen in the right-hand figure below.

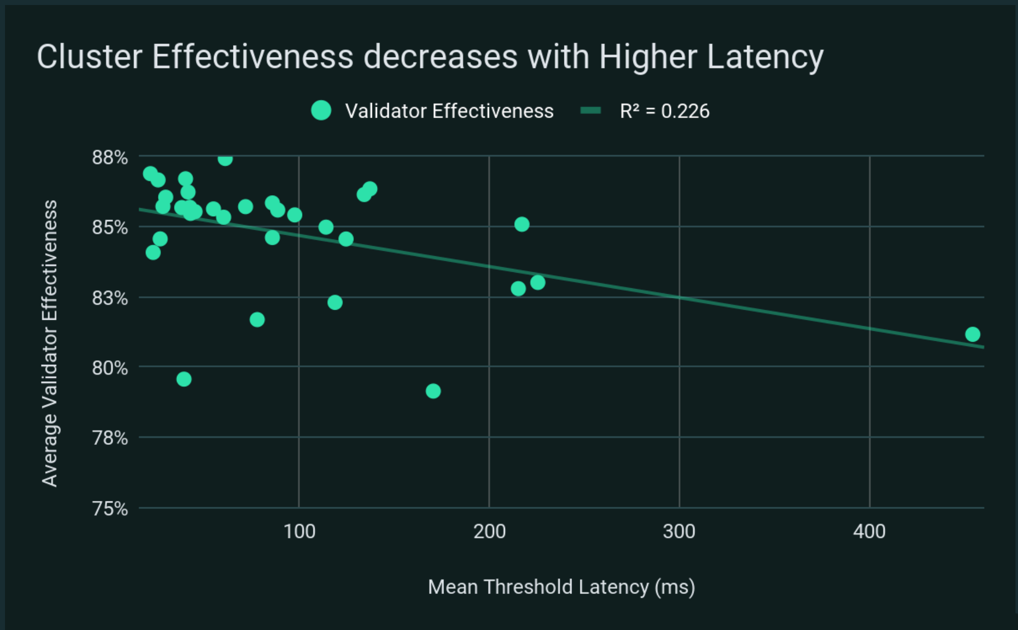

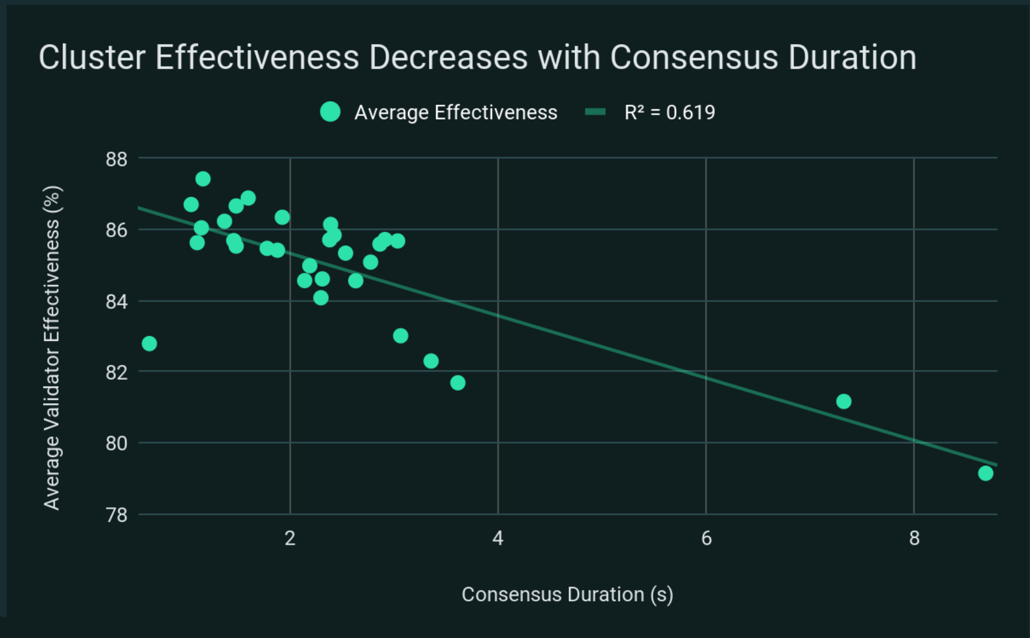

Based on the data we advise to adhere to a **latency threshold of 235 milliseconds for optimal performance**. This roughly corresponds to typical latencies between continents such as Europe and North America. Ideally, no peer should exceed that latency to any other peer, as it introduces weakness in the cluster.

One optimization introduced during the testnet was to start consensus before all peers received the header from the MEV-Boost relay. This method shortened the average blinded block duty consensus duration from 2.38 seconds to 2.16 seconds, but did not improve the duration for all clusters. More research is needed here. We will cover MEV-Boost in depth in section 4 of this report.

Nodes were split between all continents with the majority (80%) being located in Europe and a minority (12%) in Asia. While DVT allows for the inclusion of nodes in less well-connected regions within clusters, nodes can only contribute to the cluster's consensus without introducing degrading performance if they have sufficiently low latency to the other nodes. Obol is exploring other consensus mechanisms to improve this situation (e.g: HotStuff).

The map below shows the geographical distribution of nodes that participated in this testing effort. The majority of nodes were hosted in Germany (72), followed by Finland (22).

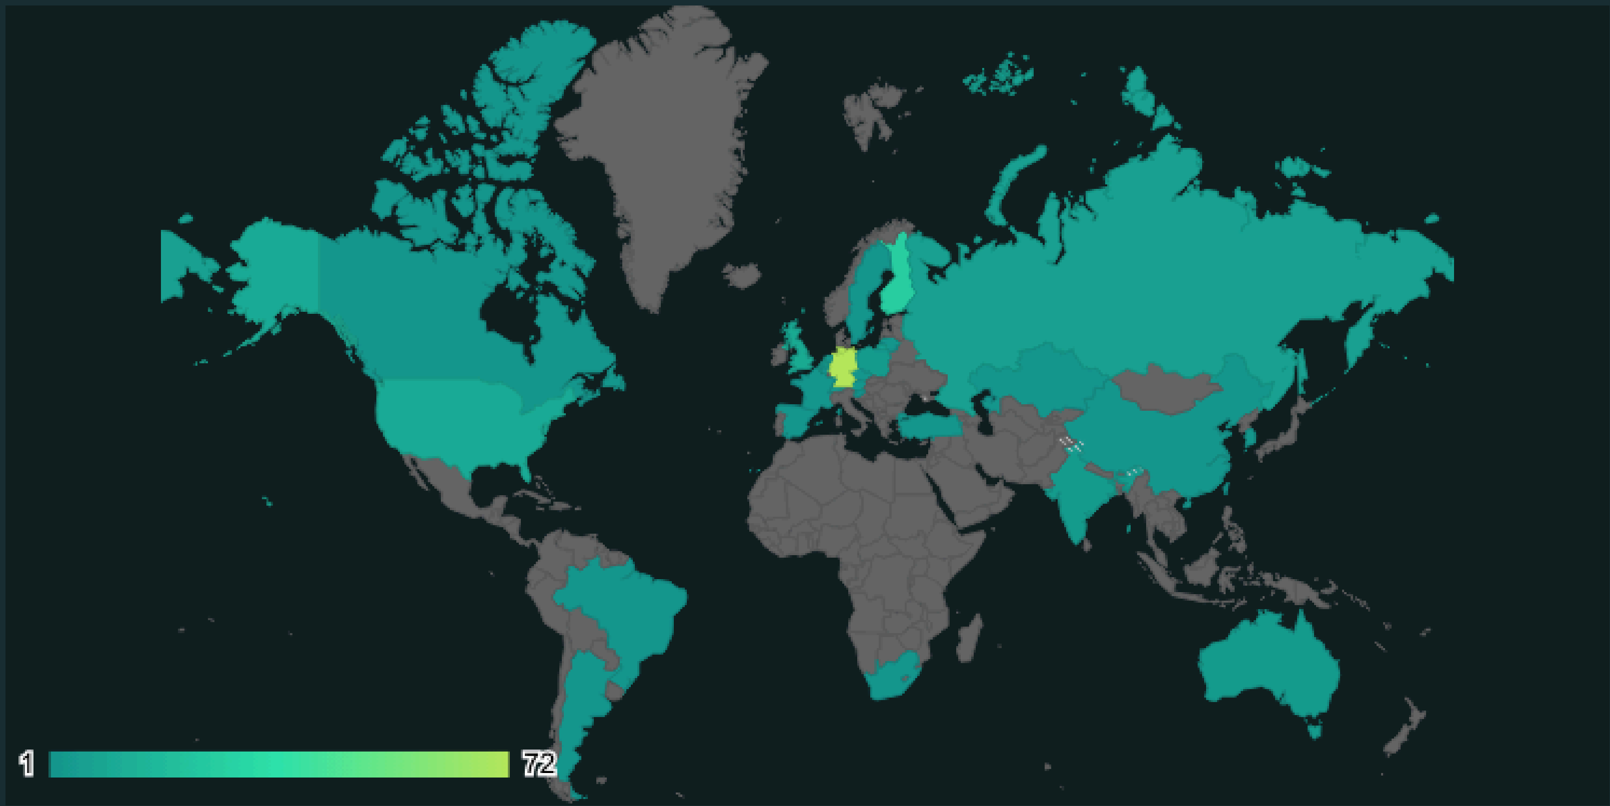

The impact of latency becomes a critical factor in geo-diverse clusters. Clusters including nodes spread across multiple continents experienced notably higher latency compared to those with all peers situated within a single region.

Even in clusters where all nodes are confined to a single continent, latency remains a significant consideration. Clusters containing one or more nodes in Asia, in particular, encountered challenges due to the region's substantial size and comparatively worse residential internet connectivity compared to Europe or North America. (Residential internet connectivity issues could also be seen in clusters with at-home validator nodes running in Africa and South America.) For example, while the distance from Lisbon to Moscow is approximately 3,900 km, the distance between Seoul and Sydney is over 8,300 km.

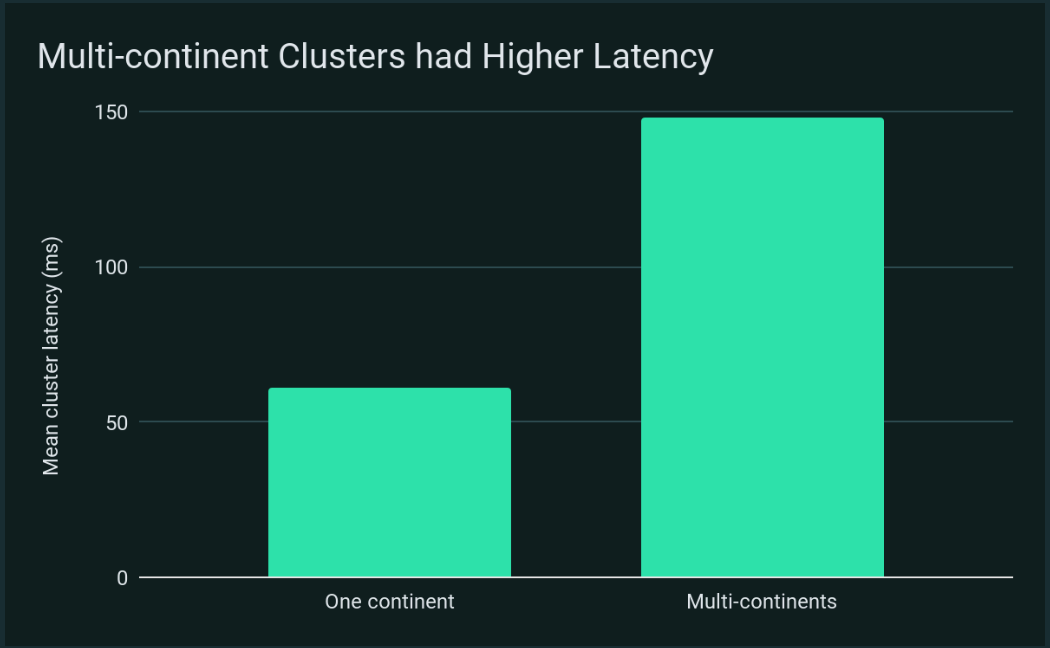

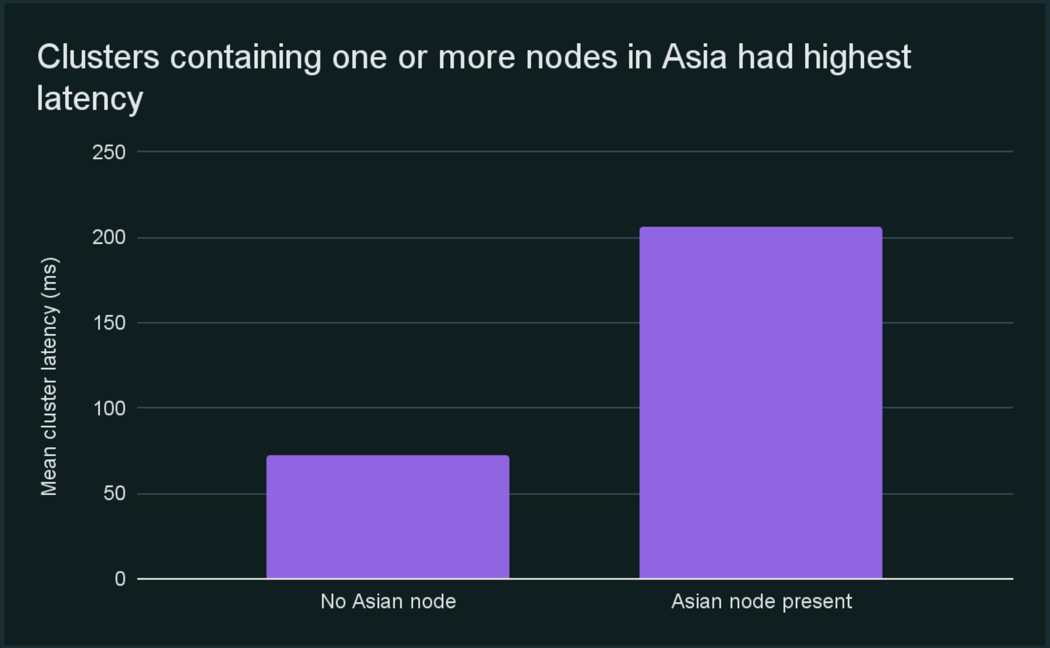

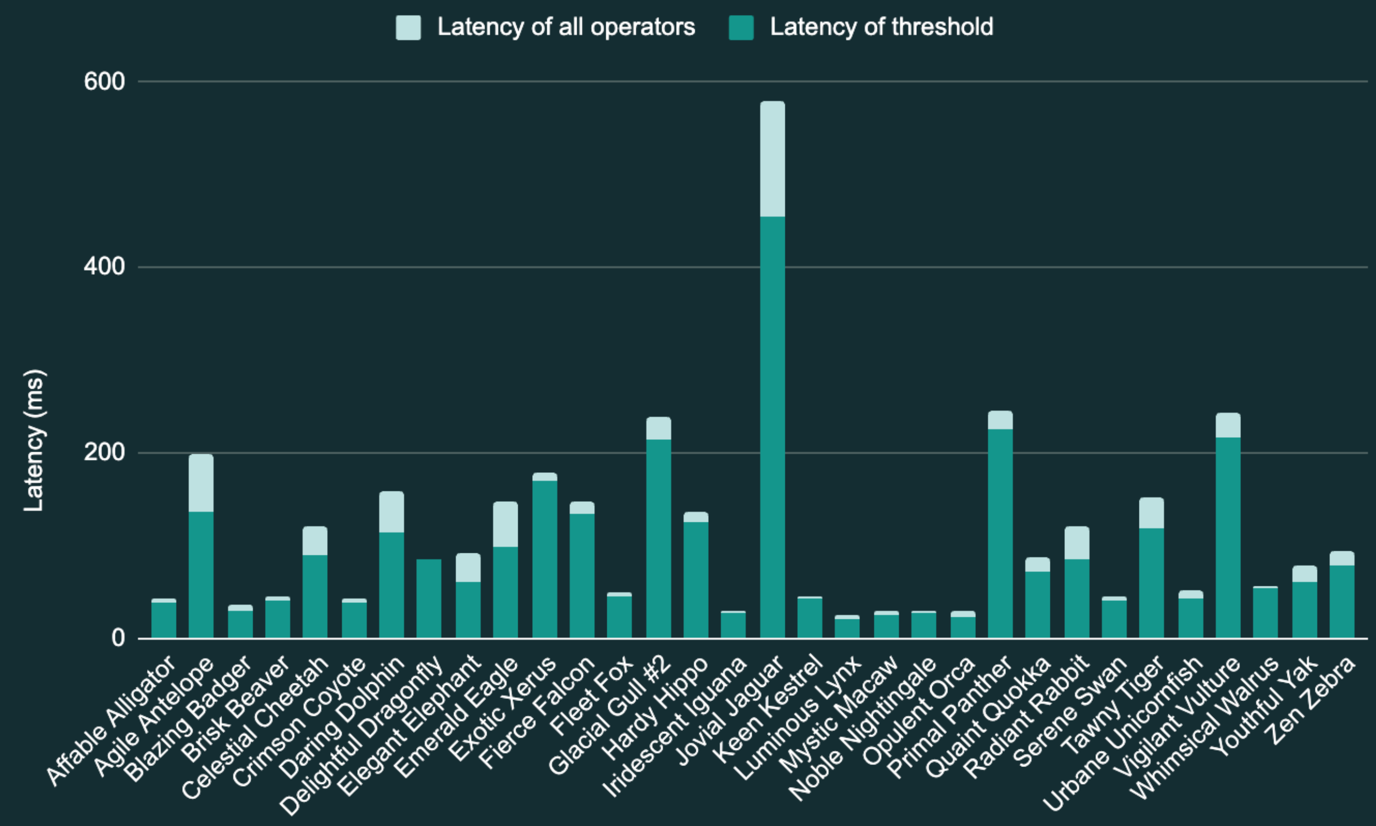

The above chart shows how the latency of the required threshold is lower, often significantly lower, than the latency of all operators including the worst node. This should warn against the inclusion of a single node in a cluster which has significantly higher latency than the others, due to geographical distance or otherwise.

4. MEV-Boost and Relays

MEV-boost offered limited support for Holesky at the beginning of the testing. Once supported, only a few relays were available, some of which engaged in overbidding and were not optimally located. This experimentation phase on the side of MEV-Boost introduced noise into the testing results, however, we were able to find some interesting insights. The relays used were: Titan relay in Frankfurt, Germany, Titan relay in Ashburn, USA, Flashbots relay in Ohio, USA, and Eden relay in the central US.

MEV-boost adds latency and can introduce significant on-chain inclusion delays in the duty process. The increase in latency is due to the requirement for every node in the cluster to fetch the block header from MEV-Boost before Charon can reach consensus. (This requires waiting up to 1 second for an answer from all the specified relays.) Then, the nodes can send the signed block back to MEV-boost. The entire process can take several seconds if nodes are far away from the relays or if relays are non-responsive. But if the process takes more than 4s, the block will not achieve chain inclusion. This can explain more than 70% of the blocks that were missed, as previously shown in the pie chart.

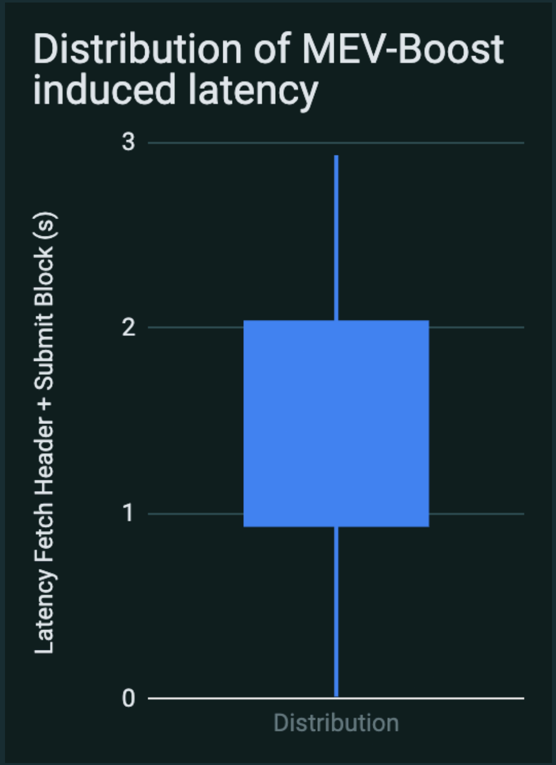

As shown in the box plot on the right, in the worst 25% of cases, MEV-Boost introduced at least 2.03 seconds of latency. In at least 75% of cases, MEV Boost introduced 0.94 seconds of latency.

The figure below shows how the increased latency of MEV-Boost resulted in lower average block proposer effectiveness. This result is likely to be more exaggerated on Holesky than Mainnet, due to the lack of well developed MEV-Boost and relay infrastructure at the time of the testing. Additionally, relays were chosen without precise knowledge of their geographic location which, as we discovered during testing, is critical to avoiding latency issues between nodes and MEV-Boost relays.

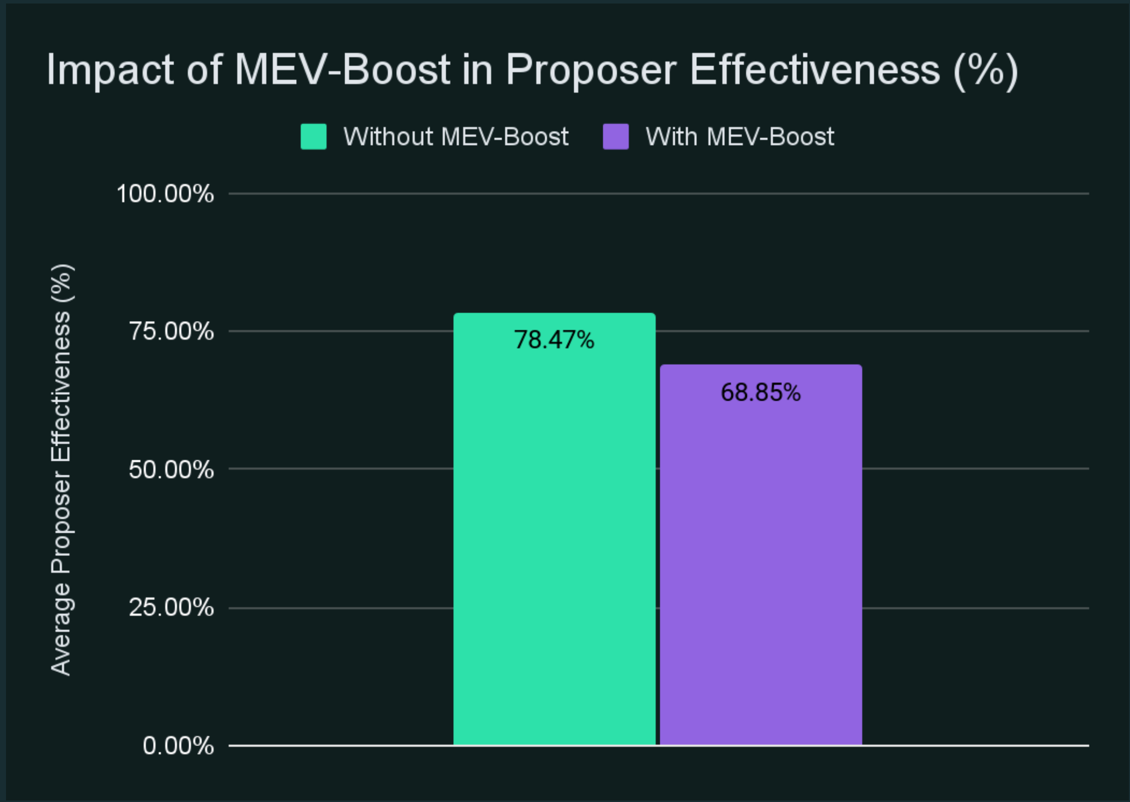

More research is needed, but one way to potentially improve latency could be to build a MEV-boost equivalent inside of Charon, and connect relays directly to the Charon DV client.

5. Hardware

Hardware is an important factor in node performance, especially on high total validator count networks such as Holesky. As shown in the bar chart, the nodes running on bare-metal were, on average, slightly more performant than nodes hosted with cloud providers. Roughly half of all operators ran in the cloud. Node operators should ensure that their nodes are properly spec'd, whether in the cloud or on bare metal, as we'll explore in the next section. More research will need to be conducted before drawing strong conclusions, and it will be important to educate operators on the most effective setups.

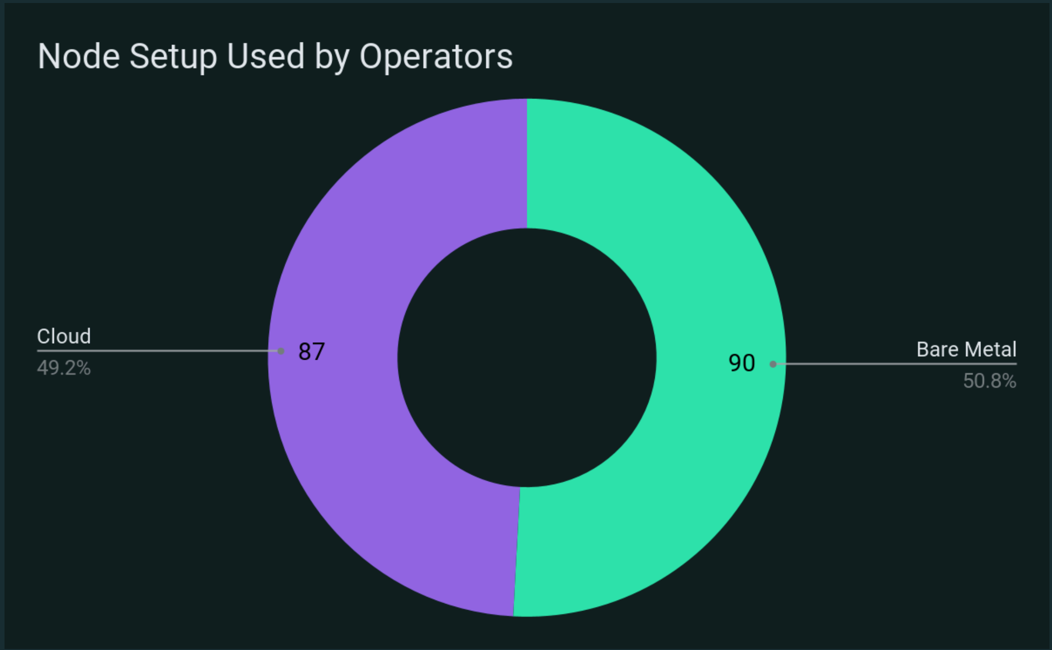

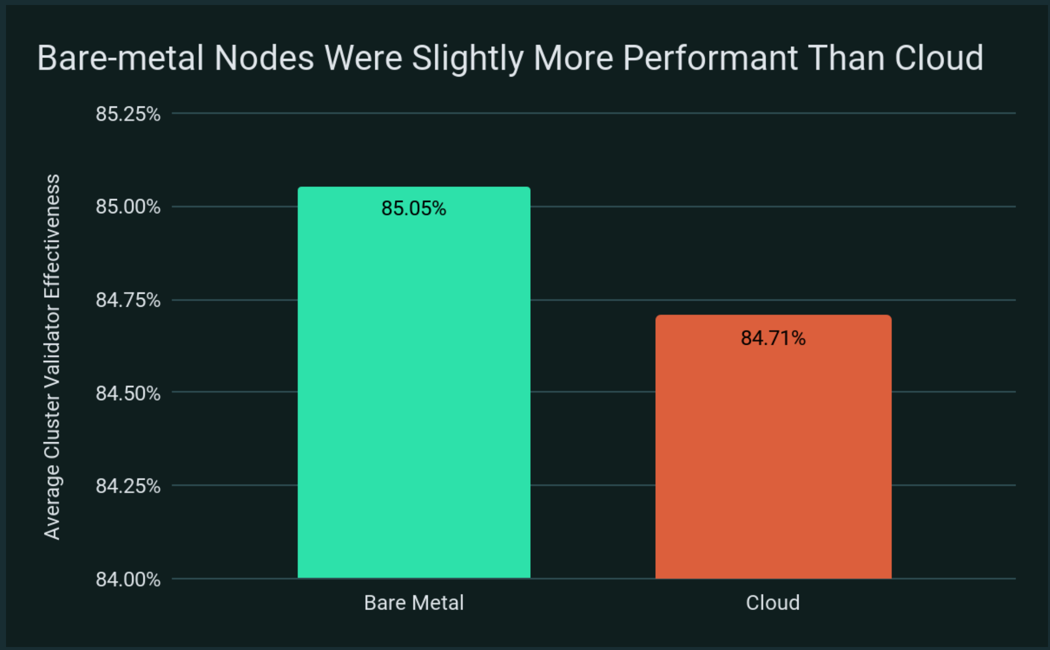

Validator effectiveness increased with increasing CPU core count, as shown in the left-hand chart below. *Note: we've removed data points for which we have less than 10 samples, as shown in the left-hand charts.*

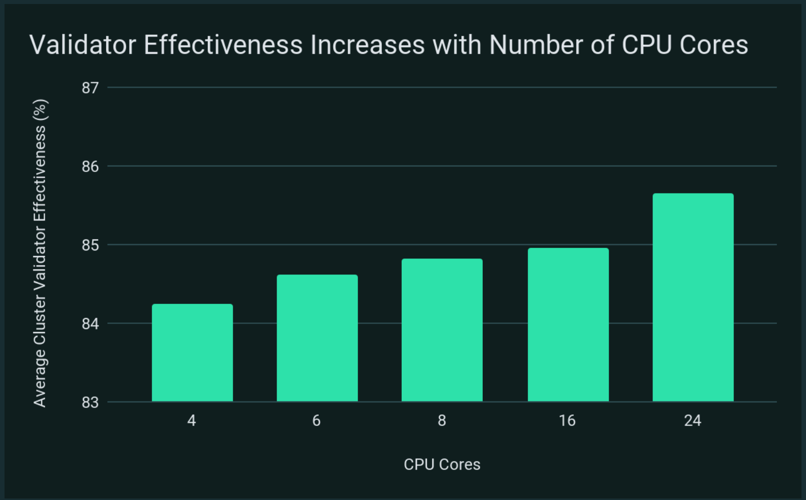

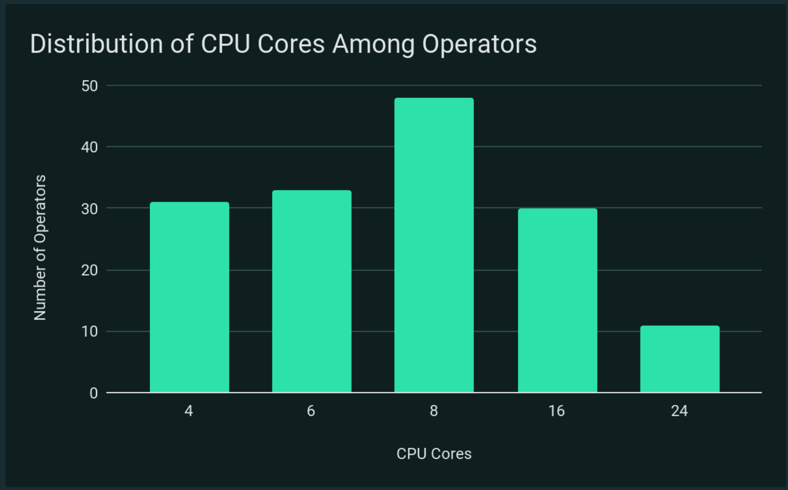

Amount of RAM did not seem to impact performance. Note: we've removed data points for which we have less than 10 samples, as shown in the left-hand chart below.

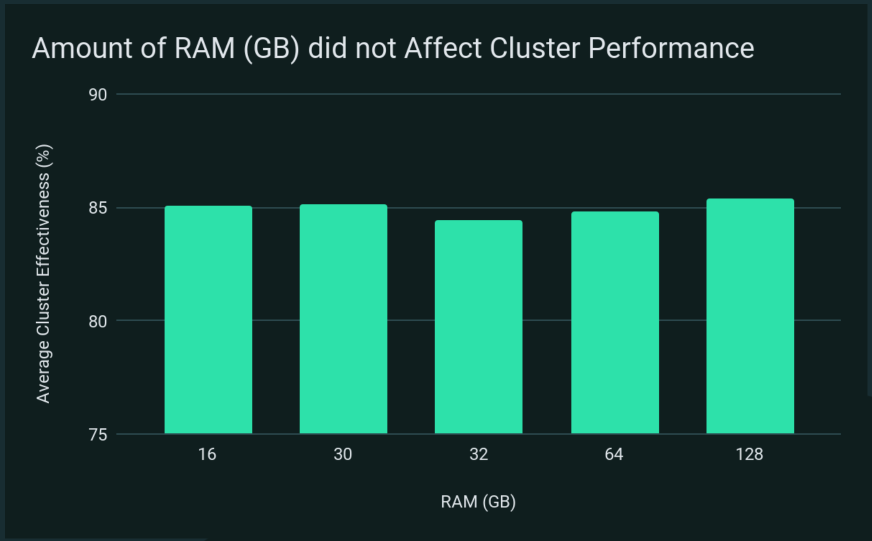

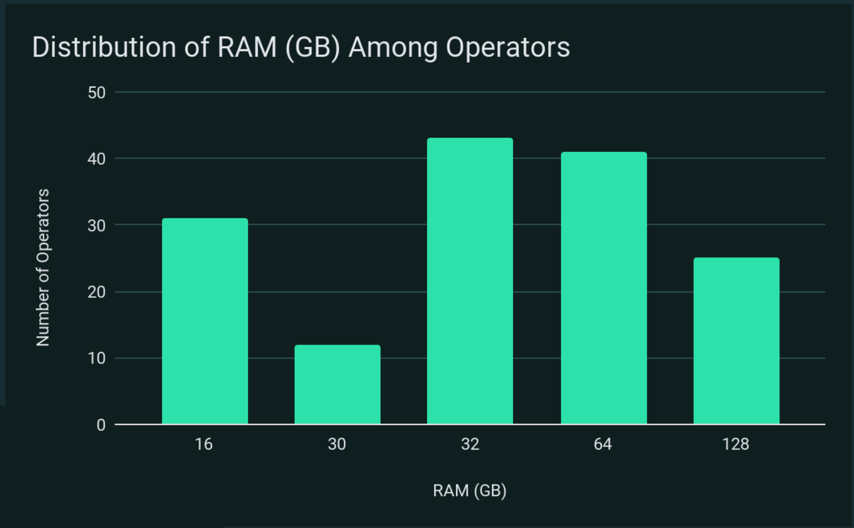

6. Notes and References

- *Data was retrieved using Rated Network's API for days 61 to 105 since genesis on Holesky.*
- *All metrics are as-is from Rated Network or deduced from simple math (average, median, etc.) definitions. See Rated docs and definitions here.*
- *Pubkeys and Raw Data can be accessed here.*