

Advanced Database Systems

Final Presentation

Golatac

May 2, 2024

Yen-Ju Wu, Chien-Yu Liu, Zilong Zhou

Project Goals

- **75% → Complete the catalog implementation following rest-catalog-open-api.yaml**
- **100% → Unit test & Performance evaluation**
- 125% → Performance optimization

DONE

DONE

DONE?

Unit Testing - Namespace

- Namespace
 - list_namespace
 - non_exist (404) / exist (200)
 - create_namespace
 - success (200)
 - get_namespace
 - non_exist (404) / exist (200)
 - check_namespace
 - non_exist (404) / exist (204)
 - delete_namespace
 - non_exist (404) / exist (204)

Unit Testing - Table

- Table
 - get_table_by_namespace
 - empty_result (404) / result_found (200)
 - post_table_by_namespace
 - new_table (200) / conflict (409)
 - delete_table
 - table_exists (204) / table_not_exists (404)
 - head_table
 - table_exists (204) / table_not_exists (404)
 - rename_table
 - (204)

Unit Testing

```
running 18 tests
test server::routes::namespace::test::test_list_namespace ... ok
test server::routes::namespace::test::test_create_namespace ... ok
test server::routes::namespace::test::test_check_namespace_non_exist ... ok
test server::routes::namespace::test::test_delete_non_exist_namespace ... ok
test server::routes::namespace::test::test_check_namespace_exist ... ok
test server::routes::namespace::test::test_get_non_exist_namespace ... ok
test server::routes::namespace::test::test_delete_exist_namespace ... ok
test server::routes::namespace::test::test_get_namespace ... ok
test server::routes::table::test::test_get_table_by_namespace_empty_result ... ok
test server::routes::table::test::test_delete_table_that_not_exists ... ok
test server::routes::table::test::test_delete_table_that_exists ... ok
test server::routes::table::test::test_post_table_by_namespace_conflict ... ok
test server::routes::table::test::test_post_table_by_namespace_new_table ... ok
test server::routes::namespace::test::test_list_non_exist_namespace ... ok
test server::routes::table::test::test_head_table ... ok
test server::routes::table::test::test_get_table_by_namespace_result_found ... ok
test server::routes::table::test::test_rename_table ... ok
test test::test_create_server ... ok

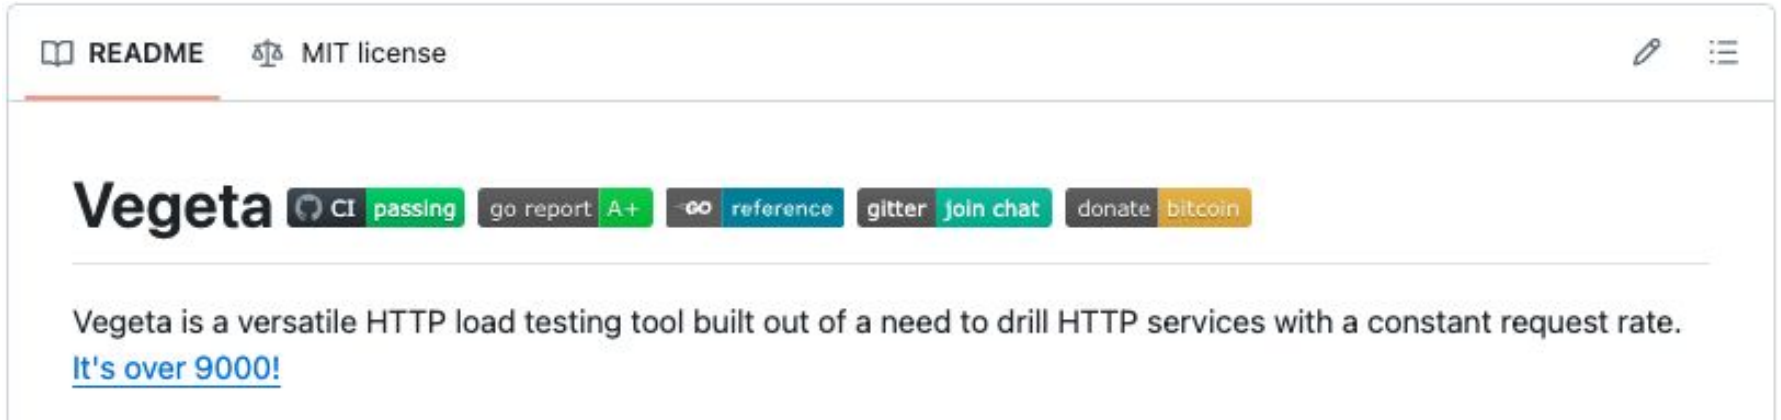
test result: ok. 18 passed; 0 failed; 0 ignored; 0 measured; 0 filtered out; finished in 0.04s
```

Assessment

- Pros
 - Modularity and Encapsulation
 - Serve http requests strictly following rest-catalog-open-api.yaml
 - Robust Error Handling based on Iceberg
- Need more work
 - Higher unit test coverage & integration tests
 - Performance optimizations

Benchmark Results

- Key performance metrics: **Latency**
- Framework: [Vegeta](#), which is a HTTP load testing tool



The screenshot shows the top portion of a GitHub README page for the 'Vegeta' project. At the top left, there are links for 'README' and 'MIT license'. Below this, the word 'Vegeta' is displayed in a large, bold font. To its right is a row of utility buttons: 'CI passing' (with a green checkmark), 'go report A+' (with a green checkmark), 'reference' (with a blue checkmark), 'gitter join chat' (with a green checkmark), and 'donate bitcoin' (with a yellow checkmark). Below the buttons, a paragraph of text reads: 'Vegeta is a versatile HTTP load testing tool built out of a need to drill HTTP services with a constant request rate. [It's over 9000!](#)'

Benchmark Design

- Single endpoint stress tests
 - `get_table`
 - `list_table`
 - `get_namespace`
 - `list_namespace`
- Multiple random endpoints stress test
 - `get_randomized_table`

Benchmark Settings

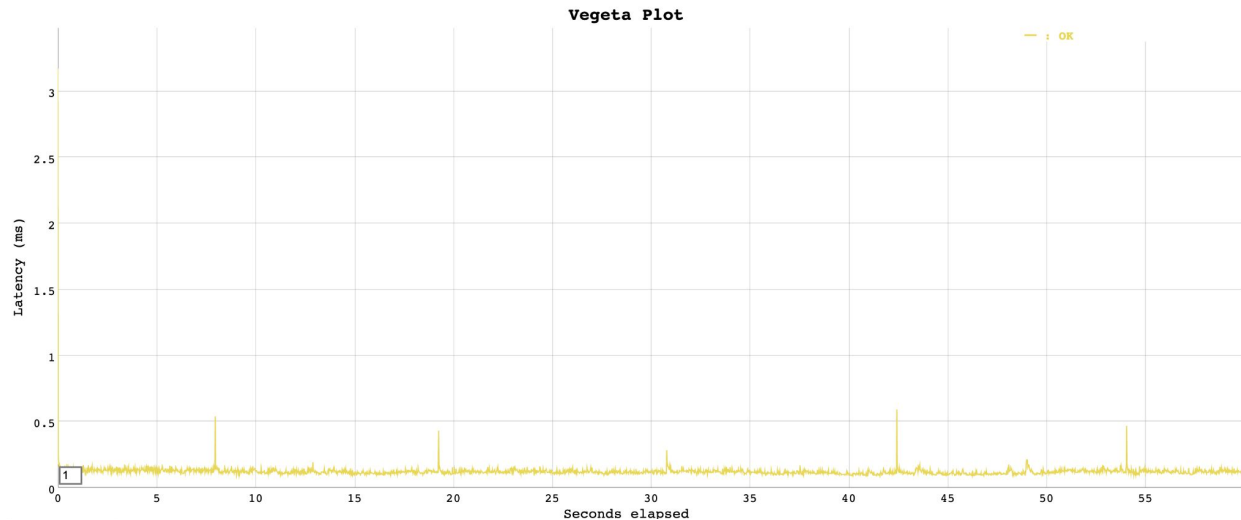
- AWS EC2 Instance
 - C5.2xlarge
 - Image: ubuntu
 - 8 vCPU
 - 16 GiB Memory
 - 20 GiB Disk
- Parameters
 - Empty Metadata
 - Number of namespaces: 100
 - Number of tables: 10000
 - Request rate: 50 per second
 - Duration: 60 seconds



Single endpoint – get_table

Latencies

- min: 91.112 μ s
- mean: 122.644 μ s
- 50: 118.992 μ s
- 90: 139.968 μ s
- 95: 147.852 μ s
- 99: 168.578 μ s
- max: 3.172ms

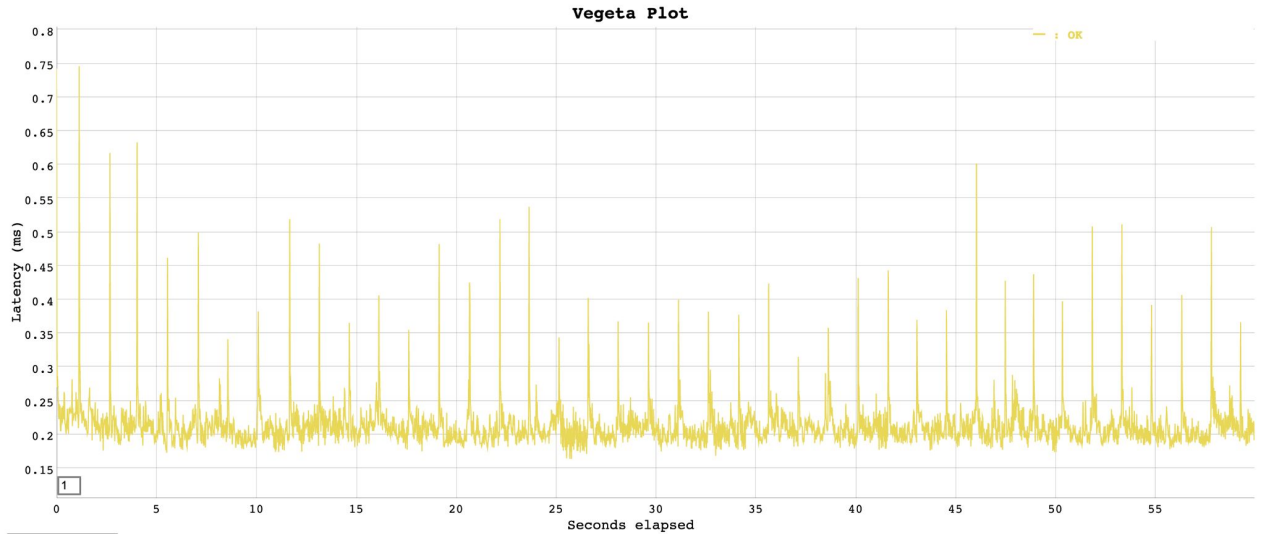


Throughput (RPS): 13281.55

Single endpoint – list_table

Latencies

- min: 163.798 μ s
- mean: 213.732 μ s
- 50: 208.01 μ s
- 90: 237.749 μ s
- 95: 253.864 μ s
- 99: 382.017 μ s
- max: 746.229 μ s




Throughput (RPS): 12824.26

Single endpoint – get_namespace

Latencies

- min: 100.617 μ s
- mean: 139.855 μ s
- 50: 136.852 μ s
- 90: 158.902 μ s
- 95: 170.513 μ s
- 99: 200.55 μ s
- max: 685.701 μ s




Throughput (RPS): 7470.60

Single endpoint – list_namespace

Latencies

- min: 86.401 μ s
- mean: 116.493 μ s
- 50: 113.343 μ s
- 90: 132.67 μ s
- 95: 140.342 μ s
- 99: 169.723 μ s
- max: 590.968 μ s

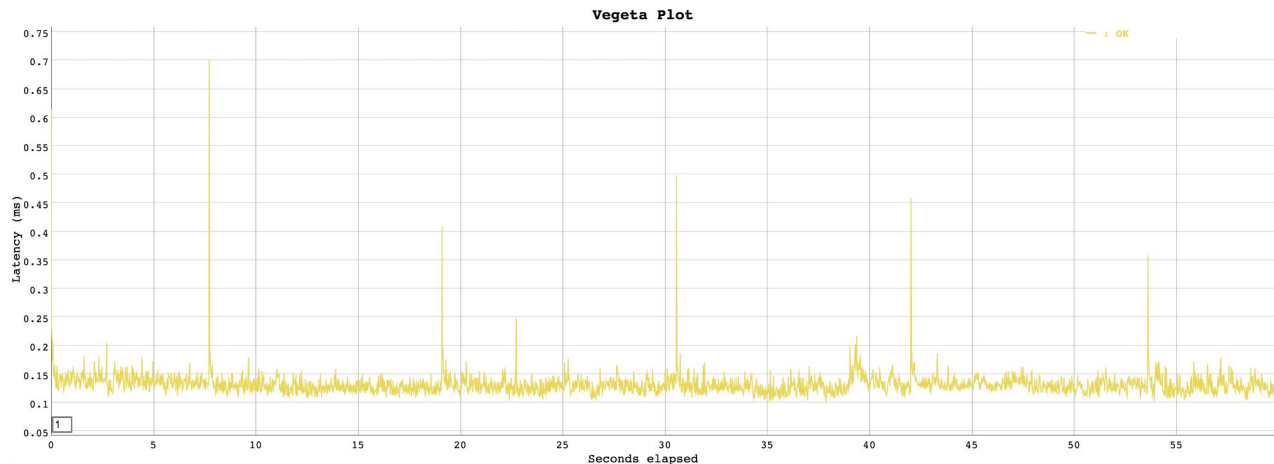


Throughput (RPS): 13641.37

Single endpoint – get_randomized_table

Latencies

- min: 102.25 μ s
- mean: 132.332 μ s
- 50: 130.002 μ s
- 90: 147.739 μ s
- 95: 155.26 μ s
- 99: 176.419 μ s
- max: 699.689 μ s



Throughput (RPS): 12966.11

Benchmark Insights

- P99
 - < 1ms, lower than network latency
- Mean Latency
 - The "List table's result" benchmark test has the highest mean latency at approximately 213.732 μ s
 - The "Get table's result" benchmark test has the lowest mean latency at approximately 122.644 μ s
- Latency deviation
 - In all stress tests, the 95th percentile latency remains within a 20% deviation from the mean latency

Future Work

- Performance optimizations
- Support of optional fields in schema according to other teams' needs
- Documentation / Comments
- Consistent coding style

Thank You