



# Quest for the Best Tests

A retrospective on #TestingTheMerge

**Parithosh Jayanthi**

DevOps Engineer, Ethereum Foundation


Section 1

# #TestingTheMerge Assemble

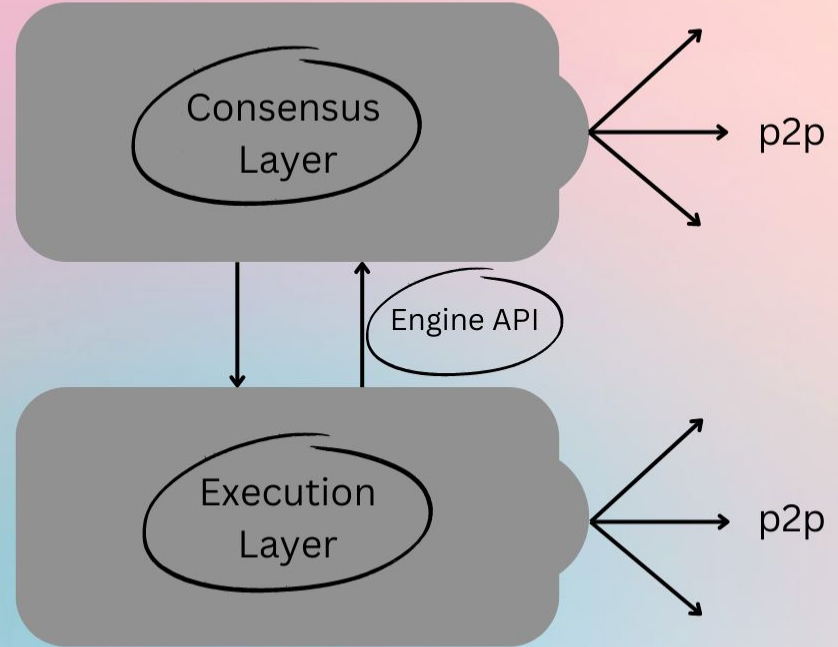
# Why is the Merge complicated?

- >20 client combinations need to be tested & regressions can sneak in very easily
- Specification is under active development -> Harder to track subtle differences
- Communicating and debugging various client combinations
- Figuring out how to test this in a reliable manner!
- All future upgrades will inherit some of the complexity - build once, use many
- Debug knowledge needed for ELs and CLs are quite different

STARECAT.COM



Forked network  
Slashing  
Unhealthy network



Reorgs  
Reduced throughput  
State corruption

# What tests can we have?

- Unit tests:
  - Handled by client teams internally
  - Usually runs on every PR
  - Reduces chance of regressions
- Integration tests:
  - Handled partially by teams
  - Involves local testnets or interop tests
  - Ensures interop at a high level

# What tests can we have?

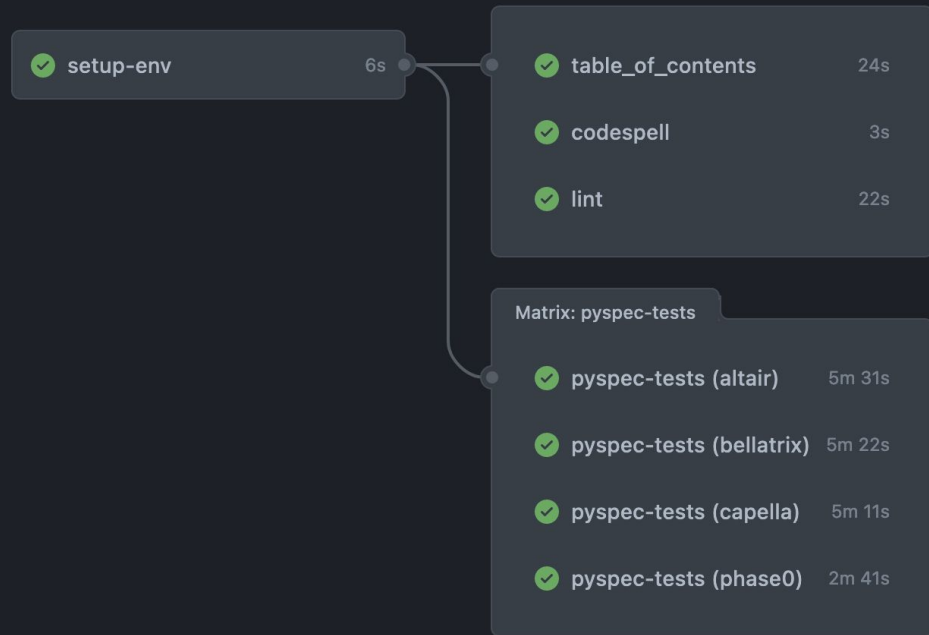
- System tests:
  - Tests end-to-end functionality
  - Involves external parties and the community
- Production tests:
  - Tests performance on a prod-like environment
  - Public testnets involving everyone
  - Finds issues that happen only at real-world loads

Section 2

# #TestingTheMerge: The Infinity War

## Spec tests

- The CI runs on every commit to the specs repo, ensuring that the specs pass tests
- Client teams import the specs and test it in their local CIs as well
- Acts as a sanity check to make sure client aren't implementing a spec that won't pass tests





# Hive tests

- Hive tests run using a simulator that starts up the clients and runs the tests against a pre-defined interface
- Acts as a integration and regression check to make sure client aren't failing defined edge cases
- e.g: Feed a Nethermind node two terminal blocks, assert how it transitions
- Shoutout to @elbuenmayini

Start time	Suite	Clients	Pass
2022-10-10T13:24:58.444Z	engine-api	nethermind	✓ (141)
2022-10-10T11:19:45.834Z	engine-api	go-ethereum	✓ (141)
2022-10-10T08:32:22.169Z	engine-api	erigon	✓ (141)
2022-10-10T04:19:55.077Z	engine-api	nethermind	✓ (141)
2022-10-10T03:25:14.198Z	engine-api	go-ethereum	✓ (141)
2022-10-09T20:19:15.930Z	engine-api	nethermind	✓ (141)
2022-10-09T19:24:28.663Z	engine-api	go-ethereum	✓ (141)
2022-10-09T16:35:01.964Z	engine-api	erigon	✓ (141)
2022-10-09T11:41:45.059Z	engine-api	nethermind	✓ (141)
2022-10-09T10:46:59.564Z	engine-api	go-ethereum	✓ (141)
2022-10-09T07:59:43.300Z	engine-api	erigon	✓ (141)
2022-10-09T03:05:33.754Z	engine-api	nethermind	✓ (141)
2022-10-09T02:10:57.615Z	engine-api	go-ethereum	✓ (141)
2022-10-08T18:25:21.055Z	engine-api	nethermind	✓ (141)
2022-10-08T17:30:35.247Z	engine-api	go-ethereum	✓ (141)
2022-10-10T10:24:07.968Z	eth2-testnet	teku-vc,go-ethereum,teku-bn	✓ (2)

# Kurtosis tests

- Kurtosis spins up a local testnet with the required EL/CL combinations and then allows them to transition/merge. It then asserts some “happy case” conditions.
- An integration test make sure client are compatible
- Useful to rapidly iterate ideas
- e.g: Are blocks being produced, are there tx's...

The screenshot shows a CI pipeline for `cl-nethermind.yml` on the `workflow_dispatch` event. A job named `apt` has completed successfully in 11 seconds. A matrix job named `kurtosis` is shown with five entries, each representing a different client: `lighthouse`, `lodestar`, `nimbus`, `prysm`, and `teku`. All entries in the matrix are marked as successful with a green checkmark.

Client	Duration
kurtosis (lighthouse)	42m 29s
kurtosis (lodestar)	42m 6s
kurtosis (nimbus)	38m 15s
kurtosis (prysm)	38m 30s
kurtosis (teku)	39m 28s

**Artifacts**  
Produced during runtime

Name	Size
<code>lighthouse-nethermind</code> <span>Expired</span>	33.7 MB
<code>lodestar-nethermind</code> <span>Expired</span>	17.5 MB
<code>nimbus-nethermind</code> <span>Expired</span>	38.2 MB
<code>prysm-nethermind</code> <span>Expired</span>	34 MB
<code>teku-nethermind</code> <span>Expired</span>	714 MB

## Sync tests

- The sync test co-ordinator spins up every client combination daily and syncs to head on various testnets. Both genesis sync as well as Checkpoint sync are performed.
- Edge case sync tests are also performed: EL down, CL down, etc
- Acts as a integration test make sure users can always sync the network
- Shoutout to @samcmAU

```
test:
  name: "basic"

  tasks:
  - name: run_command
    config:
      command:
      - "echo"
      - "hello!"
  - name: execution_is_healthy
  - name: consensus_is_healthy
  - name: both_are_synced
    config:
      consensus:
        percent: 100
      execution:
        percent: 100
  - name: run_command
    config:
      command:
      - "echo"
      - "done!"

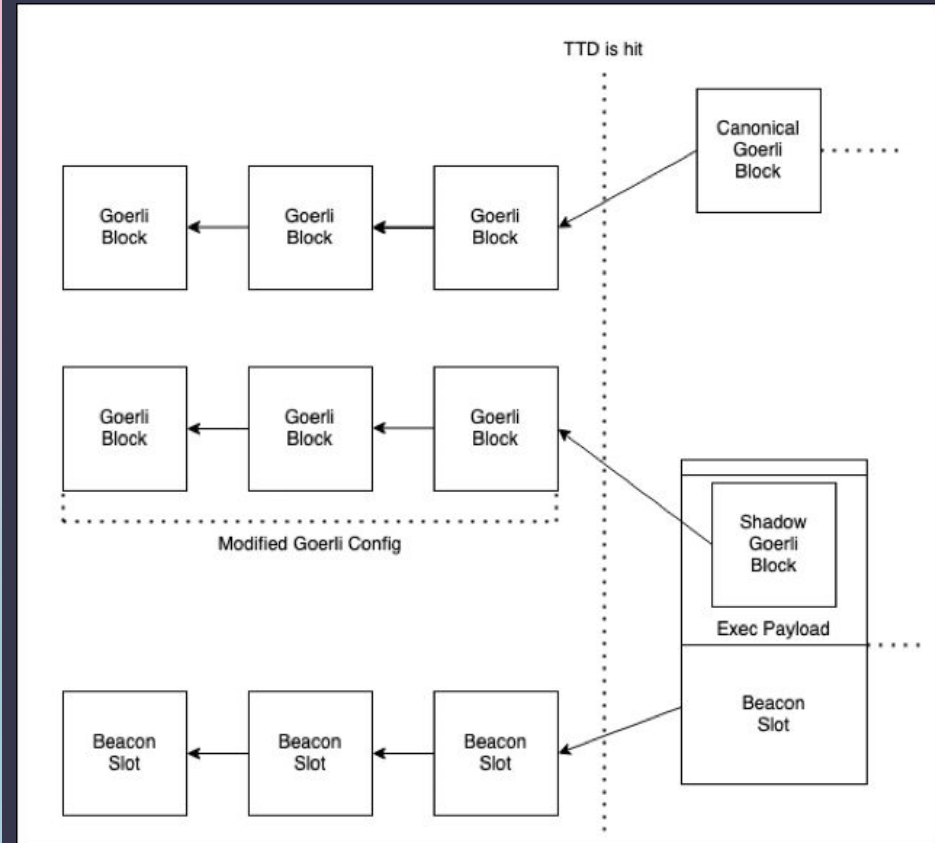
  execution:
    url: http://localhost:8545

  consensus:
    url: http://localhost:5052
```

- ✓ run-test (lighthouse, geth, ropsten, i...
- ✓ run-test (lighthouse, besu, ropsten, ...
- ✓ run-test (lighthouse, nethermind, ro...
- ✓ run-test (lighthouse, erigon, ropsten...
- ✓ run-test (teku, geth, ropsten, is-heal...
- ✓ run-test (teku, besu, ropsten, is-hea...
- ✓ run-test (teku, nethermind, ropsten,...
- ✓ run-test (teku, erigon, ropsten, is-h...
- ✓ run-test (prysm, geth, ropsten, is-h...
- ✓ run-test (prysm, besu, ropsten, is-h...
- ✓ run-test (prysm, nethermind, ropste...
- ✓ run-test (prysm, erigon, ropsten, is-...
- ✓ run-test (nimbus, geth, ropsten, is-h...
- ✓ run-test (nimbus, besu, ropsten, is-...
- ✓ run-test (nimbus, nethermind, ropst...
- ✓ run-test (nimbus, erigon, ropsten, is...
- ✓ run-test (lodestar, geth, ropsten, is-...
- ✓ run-test (lodestar, besu, ropsten, is-...
- ✓ run-test (lodestar, nethermind, rop...
- ✓ run-test (lodestar, erigon, ropsten, i...

# Shadow Forks & Testnets

- Allows us to check compatibility across all clients through the entire lifecycle
- Fresh testnets allow us to check assumptions across client pairs without much overhead
- Shadow forks allow us to stress test the clients with real state and transaction load
- We can invite participants in a controller manner to take part in the tests
- Acts as release test which triggers real world edge cases, before we recommend the releases to the general public



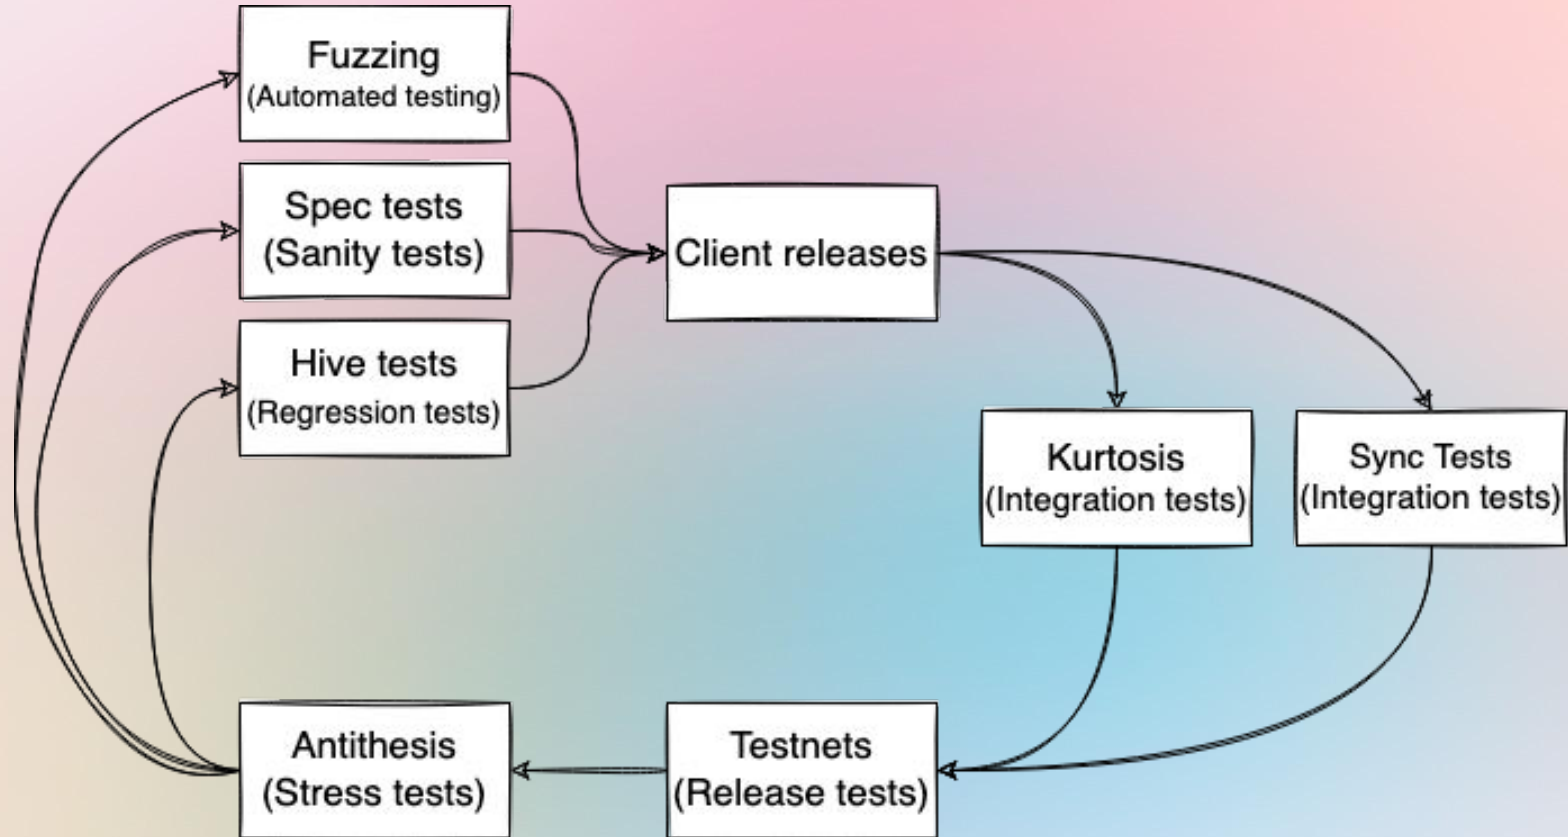
# Antithesis & Fuzzers

- Antithesis offers a deterministic hypervisor which allows us to perform network splits, packet loss while fuzzing clients. The deterministic hypervisor allows us to re-trigger the issue, allowing for capturing the state of the client and easier debugging.
- Various fuzzers are run against different layers of the stack to find bugs.
- These bugs also allow us to re-evaluate if changes need to be made in the specs or if the bug is an implementation level issue.

The screenshot shows a debugger window with three memory locations. Each location displays a list of memory addresses and their corresponding hex values. The first location starts at 01110804 and ends at 01110844. The second location starts at 01110848 and ends at 01110888. The third location starts at 01110890 and ends at 01110930. The hex values are displayed in a grid-like format, with some values appearing to be part of a larger data structure or array.

Address	Value
01110804	49 01110820 9F 01110778 45E11175.1
01110808	28 01110824 29 01110808 00000000.0
0111080C	55 01110808 9F 0111079C 44011155.6
01110810	33 0111079C 44 0001 0000 0000 0000.0
01110814	54 0111079C 30 01110798 40011172.0
01110818	43 01110798 40 0001 0000 0000 0000.0
0111081C	56 01110798 40 0001 0000 0000 0000.0
01110820	33 0111079C 40 0001 0000 0000 0000.0
01110824	59 0111079C 40 0001 0000 0000 0000.0
01110828	48 0111079C 30 0001 0000 0000 0000.0
0111082C	62 0111079C 40 0001 0000 0000 0000.0
01110830	45 0111079C 40 0001 0000 0000 0000.0
01110834	61 0111079C 40 0001 0000 0000 0000.0
01110838	60 0111079C 40 0001 0000 0000 0000.0
0111083C	69 0111079C 30 0001 0000 0000 0000.0
01110840	68 0111079C 40 0001 0000 0000 0000.0
01110844	70 0111079C 40 0001 0000 0000 0000.0
01110848	71 0111079C 40 0001 0000 0000 0000.0
0111084C	72 0111079C 40 0001 0000 0000 0000.0
01110850	73 0111079C 40 0001 0000 0000 0000.0
01110854	74 0111079C 40 0001 0000 0000 0000.0
01110858	75 0111079C 40 0001 0000 0000 0000.0
0111085C	76 0111079C 40 0001 0000 0000 0000.0
01110860	77 0111079C 40 0001 0000 0000 0000.0
01110864	78 0111079C 40 0001 0000 0000 0000.0
01110868	79 0111079C 40 0001 0000 0000 0000.0
0111086C	80 0111079C 40 0001 0000 0000 0000.0
01110870	81 0111079C 40 0001 0000 0000 0000.0
01110874	82 0111079C 40 0001 0000 0000 0000.0
01110878	83 0111079C 40 0001 0000 0000 0000.0
0111087C	84 0111079C 40 0001 0000 0000 0000.0
01110880	85 0111079C 40 0001 0000 0000 0000.0
01110884	86 0111079C 40 0001 0000 0000 0000.0
01110888	87 0111079C 40 0001 0000 0000 0000.0
0111088C	88 0111079C 40 0001 0000 0000 0000.0
01110890	89 0111079C 40 0001 0000 0000 0000.0
01110894	90 0111079C 40 0001 0000 0000 0000.0
01110898	91 0111079C 40 0001 0000 0000 0000.0
0111089C	92 0111079C 40 0001 0000 0000 0000.0
011108A0	93 0111079C 40 0001 0000 0000 0000.0
011108A4	94 0111079C 40 0001 0000 0000 0000.0
011108A8	95 0111079C 40 0001 0000 0000 0000.0
011108AC	96 0111079C 40 0001 0000 0000 0000.0
011108B0	97 0111079C 40 0001 0000 0000 0000.0
011108B4	98 0111079C 40 0001 0000 0000 0000.0
011108B8	99 0111079C 40 0001 0000 0000 0000.0
011108BC	00 0111079C 40 0001 0000 0000 0000.0
011108C0	01 0111079C 40 0001 0000 0000 0000.0
011108C4	02 0111079C 40 0001 0000 0000 0000.0
011108C8	03 0111079C 40 0001 0000 0000 0000.0
011108CC	04 0111079C 40 0001 0000 0000 0000.0
011108D0	05 0111079C 40 0001 0000 0000 0000.0
011108D4	06 0111079C 40 0001 0000 0000 0000.0
011108D8	07 0111079C 40 0001 0000 0000 0000.0
011108DC	08 0111079C 40 0001 0000 0000 0000.0
011108E0	09 0111079C 40 0001 0000 0000 0000.0
011108E4	0A 0111079C 40 0001 0000 0000 0000.0
011108E8	0B 0111079C 40 0001 0000 0000 0000.0
011108EC	0C 0111079C 40 0001 0000 0000 0000.0
011108F0	0D 0111079C 40 0001 0000 0000 0000.0
011108F4	0E 0111079C 40 0001 0000 0000 0000.0
011108F8	0F 0111079C 40 0001 0000 0000 0000.0
01110900	10 0111079C 40 0001 0000 0000 0000.0
01110904	11 0111079C 40 0001 0000 0000 0000.0
01110908	12 0111079C 40 0001 0000 0000 0000.0
0111090C	13 0111079C 40 0001 0000 0000 0000.0
01110910	14 0111079C 40 0001 0000 0000 0000.0
01110914	15 0111079C 40 0001 0000 0000 0000.0
01110918	16 0111079C 40 0001 0000 0000 0000.0
0111091C	17 0111079C 40 0001 0000 0000 0000.0
01110920	18 0111079C 40 0001 0000 0000 0000.0
01110924	19 0111079C 40 0001 0000 0000 0000.0
01110928	1A 0111079C 40 0001 0000 0000 0000.0
0111092C	1B 0111079C 40 0001 0000 0000 0000.0
01110930	1C 0111079C 40 0001 0000 0000 0000.0

# Testing lifecycle for the Merge

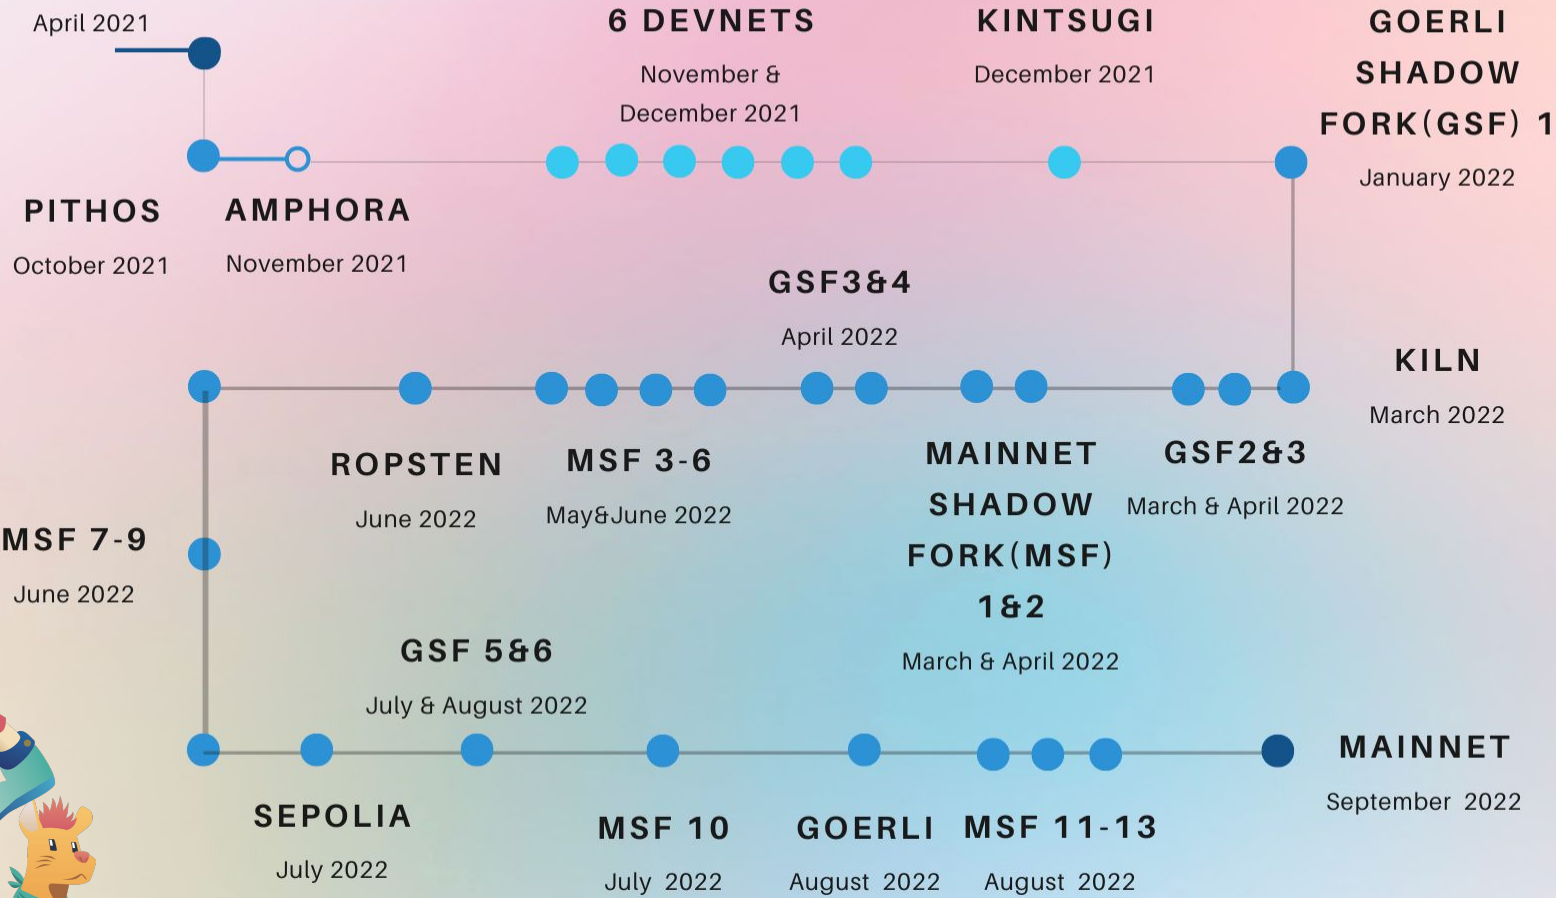


Section 3

# #TestingTheMerge: The Endgame



# RAYONISM





## So what did we still miss?

- In-memory database too low to process mainnet blocks
- Non-optimal block production: Random production of 0/few tx blocks
- Multiple terminal blocks (in a specific condition) caused missing receipts and caused failed proposals
- Lots of constant syncing nodes on mainnet led to unexpected performance degradation when compared to shadow forks
- Failover beacon node scenario -> some requests sent just to the primary

# What can we reuse?

Running testnets helped show us tooling blind spots in the DevOps ecosystem:

- Metrics exporter: <https://github.com/samcm/ethereum-metrics-exporter>
- Sync testing: <https://github.com/samcm/ethereum-sync-testing/actions>
- Genesis gen.: <https://github.com/skylenet/ethereum-genesis-generator>
- Client automation: <https://github.com/ethPandaOps/ethereum-helm-charts>
- Scalable testnets: <https://github.com/ethPandaOps/ethereum-k8s-testnets>
- Easy testnets: <https://github.com/ethereum/consensus-deployment-ansible>
- Faucet: <https://github.com/komputing/FaucETH>
- Checkpoint Sync Provider: <https://github.com/samcm/checkpointz>
- PRs to explorers, validator key generation tools, load balancer

**If you want to join the testing efforts  
contact [mario.vega@ethereum.org](mailto:mario.vega@ethereum.org)**



**Thank you!**  
**Join #TestingThe{Surge,Verge,Purge}!**

**Parithosh Jayanthi**

DevOps Engineer

parithosh@ethereum.org



@parithosh\_j