



Workshop on Load Testing & Benchmarking  
LTB 2024  
London | May 7, 2024



**LTB 2024**  
Load Testing &  
Benchmarking

# Performance Testing Transformation

Alexander Podelko  
Sr. Performance Engineer  
Amazon Web Services

© 2024, Amazon Web Services, Inc. or its affiliates.

1

## PERFORMANCE TESTING TRANSFORMATION

### Alex Podelko

- Has specialized in performance since 1997
- Senior Performance Engineer at AWS – Amazon Aurora
  - Before worked for MongoDB, Oracle/Hyperion, Intel, and Aetna
- SPEC RG Steering Committee Member



Disclaimer: The views expressed here are my personal views only and do not necessarily represent those of my current or previous employers. All brands and trademarks mentioned are the property of their owners. All products are mentioned as examples only, not as recommendations.



© 2024, Amazon Web Services, Inc. or its affiliates.

2

# Adjusting Performance Testing to Industry Trends:

Adding early and continuous performance testing



© 2024, Amazon Web Services, Inc. or its affiliates.

3

3

## PERFORMANCE TESTING TRANSFORMATION

### Industry Trends

- Web
  - Centralization, open / unlimited workload
- Cloud
  - Further centralization, price tag (FinOps)
  - Dynamic configurations / Self-Management
- Agile / iterative development
  - Continuous Integration / Delivery / Deployment
  - DevOps / SRE

[The Past, Present, and Future of Performance Engineering](#)



© 2024, Amazon Web Services, Inc. or its affiliates.

4

4

## Industry Trends

### Centralization

- => Control over deployments
- => Ability to deploy small changes
- => Agile development
- => Fuzzier line between Dev and Ops (DevOps, SRE)
- => Need for continuous performance engineering

[The Past, Present, and Future of Performance Engineering](#)



© 2024, Amazon Web Services, Inc. or its affiliates.

5

## Continuous Performance Testing

- Continuous performance testing
  - To catch regressions early
- Collecting all info needed to investigate regressions
  - In the form convenient for further analysis
- Foundation to build further automation on the top of it
  - For further performance optimization
- All context-dependent
  - Don't wait for an exact recipe, figure it out depending on your needs



© 2024, Amazon Web Services, Inc. or its affiliates.

6

6

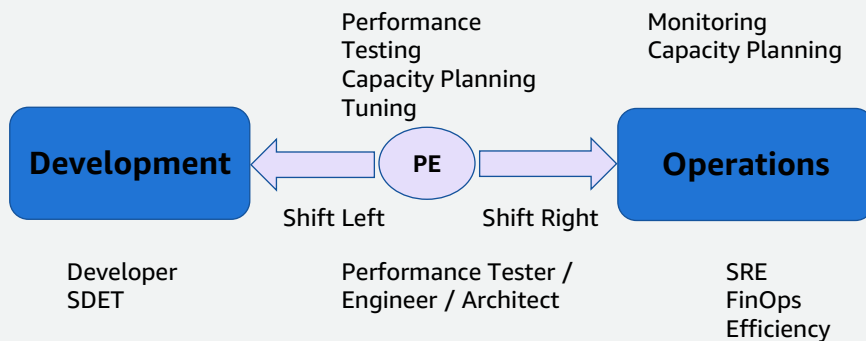
## Performance Testing Traditional vs Continuous

- Before releases
- Realistic Mix
  - As close to production as possible
- Checking Service Level Objectives (SLOs)
- Using a load testing tool or harness
- The approach is relatively consistent and well described
- Often (maybe even each build)
- Different tests
  - To maximize coverage
- Checking the difference between builds
- Using an additional layer of automation on the top of load testing tool
- All context-dependent



7

## Integrating Performance Engineering into DevOps



8

## Challenges of Continuous Performance Testing

- Integration
- Decomposition
- Coverage Optimization
- Variability / Noise Reduction
- Change Detection
- Advanced Analysis
- Operations / Maintenance



## The Challenge of Integration



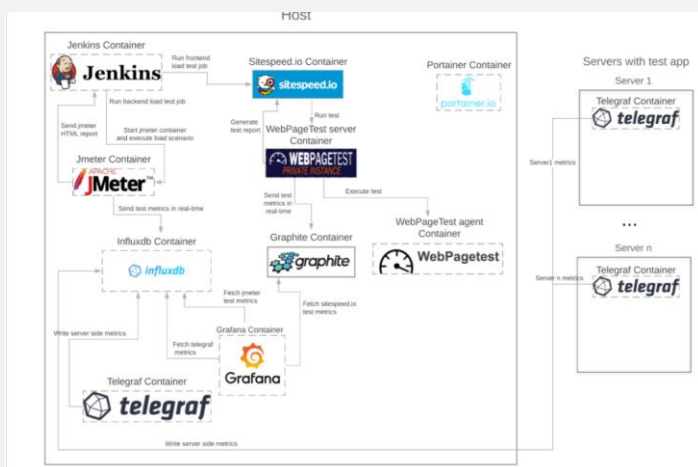
## Continuous Integration: Load Testing Tools

- CI support in load testing tools
  - Integration with CI Servers (Jenkins, Hudson, etc.)
  - Automation support
- CI tools support for performance testing
  - [Jenkins Performance Plugin](#) 
- Performance Testing Frameworks
  - Combining multiple tools



11

## A Performance Testing Framework



12

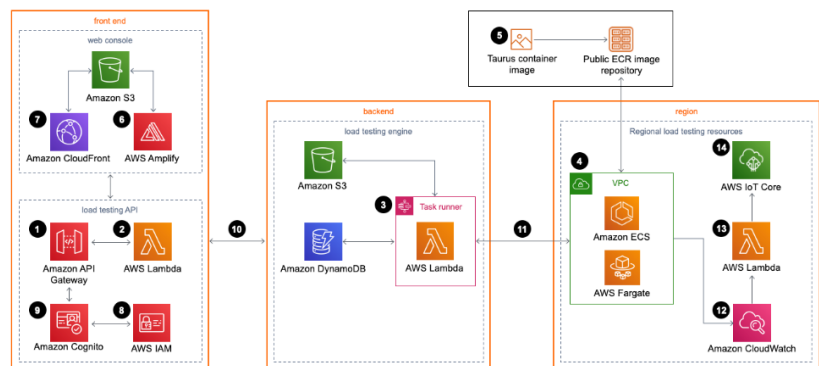
An example:

<https://github.com/serputko/performance-testing-framework>





## Distributed Load Testing on AWS



From AWS Solutions Library

<https://aws.amazon.com/solutions/implementations/distributed-load-testing-on-aws/>



© 2024, Amazon Web Services, Inc. or its affiliates.

13

13

## Closely Integrated Systems

- Sophisticated, but proprietary closely integrated systems
  - [Creating a Virtuous Cycle in Performance Testing at MongoDB](#)
  - [Fallout: Distributed Systems Testing as a Service \(DataStax\)](#)
  - [Tracking Performance of the Graal Compiler on Public Benchmarks \(Charles University / Oracle Labs\)](#)
  - [Introducing Ballast: An Adaptive Load Test Framework \(Uber\)](#)



© 2024, Amazon Web Services, Inc. or its affiliates.

14

14

# The Challenge of Decomposition



© 2024, Amazon Web Services, Inc. or its affiliates.

15

15

## PERFORMANCE TESTING TRANSFORMATION

### Decomposition

- For most complex systems, continuous performance testing should be done on component level / limited scale
  - To align with development
  - System-level requirements -> Component-level requirements
  - Record/playback approach -> Programming
    - Custom Load generation
    - Stubbing/Mocking/Service Virtualization



© 2024, Amazon Web Services, Inc. or its affiliates.

16

16



## Result Interpretation [Modeling]

- If the results are for component / small-scale environment, changes should be modeled into end-to-end performance
  - [Performance Testing and Modeling for New Analytic Applications](#)
  - Or/and confirmed by full-scale end-to-end performance test



17

## The Challenge of Coverage Optimization



18

## Time / Resource Considerations

- Performance tests take time and resources
  - The larger tests, the more
- May be not an option on each commit
- Need of a tiered solution
  - Some performance measurements each commit
  - Daily mid-size performance tests
  - Periodic large-scale / uptime tests *outside CI*




## Coverage Optimization

- A multi-dimensional problem
  - Configuration
  - Workloads / Tests
  - Frequency of runs
- A trade off between coverage and costs
  - Costs of running, analyzing, maintenance, etc.



## The Challenge

- If addressed seriously, the number of workloads / tests / configurations is growing
- No good way to optimize
- One approach is to see if some results are correlated
  - If we find same problems on the same set of tests, we can use just one or few tests from this group
  - [Tracking Performance of the Graal Compiler on Public Benchmarks](#) (Charles University / Oracle Labs) 
- Combinatorial testing approaches (PairWise / Covering Arrays)
  - From functional testing



21


## The Challenge of Variability



22

## Variability - System

- Inherent to the test setup



© 2024, Amazon Web Services, Inc. or its affiliates.

23

23

## Addressing Variability

- [Methodological principles for reproducible performance evaluation in cloud computing. 2019 \(SPEC RG – Cloud\)](#)
- [Reducing variability in performance tests on EC2: Setup and Key Results \(MongoDB\)](#)
- [Tracking Performance of the Graal Compiler on Public Benchmarks](#)



© 2024, Amazon Web Services, Inc. or its affiliates.

24

24

## Addressing Variability

- Same environment / starting config
  - For example, AWS cluster placement groups
- No other load
- Multiple iterations
- Reproducible multi-user tests
  - Restarts between tests
  - Clearing caches / Warming up caches
  - Staggering / Sync points



25

## The Challenge of Change Detection



26

## Complex Results

- No easy pass/fail
  - Individual responses, monitoring results, errors, etc.
- No easy comparison
  - Against SLA
  - Between builds
- Variability



27

## Simple Comparison

### Jenkins Performance Plugin

URI	Samples	Samples diff	Average (ms)	Average diff (ms)
001 home	1	0	347	-22
005 login	1	0	2438	-66
157 views	1	0	117	-33
173 open volume view	1	0	84792	3945
261 search 1M balanced viewpoint	1	0	10964	4295
262 navigate 1M balanced viewpoint	1	0	208	-47
268 open 1M flat viewpoint	1	0	17462	-1562
272 open 1M grid	1	0	5040	572
282 search 1M grid	1	0	2247	8
283 navigate 1M grid	1	0	8343	-181
286 open 200k balanced viewpoint	1	0	16890	-3703
289 search 200k balanced viewpoint	1	0	1261	-1027
290 navigate 200k balanced viewpoint	1	0	148	10
296 validate 200k viewpoint	1	0	81126	723



28

[keptn.sh](https://keptn.sh)



```

1 ---
2 spec_version: "1.0"
3 comparison:
4   aggregate_function: "avg"
5   compare_with: "single_result"
6   include_result_with_score: "pass"
7   number_of_comparison_results: 1
8 filter:
9 objectives:
10  - sli: "response_time_p95"
11    key_sli: false
12    pass: # pass if (relative change <= 10% AND absolute value is < 600ms)
13      - criteria:
14        - "<=+10%" # relative values require a prefixed sign (plus or minus)
15        - "<600" # absolute values only require a logical operator
16      warning: # if the response time is below 800ms, the result should be a warning
17        - criteria:
18          - "<=800"
19      weight: 1
20 total_score:
21   pass: "90%"
22   warning: "75%"

```

Quality Gates  
SLIs / SLOs as code



29

## Change Point Detection

- Statistical methods taking noise in consideration
- E-Divisive means algorithm
  - ICPE Paper: [Change Point Detection in Software Performance Testing](#)
  - [Fixing Performance Regressions Before they Happen](#), Netflix Technology Blog
  - <https://github.com/mongodb/signal-processing-algorithms>
    - Open sourced, generic
  - Need several data points. May miss a gradual degradation.



30

# The Challenge of Advanced Analysis



© 2024, Amazon Web Services, Inc. or its affiliates.

31

31

## PERFORMANCE TESTING TRANSFORMATION

### Keep All Artifacts for Further Analysis

- Get all metrics
  - Throughputs, latencies, resource utilizations, etc.
- Save all related artifacts
  - Exact code / configuration
  - Logs, etc.
- Ability to re-run the test in the exactly same configuration is helpful



© 2024, Amazon Web Services, Inc. or its affiliates.

32

32



## Root Cause Analysis


- Collecting artifacts to do root cause analysis
- Insights snapshots
  - Flamegraphs (perf, eBPF)
- Continuous Profiling
  - Java Flight Recorder
  - APM
  - Tracing
  - Observability
  - eBPF-based tools



33

## Visualization

- [Visualizing systems and software performance - Report on the GI-Dagstuhl](#)
- Sometimes helps to catch an issue



34

# Looking Beyond Aggregate Info



# The Challenge of Operations and Maintenance

## Operations

- Scheduling / execution tests
- Results analysis
- Triaging and escalating issues
- Maintenance



37

## Coverage / Maintenance Trade-Off



38

## Catching / Troubleshooting Errors

- Catching errors is not trivial
  - Building in checks
  - Depends on interfaces used
    - Protocol-level [recording]
    - GUI
    - API/Programming
    - Production Workloads
- Keeping logs / all info needed to investigate issues



## Changing Interfaces

- If using protocol-level or GUI scripts, minor changes may break them
  - It may be not evident
  - If recording used, a change in interfaces may require to recreate the whole script
- API / Programming is usually more stable / easier to fix
- AI to catch the changes / self-healing scripts



## Who Is Doing Maintenance?

- Who is responsible for what?
- Infrastructure Code
  - Tools, plumbing code, integration
- Specific tests
- Integrated workloads
  - Covered multiple functional areas



41

## SUMMARY

- Adjusting performance testing to industry trends
- Specific challenges should be addressed:
  - Integration
  - Coverage Optimization
  - Variability / Noise Reduction
  - Change Detection
  - Advanced Analysis
  - Operations / Maintenance
- Performance engineering gets more integrated, context-dependent
  - Integrated into both **Development** and **Operations**



42



Thank you!

Alex Podelko  
podealex@amazon.com

© 2024, Amazon Web Services, Inc. or its affiliates.