



Continuous Performance Testing Challenges and Approaches

Alexander Podelko
Staff Performance Engineer
MongoDB

1

ABOUT THE SPEAKER



Alex Podelko

- ▶ Specializing in performance since 1997
- ▶ Staff Performance Engineer at MongoDB
 - ▶ Before worked for Oracle/Hyperion, Intel, and Aetna
- ▶ Board director at CMG

apodelko@yahoo.com

@apodelko

<https://alexanderpodelko.com/blog/>

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.

2

2

Continuous Performance Testing

- ▶ Big and not formalized topic
 - ▶ We just starting to see advances here
- ▶ Covering some challenges here
 - ▶ Not all
 - ▶ And just some possible approaches
- ▶ The main point is that it all context-dependent
 - ▶ Don't wait for exact recipe, you need to figure out your own depending on your needs

3

3

Performance Testing @MongoDB

- ▶ Used here to illustrate concepts
- ▶ David Daly and others talked about it in details
 - ▶ <https://www.daviddaly.me/p/recent-presentations.html>
 - ▶ Many MongoDB-related slides here are adopted from David's presentations
- ▶ Very advanced implementation
 - ▶ Highly optimized for MongoDB



4

4

Why Do We Need Performance Testing to Be Continuous ?

5

5


Agile Development

- ▶ Agile development should be rather a trivial case for performance testing
 - ▶ You have a working system each iteration to test early by definition.
 - ▶ You may need performance testing during the whole project
 - Savings come from detecting problems early

6

6

Cost of fixing defects during earlier phases of application life cycle is significantly lower



Jones, Capers. *Applied Software Measurement: Global Analysis of Productivity and Quality*.

7

7

Paradigm Change

- ▶ Traditional performance approach don't scale well
 - ▶ Even having competent and effective engineers
- ▶ Increased volume exposes the problem
 - ▶ Early testing
 - ▶ Each iteration
- ▶ Remedies: **automation, making performance everyone's job**

8

8

Early Testing - Mentality Change

- ▶ Making performance everyone's job
- ▶ Late record/playback performance testing -> Early Performance Engineering
- ▶ System-level requirements -> Component-level requirements
- ▶ Record/playback approach -> Programming to generate load/create stubs
- ▶ "Black Box" -> "Grey Box"

9

9

What Does Continuous Performance Testing Really Mean?

10

10

My View of Notions

- ▶ Performance testing
 - ▶ Automation
 - Continuous performance testing
 - Automated regression performance testing
 - Can't be done without automation
 - Just one kind of performance testing

11

11

MongoDB Performance Use Cases

- **Detect performance impacting commits (Waterfall)**
- Test impact of proposed code change (Patch Test)
- Diagnose performance regressions (Diagnostics, Profiling)
- **Release support (how do we compare to previous stable?)**
- Performance exploration



From David Daly's presentations

12

12

Many Parts of the Puzzle

- ▶ System Under Test
 - ▶ Usually distributed with meaningful data sets
- ▶ Load Testing Tool / Harness
- ▶ CI plumbing
- ▶ Results analysis / alerting
- ▶ And everything may go wrong

13

13

Performance Testing in Continuous Integration @MongoDB

- ▶ Setup a system under test
- ▶ Run a workload
- ▶ *Report the results*
- ▶ *Visualize the result*
- ▶ Decide (and alert) if the performance changed
- ▶ Automate everything / Keep Noise Down



From David Daly's presentations

14

14

The Challenge of Coverage Optimization

15

15

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

16

16

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.

17

17

The Challenge

- ▶ If addressed seriously, the number of workloads / tests / configurations is growing
 - ▶ As we extend functionality / find gaps in coverage / etc.
 - ▶ If each dev team indeed is working on it, it adds quickly
- ▶ 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, perhaps we can use just one or few tests from this group.

18

18



The Challenge of Integration

19

19

Continuous Integration: Load Testing Tools

- ▶ CI support becoming the main theme
- ▶ Integration with Continuous Integration Servers
 - ▶ Jenkins, Hudson, etc.
 - ▶ Making a part of automatic build process
- ▶ Automation support
- ▶ Cloud support
- ▶ Support of newest technologies



20

20

Jenkins Performance Plugin

Standard Mode

Select mode: Relative Threshold Error Threshold

Build result: Fail build when result files are not present

Use Error thresholds on single build:

Unstable:

Failed:

Average response time threshold:

Use Relative thresholds for build comparison:

(-) (+)

Unstable % Range:

Failed % Range:

Compare with previous Build Compare with Build number


Compare based on:

<https://plugins.jenkins.io/performance/>

21

21

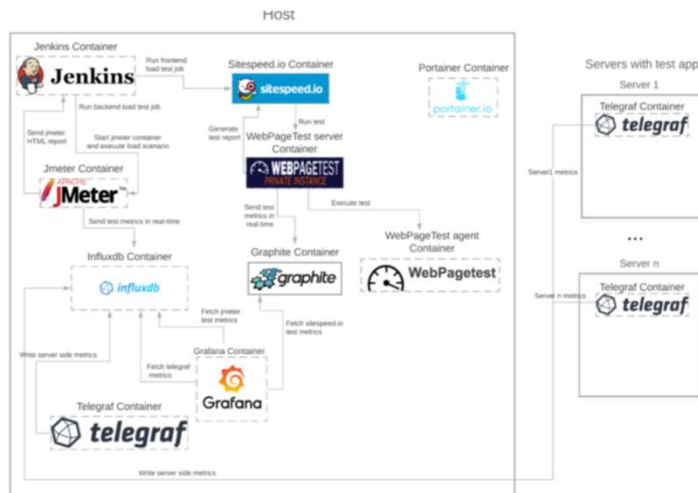
Performance Plugin: Visualization



22

22

A Performance Testing Framework



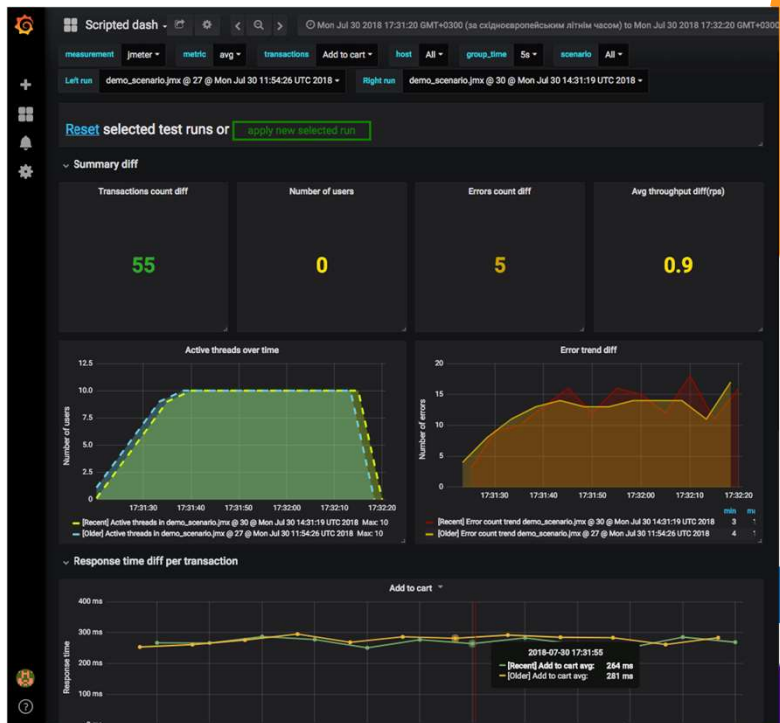
<https://github.com/serputko/performance-testing-framework>

23

23

Performance Testing Framework: Visualization

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



24

24

MongoDB

- ▶ Close integrations of all parts
 - ▶ **CI – Evergreen**
 - ▶ **DSI (Distributed Systems Infrastructure)**
 - ▶ Workload Generation
 - **benchRun, Genny**, industry benchmarks
 - ▶ Git, compilers, Terraform, etc.



25

25

DSI Goals

- Full end-to-end automation
- Support both CI and manual testing
- Elastic, public cloud infrastructure
- Everything configurable
- All configuration via YAML
- Diagnosability
- Repeatability



From David Daly's presentations

26

26

DSI Modules

- Bootstrap
- Infrastructure provisioning
- System setup
- Workload setup
- MongoDB setup
- Test Control
- Analysis
- Infrastructure teardown



From David Daly's presentations

27

27

Configuration Files

```
1 mongod_config_file:
2 storage:
3   engine: wiredTiger
4 replication:
5   replSetName: rs0
6
7 topology:
8   - cluster_type: replset
9     id: rs0
10    mongod:
11      - public_ip: ${infrastructure_provisioning.out.mongod.0.public_ip}
12        - public_ip: ${infrastructure_provisioning.out.mongod.1.public_ip}
13        - public_ip: ${infrastructure_provisioning.out.mongod.2.public_ip}
14
15 # Meta data about this mongodb setup
16 meta:
17   # The list of hosts that can be used in a mongodb connection string
18   hosts: ${mongodb_setup.topology.0.mongod.0.private_ip}:27017
19   hostname: ${mongodb_setup.topology.0.mongod.0.private_ip}
20   mongodb_url: mongodb://${mongodb_setup.meta.hosts}/test?replicaSet=rs0
21   is_replset: true
```

```
1 run:
2   - id: ycsb_load
3     type: ycsb
4     cmd: ./bin/ycsb load mongodb -s -P ../../workloadEvergreen -threads 8
5     config_filename: workloadEvergreen
6     workload_config: |
7       mongodb.url=${mongodb_setup.meta.mongodb_url}
8       recordcount=5000000
9       workload=com.yahoo.ycsb.workloads.CoreWorkload
10    - id: ycsb_100read
11      type: ycsb
12      cmd: ./bin/ycsb run mongodb -s -P ../../workloadEvergreen_100read -threads 32
13      config_filename: workloadEvergreen_100read
14      workload_config: |
15        mongodb.url=${mongodb_setup.meta.mongodb_url}
16        recordcount=5000000
17        maxexecutiontime=240
18        workload=com.yahoo.ycsb.workloads.CoreWorkload
19        readproportion=1.0
```

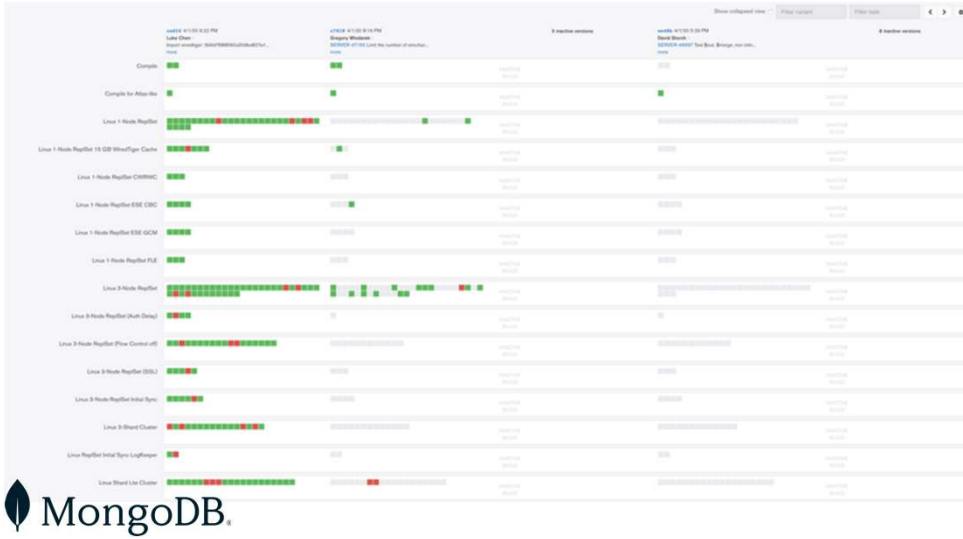


From David Daly's presentations

28

28


Evergreen Integration



The Challenge of Variability

Variability - Environment

- ▶ Due to difference in environments

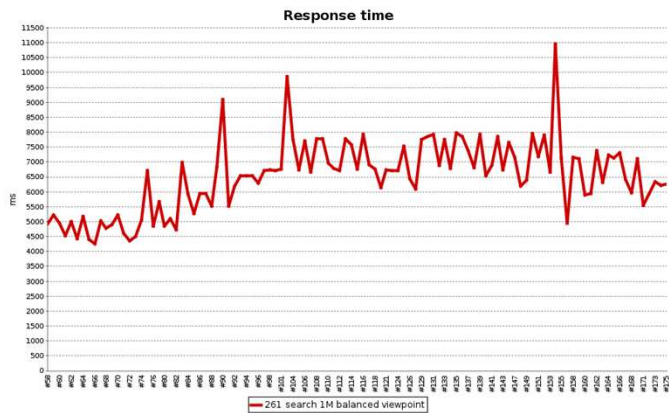


31

31

Variability - System

- ▶ Inherent to the test setup



32

32

Same DevOps concept : Cattle vs Pets



Do everything in repeatable, reproducible way – no manual steps

33

33

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

34

34

The Challenge of Change Detection

35

35

Complex Results

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

36

36

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 viewpont	1	0	81126	723

37

keptn.sh

Quality Gates
SLIs / SLOs as code

```
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%"
```

38

Change Point Detection

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



39

39

Change Point Detection - Visualization

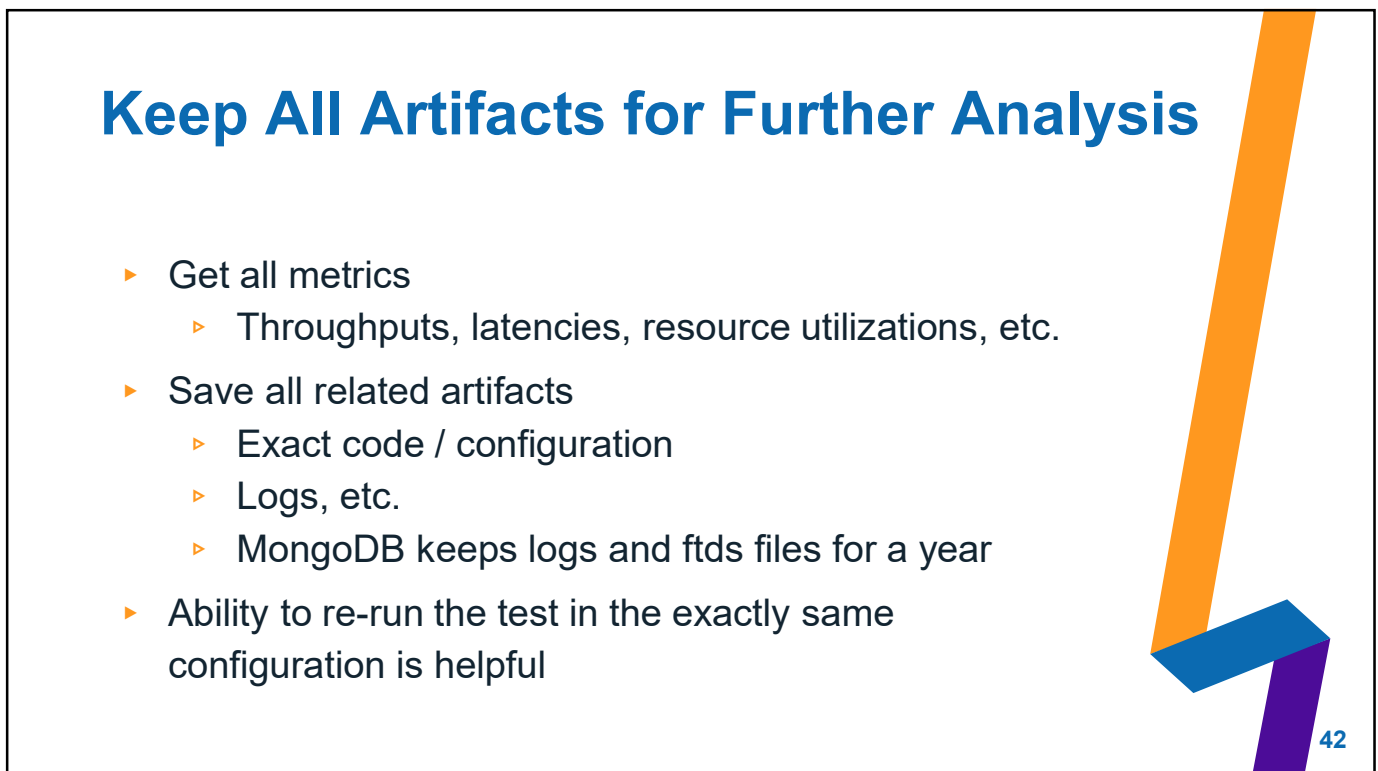


40

40



41



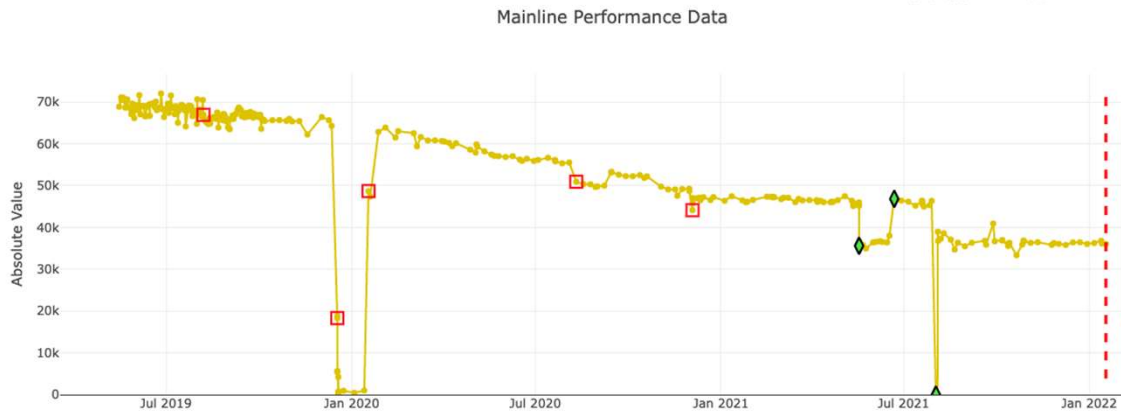
42

Keep All Artifacts for Further Analysis

- ▶ Get all metrics
 - ▶ Throughputs, latencies, resource utilizations, etc.
- ▶ Save all related artifacts
 - ▶ Exact code / configuration
 - ▶ Logs, etc.
 - ▶ MongoDB keeps logs and ftds files for a year
- ▶ Ability to re-run the test in the exactly same configuration is helpful

Visualization


- ▶ Sometimes helps to catch an issue
 - ▶ For example, gradual degradation:



43

43

Looking Beyond Aggregate Info




44

44

Looking at Individual Results Patterns

Scatter charts – a “banding” pattern from <http://www.perftestplus.com/resources/BPT6.pdf>



45

45

The Challenge of Maintenance

46

46

Coverage / Maintenance Trade-Off



47

47

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

48

48

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

49

49

The Challenge of Organization

50

50

Different Roles

- ▶ Consultant: need to test the system
 - ▶ In its current state
 - ▶ External or internal (centralized team)
 - ▶ Why bother about automation?
- ▶ Performance Engineer
 - ▶ On an agile team
 - ▶ Need to test it each build/iteration/sprint/etc.
- ▶ Automation Engineer / SDET / etc.
- ▶ Performance Engineer / Team of the future
 - ▶ TBD

51

51

Performance Engineer / Team of the Future

- ▶ The center of performance expertise (?)
 - ▶ Helping dev teams to create / run tests
 - ▶ Coordinating efforts
 - ▶ Sorting out complex issues
 - ▶ Doing sophisticated investigations

52

52

Who Is Doing Maintenance?

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

53

53

SUMMARY

- ▶ Integrating into agile development is a must
 - ▶ When performance risks need to be mitigated
- ▶ May be implemented in different ways
- ▶ Specific challenges should be addressed:
 - ▶ Optimizing coverage (cost / benefit ratio)
 - ▶ Integration with other CI / DevOps tools
 - ▶ Noise Reduction
 - ▶ Change point detection
 - ▶ Advanced analysis
 - ▶ Role of performance team

54

54

MongoDB is Open to the Community

David Daly and others discussed implementation in detail

- ▶ <https://www.daviddaly.me/p/recent-presentations.html>

Our code is open source: [signal-processing-algorithms](#),
[infrastructure code](#)

Our [regression environment](#) is open, and [the platform](#) is open source



55

55

Data Challenge

- [CFP](#)
- [Notes](#)
- [Data](#)

The screenshot shows the ICPE 2022 website with the following content:

- ICPE 2022**
13th ACM/SPEC International Conference on Performance Engineering
- Beijing, China
April 9-13, 2022
- Data Challenge Track**
Data is the foundation of many important decision-making processes in performance engineering tasks of modern systems. Data can tell us about the past and present of a system's performance, and it can help us to make predictions about the system's performance. Therefore, ICPE 2022 will for the first time be hosting a data challenge track, inspired by several other conferences, such as MSR and PROMISE.
- Research Papers Track**
In this track, an industrial performance dataset will be provided. The participants are invited to come up with research questions about the dataset, and study those. The challenge is open-ended: participants can choose the research questions that they find most interesting. The proposed approaches and/or tools and their findings are discussed in short papers, and presented in the main conference.

Navigation menu on the left includes: Home, Program, Workshops, Venue, Important Dates, Call for Contributions, Tracks and Submissions, Research Papers Track, Industry and Experience Track, Artifact Evaluation Track, Posters and Demonstrations Track.

Tweets by @ICPEconf are visible on the right side of the page.



56

56

Questions ?

apodelko@yahoo.com

[@apodelko](https://twitter.com/apodelko)

<https://alexanderpodelko.com/blog/>

<https://www.linkedin.com/in/alexanderpodelko/>

57