Experiences from Building the Open Database Performance Ranking with benchANT

13th Symposium on Software Performance 2022

Daniel Seybold, Jörg Domaschka
benchANT | Ulm University

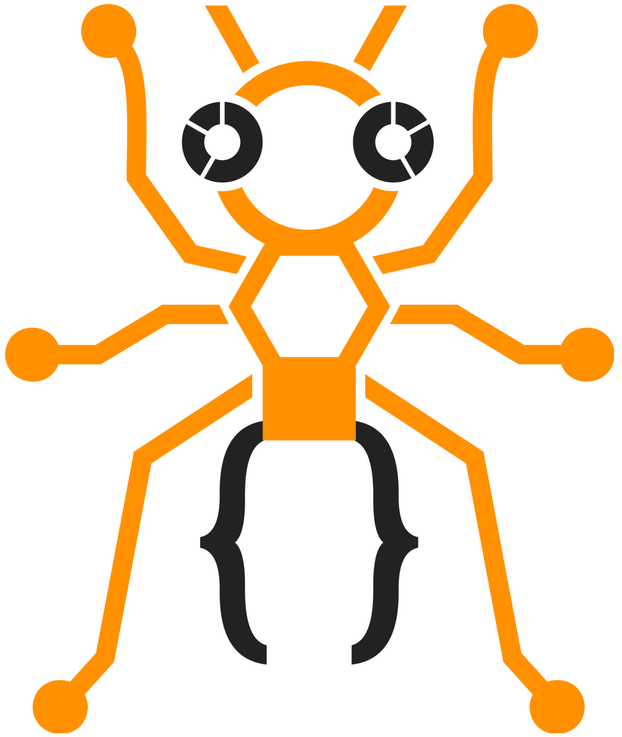

benchANT

Advances of Data Management Technologies for Data-intensive Applications

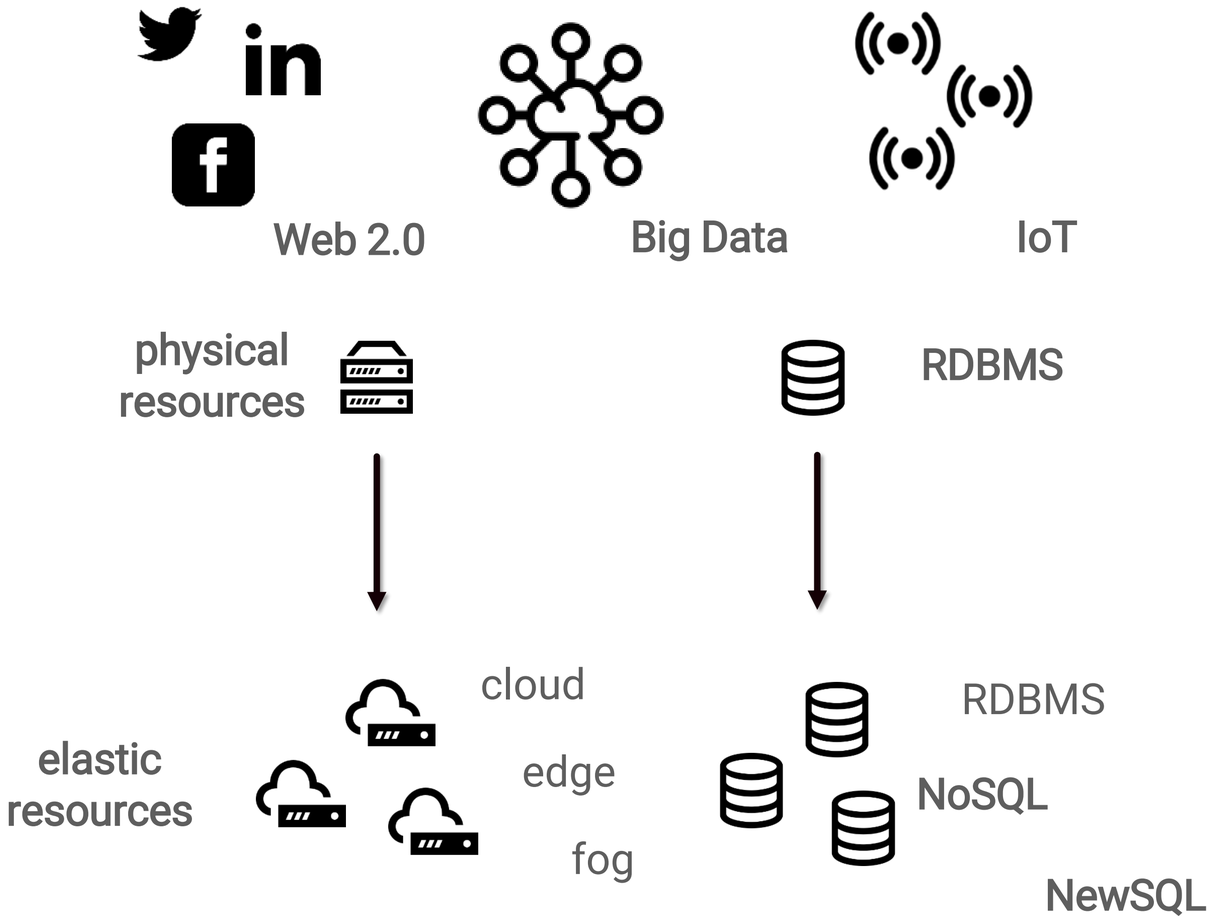

benchANT

Advances of Database Technologies for Data-intensive Applications

Web 2.0

Big Data

IoT

cloud resources have become the preferred solution to operate DBMS [1]

the idea of "one-size-fits-all" is over [2]

DBaaS reached mainstream and serverless DBaaS might be the future [3]

[1]Abadi, Daniel, et al. "The seattle report on database research." ACM SIGMOD Record (2020)
[2]Stonebraker, Michael, and Uğur Çetintemel. "One size fits all" an idea whose time has come and gone." Making Databases Work: the Pragmatic Wisdom of Michael Stonebraker. 2018
[3]Abadi, Daniel, et al. "The seattle report on database research." ACM SIGMOD Record (2022)

benchANT

Comparing Databases

How to get the required data?

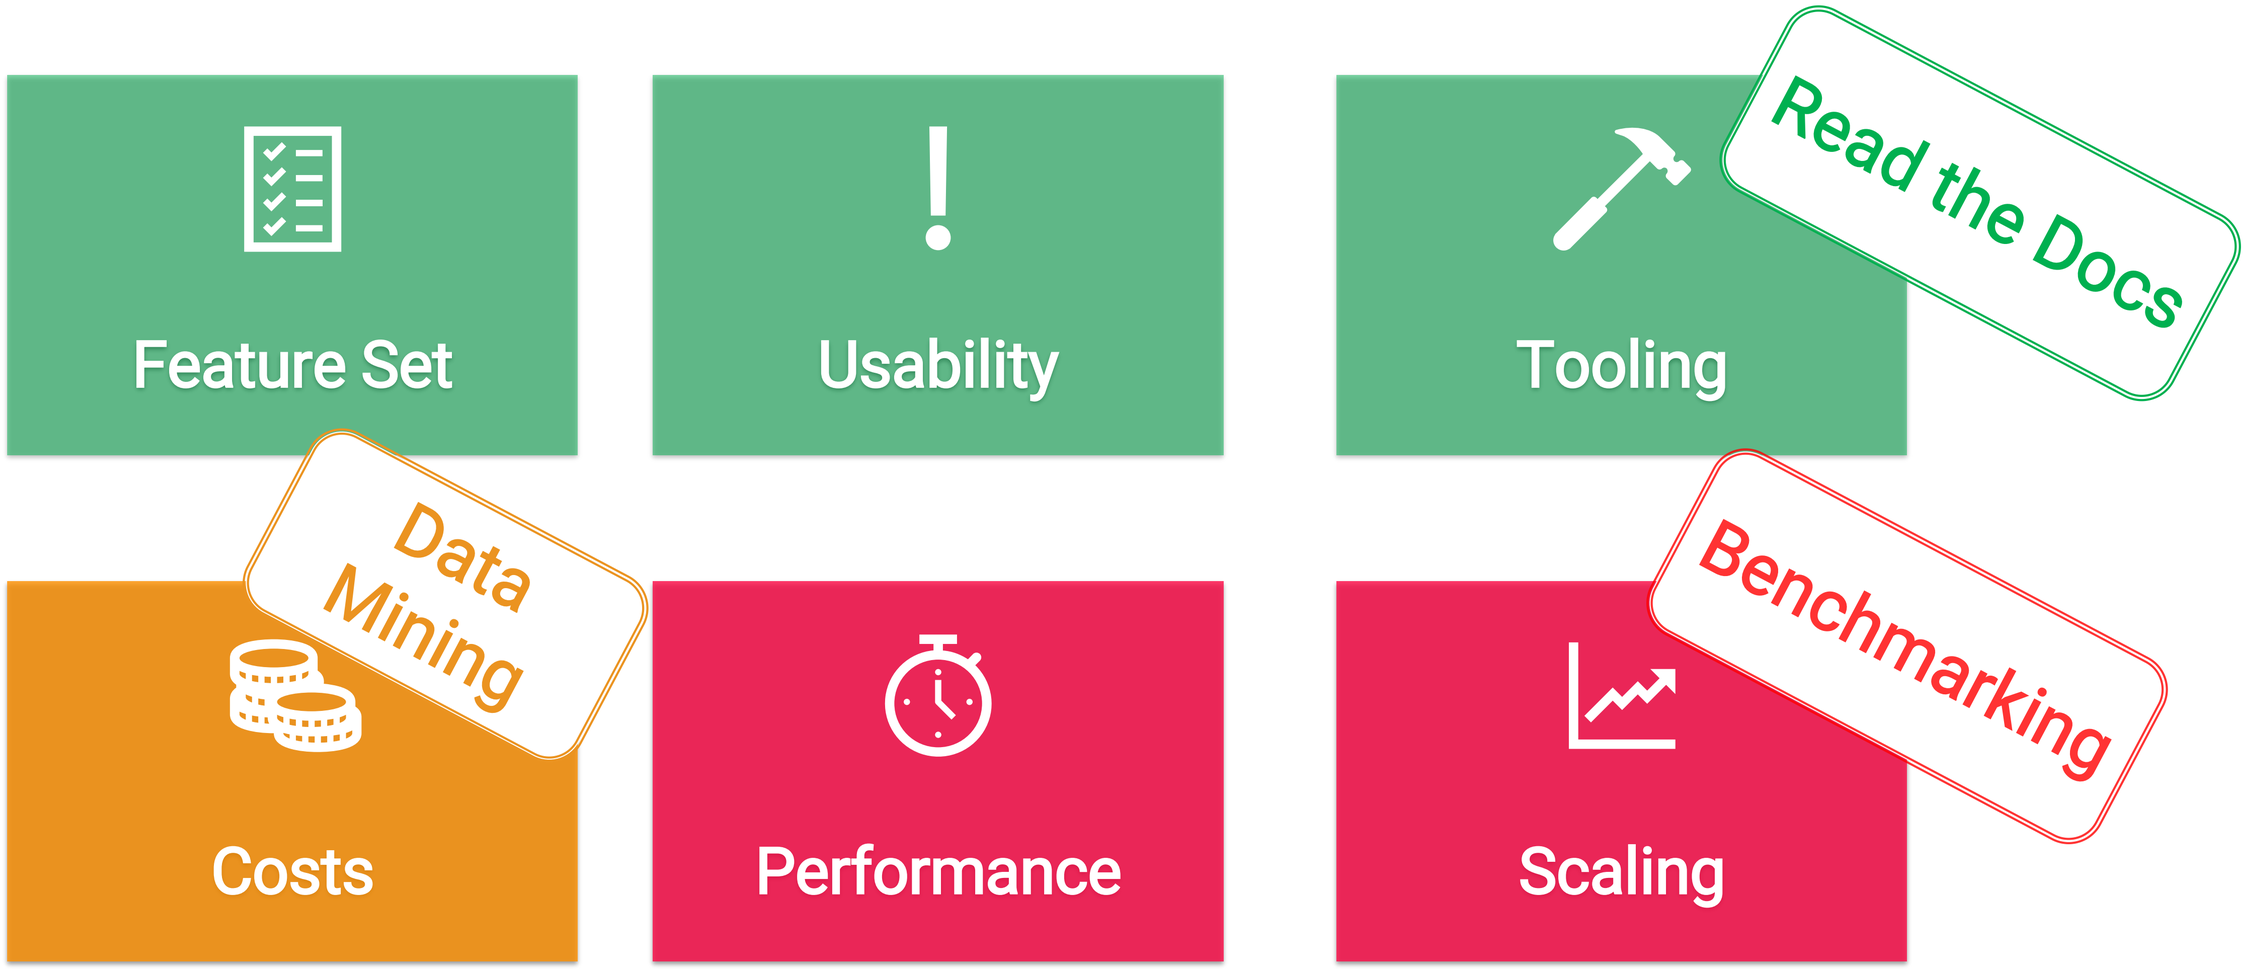

benchANT

Promises of Database Providers

Couchbase

Unparalleled performance at scale

CockroachDB

```
/* A distributed SQL database designed for speed, scale, and
survival. Trusted by thousands of innovators around the globe */
```

PostgreSQL

PostgreSQL is a powerful, open source object-relational database system with over 30 years of active development that has earned it a strong reputation for reliability, feature robustness, and performance.

influxdata

A high-performance time series engine

TIMESCALE

Accelerated performance

Achieve 10-100x faster queries than PostgreSQL, InfluxDB, and MongoDB. Native optimizations for time-series.

GridDB

High Performance

Utilizing an in-memory data architecture — along with superb parallel processing and minimal overhead — grants benchmark-shattering performance

SingleStore

Speed

Accelerate time to insight with a database built for ultra fast ingest and high performance queries

Apache CASSANDRA™

Performant

Cassandra **consistently outperforms** popular NoSQL alternatives in benchmarks and real applications, primarily because of fundamental architectural choices.

benchANT

Is Database Benchmarking still important?

"Benchmarks tremendously helped move forward the database industry and the database research community.

Moreover, without the development of appropriate benchmarking and data sets, a fair comparison … will not be feasible. Benchmarking in the cloud environment also presents unique challenges since differences in infrastructure across cloud providers makes apples to apples comparison more difficult. A closely related issue is reproducibility of performance results in database publications." -- Seattle Report on Database Research 2022 [1]

[1] Abadi, Daniel, et al. "The seattle report on database research." ACM SIGMOD Record (2022)

benchANT

Reproducible Cloud Database Benchmarking

- scientific guidelines for reproducible cloud benchmarking[1]
- scientific guidelines for reproducible cloud-hosted database benchmarking[2]
- leading database and performance engineering conferences enforce available and reproducible benchmarking data sets (VLDB, SIGMOD, ICPE, …)
- BUT: Leznik et al. show that only a very limited number of performance related research results release their benchmark results as open data sets[3]

[1]A. V. Papadopoulos et al. “Methodological Principles for Reproducible Performance Evaluation in Cloud Computing”. In: IEEE Transactions on Software Engineering 47.8 (2021)
[2]D. Seybold. “An automation-based approach for reproducible evaluations of distributed DBMS on elastic infrastructures”. PhD thesis. 2021
[3]M. Leznik et al. “Same, Same, but Dissimilar: Exploring Measurements for Workload Time-Series Similarity”. In: ACM/SPEC ICPE. ACM, 2022

benchANT

Experiences from Building a Global Database Performance Ranking

How to ensure a fully transparent and reproducible global database performance ranking?

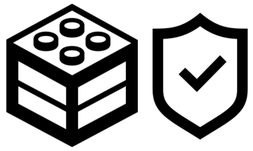

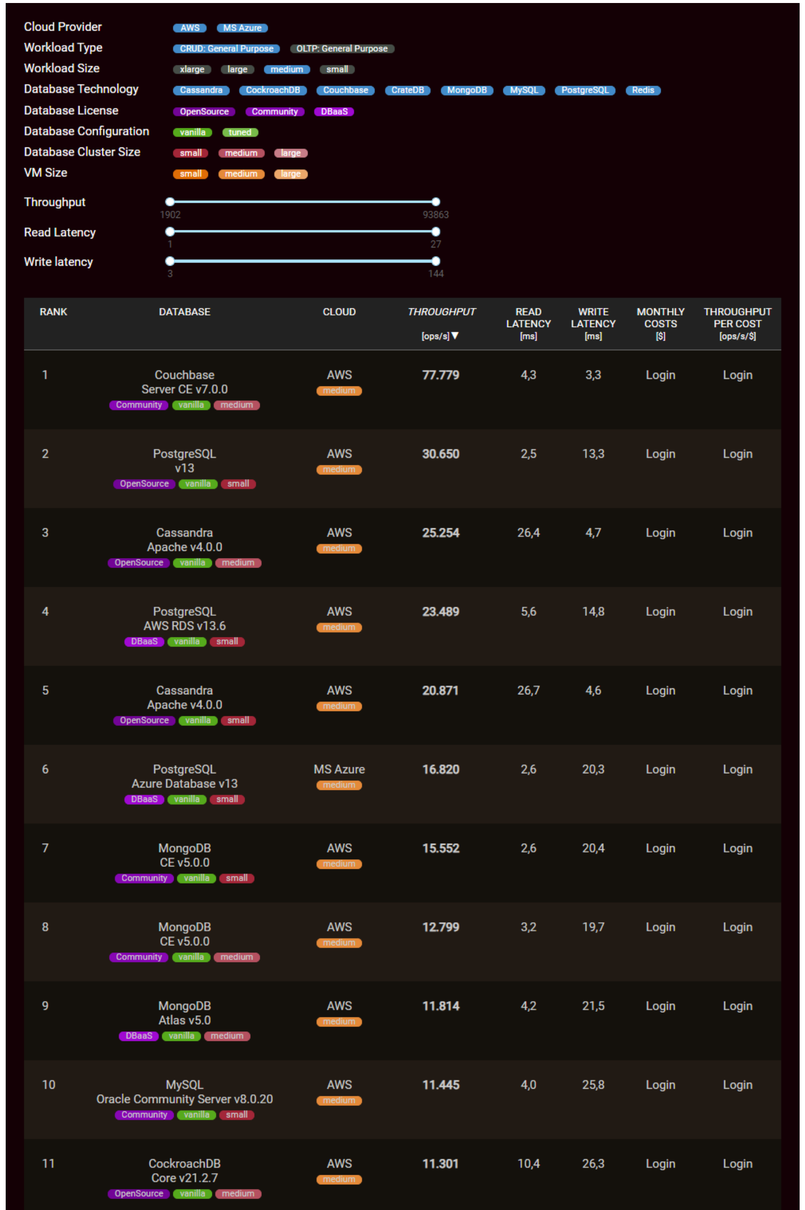

RANK	DATABASE	CLOUD	*THROUGHPUT* [ops/s]▼	READ LATENCY [ms]	WRITE LATENCY [ms]	MONTHLY COSTS [$]	THROUGHPUT PER COST [ops/s/$]
1	Couchbase Server CE v7.0.0 Community vanilla medium	AWS medium	**77.779**	4,3	3,3	Login	Login
2	PostgreSQL v13 OpenSource vanilla small	AWS medium	**30.650**	2,5	13,3	Login	Login
3	Cassandra Apache v4.0.0 OpenSource vanilla medium	AWS medium	**25.254**	26,4	4,7	Login	Login
4	PostgreSQL AWS RDS v13.6 DBaaS vanilla small	AWS medium	**23.489**	5,6	14,8	Login	Login
5	Cassandra Apache v4.0.0 OpenSource vanilla small	AWS medium	**20.871**	26,7	4,6	Login	Login
6	PostgreSQL Azure Database v13 DBaaS vanilla small	MS Azure medium	**16.820**	2,6	20,3	Login	Login
7	MongoDB CE v5.0.0 Community vanilla small	AWS medium	**15.552**	2,6	20,4	Login	Login
8	MongoDB CE v5.0.0 Community vanilla medium	AWS medium	**12.799**	3,2	19,7	Login	Login
9	MongoDB Atlas v5.0 DBaaS vanilla medium	AWS medium	**11.814**	4,2	21,5	Login	Login
10	MySQL Oracle Community Server v8.0.20 Community vanilla small	AWS medium	**11.445**	4,0	25,8	Login	Login
11	CockroachDB Core v21.2.7 OpenSource vanilla medium	AWS medium	**11.301**	10,4	26,3	Login	Login

Which insights can you get out of the global database performance ranking?

benchANT

Requirements for a Reproducible and Transparent Cloud Database Ranking

- requirements are derived from the scientific guidelines for cloud and database benchmarking
- imposed by cloud and database providers
- R1: provide raw and aggregated performance data
- R2: provide dynamic configurations for cloud, database and workload domain
- R3: provide monitoring data for all involved components
- R4: enable a performance audit -> which benchmark step X is executed at time T^x

benchANT

benchANT Background

benchANT

End-to-end Benchmark Automation with Mowgli

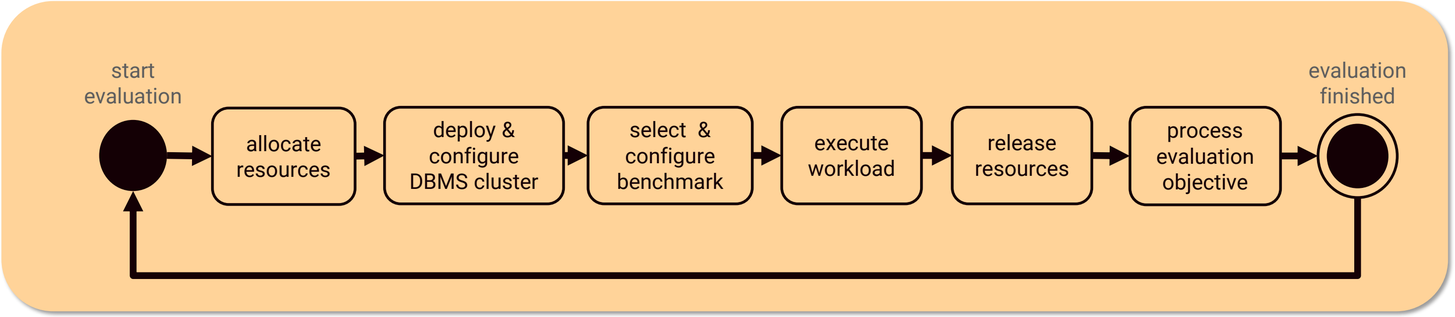

Mowgli Framework

https://research.spec.org/tools/overview/mowgli/

- fully automated benchmarking process
- guaranteed transparency by reproducibility
- comprehensive data sets

Evaluation Specification

Evaluation Execution

Multi-objective Analysis

Automating Performance & Scalability Evaluations

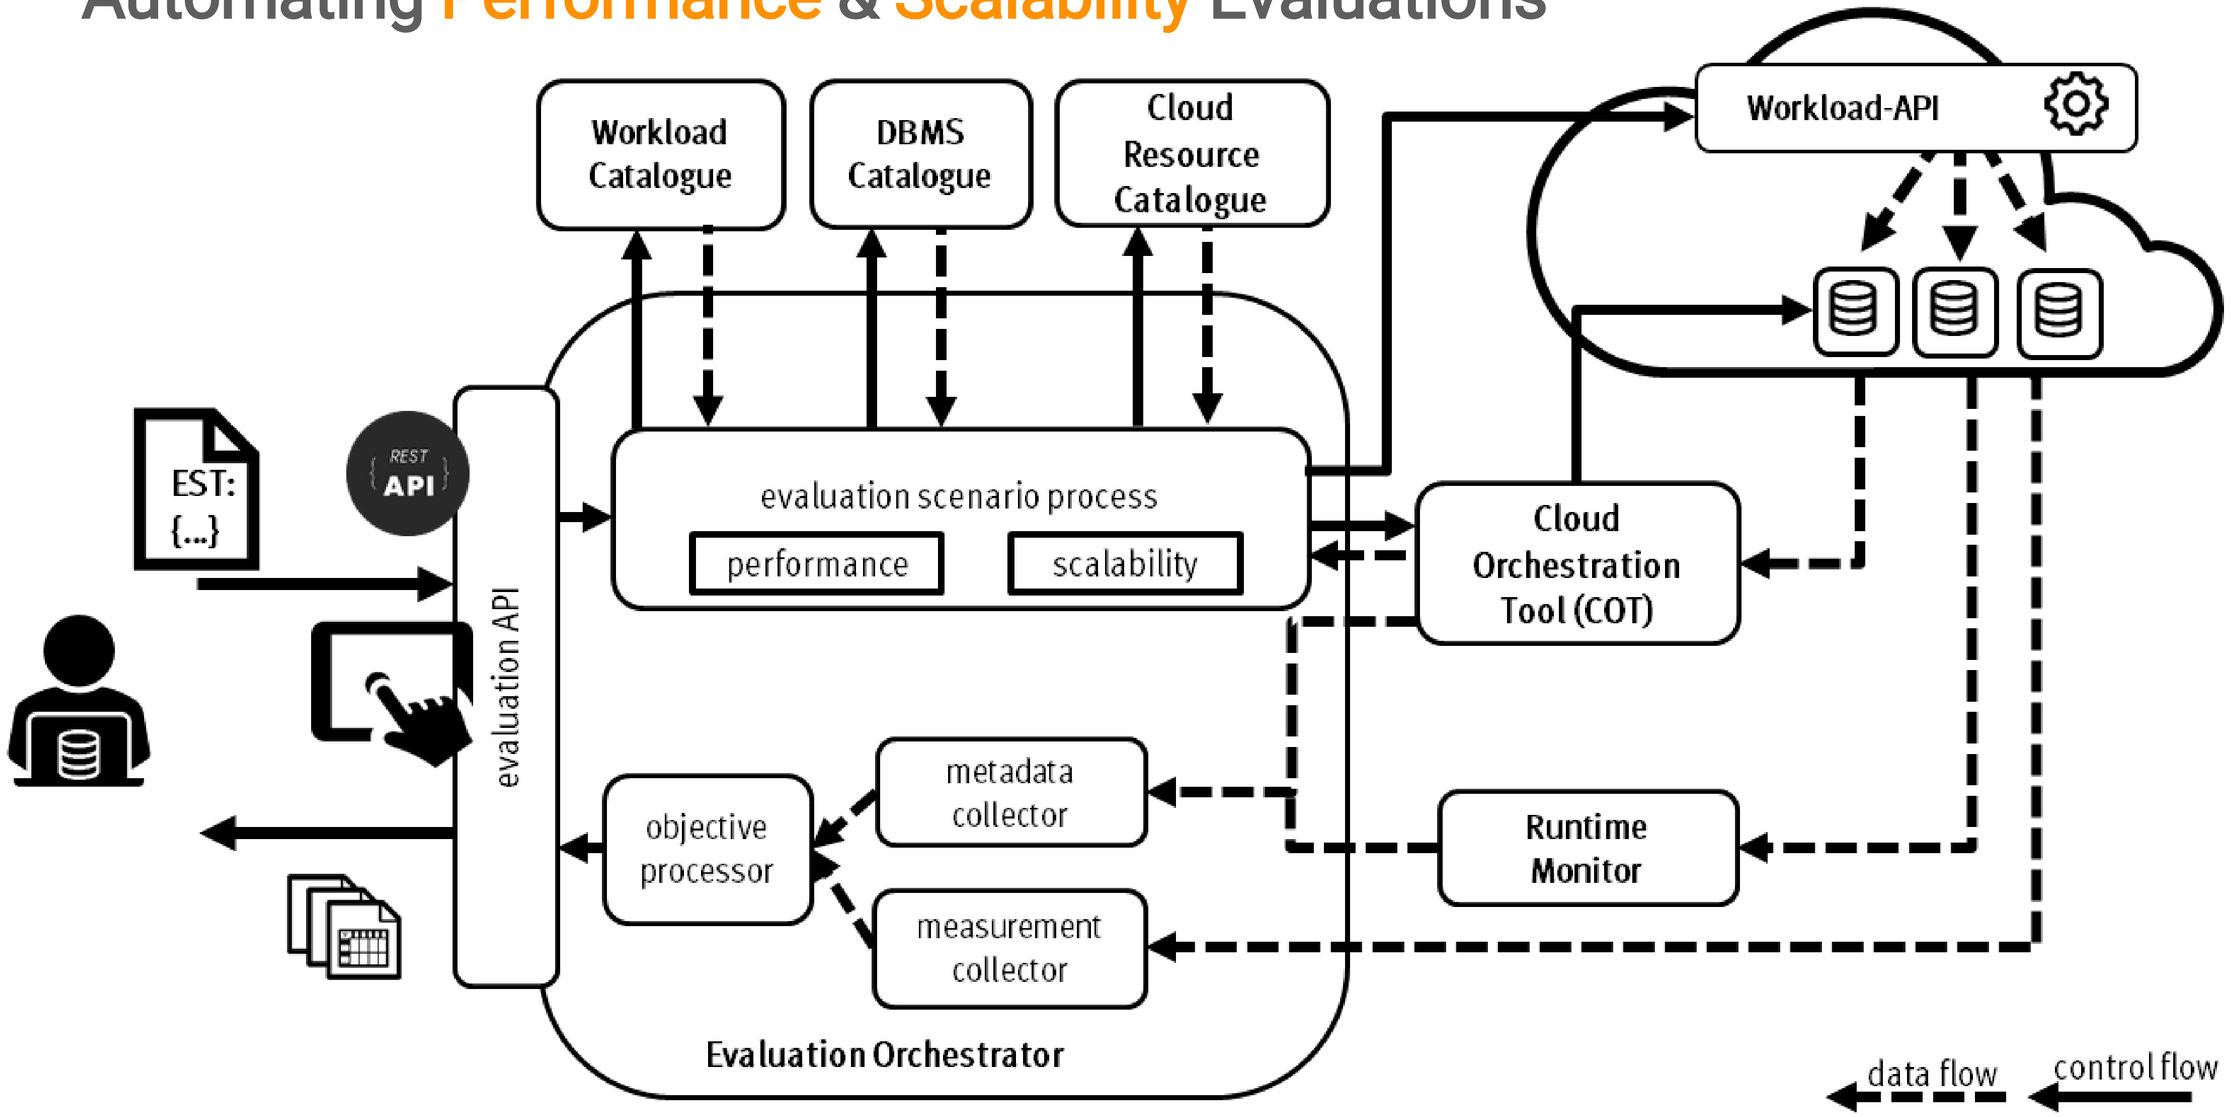

Daniel Seybold, Moritz Keppler, Daniel Gründler, and Jörg Domaschka. "Mowgli: Finding Your Way in the DBMS Jungle", ICPE 2019

From Mowgli to benchANT — Benchmarking-as-a-Service

Mowgli Framework

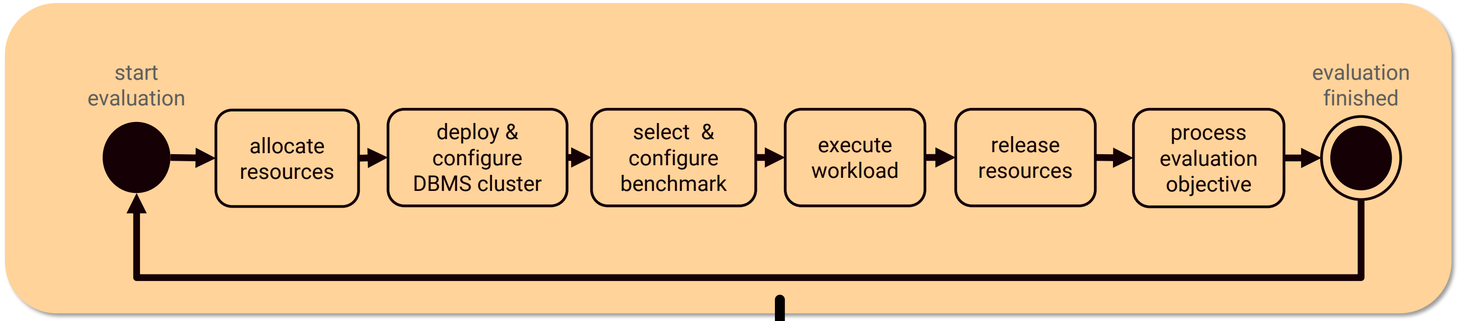

benchANT

From Mowgli to benchANT — Benchmarking-as-a-Service

Mowgli Framework

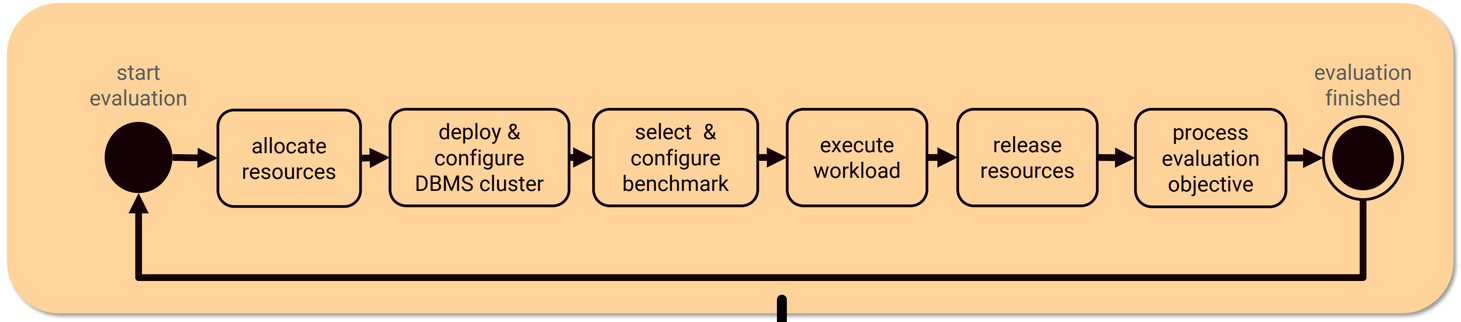

benchANT

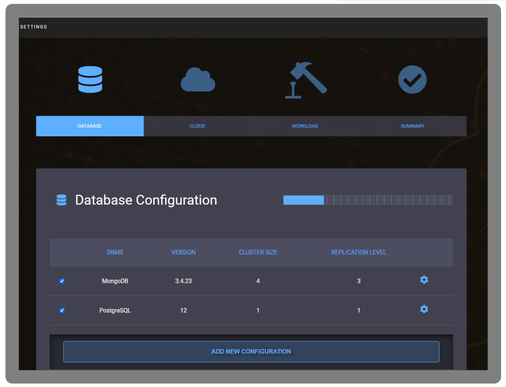

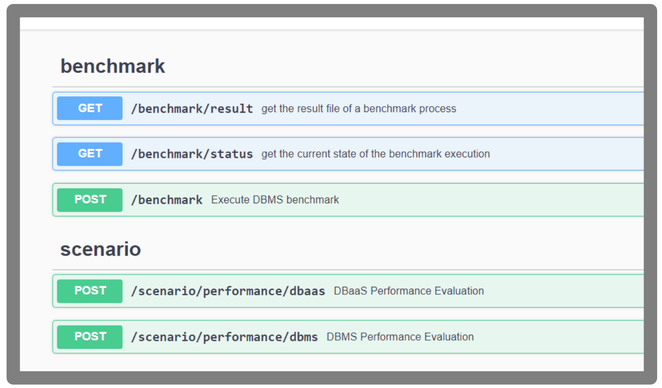

benchANT

From Mowgli to benchANT — Benchmarking-as-a-Service

Mowgli Framework

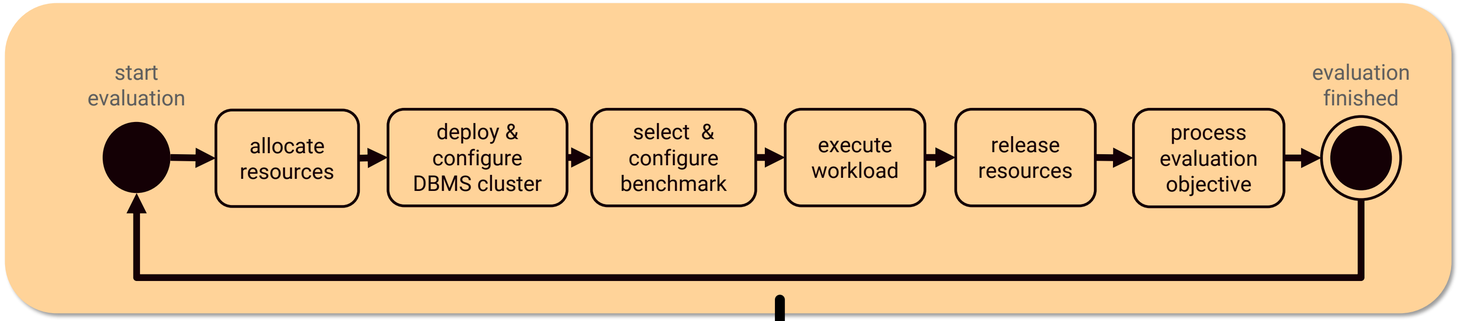

benchANT

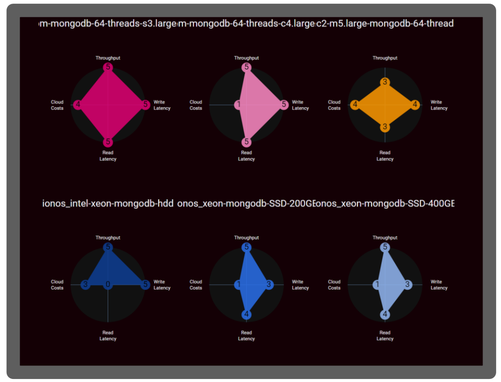

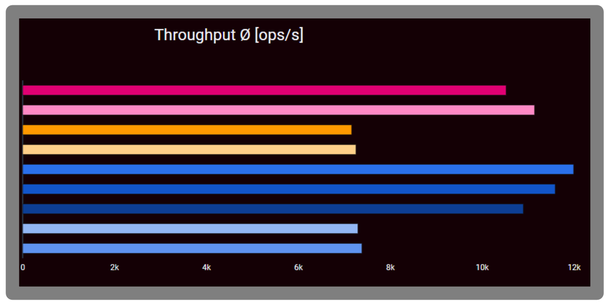

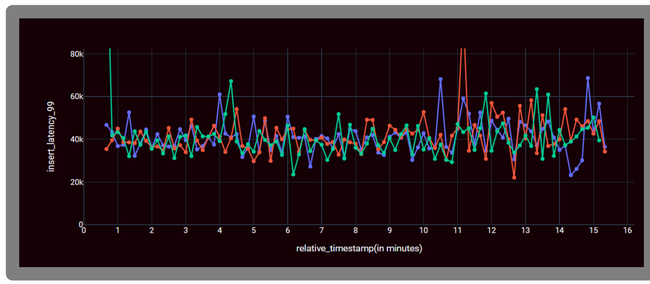

benchANT

Performance Insights

https://benchant.com/ranking/database-ranking

benchANT

Insights: Database Performance (YCSB read-write workload)

RANK	DATABASE	CLOUD	*THROUGHPUT* [ops/s]▼	READ LATENCY [ms]	WRITE LATENCY [ms]	MONTHLY COSTS [$]	THROUGHPUT PER COST [ops/s/$]
1	PostgreSQL v13 OpenSource vanilla small	AWS medium	**34.976**	2,7	12,1	178	196,90
2	Couchbase Server CE v7.0.0 Community vanilla small	AWS medium	**29.602**	36,2	2,1	178	166,70
3	PostgreSQL AWS RDS v13.6 DBaaS vanilla small	AWS medium	**23.489**	5,6	14,8	323	72,70
4	Cassandra Apache v4.0.0 OpenSource vanilla small	AWS medium	**20.871**	26,7	4,6	178	117,50
5	ScyllaDB v4.5.1 OpenSource vanilla small	AWS medium	**17.529**	31,2	2,3	178	98,70
6	MongoDB CE v5.0.0 Community vanilla small	AWS medium	**15.552**	2,6	20,4	178	87,60
7	MySQL Oracle Community Server v8.0.20 OpenSource vanilla small	AWS medium	**11.799**	2,9	25,7	178	66,40
8	CockroachDB Core v21.2.7 OpenSource vanilla small	AWS medium	**7.425**	19,6	40,7	178	41,80
9	CrateDB v4.7.0 OpenSource vanilla small	AWS medium	**3.182**	0,7	87,2	178	17,92

benchANT

Insights: Database Scalability (YCSB read-write workload)

RANK	DATABASE	CLOUD	*THROUGHPUT* [ops/s]▼	READ LATENCY [ms]	WRITE LATENCY [ms]	MONTHLY COSTS [$]	THROUGHPUT PER COST [ops/s/$]
1	Cassandra Apache v4.0.0 OpenSource vanilla large	AWS large	**139.171**	10,8	10,6	3.089	45,10
2	Cassandra Apache v4.0.0 OpenSource vanilla medium	AWS large	**62.163**	16,1	4,8	1.030	60,40
3	Cassandra Apache v4.0.0 OpenSource vanilla medium	AWS medium	**25.254**	26,4	4,7	533	47,40
4	Cassandra Apache v4.0.0 OpenSource vanilla small	AWS medium	**20.871**	26,7	4,6	178	117,50
5	Cassandra Apache v4.0.0 OpenSource vanilla small	AWS small	**12.312**	24,9	4,4	95	129,90

RANK	DATABASE	CLOUD	*THROUGHPUT* [ops/s]▼	READ LATENCY [ms]	WRITE LATENCY [ms]	MONTHLY COSTS [$]	THROUGHPUT PER COST [ops/s/$]
1	ScyllaDB v4.5.1 OpenSource vanilla large	AWS large	**204.405**	4,9	5,4	3.089	66,20
2	ScyllaDB v4.5.1 OpenSource vanilla medium	AWS large	**50.621**	4,2	3,4	1.030	49,20
3	ScyllaDB v4.5.1 OpenSource vanilla medium	AWS medium	**18.646**	6,8	2,2	533	35,00
4	ScyllaDB v4.5.1 OpenSource vanilla small	AWS medium	**17.529**	31,2	2,3	178	98,70
5	ScyllaDB v4.5.1 OpenSource vanilla small	AWS small	**11.708**	42,4	2,1	95	123,50

benchANT

Insights: IaaS Resource Performance & Costs

RANK	DATABASE	CLOUD	*THROUGHPUT* [ops/s]▼	READ LATENCY [ms]	WRITE LATENCY [ms]	MONTHLY COSTS [$]	THROUGHPUT PER COST [ops/s/$]
1	PostgreSQL v13 OpenSource vanilla small	Alibaba Cloud small	**21.652**	3,1	10,3	96	224,40
2	PostgreSQL v13 OpenSource vanilla small	IONOS Cloud small	**20.834**	7,1	9,2	118	176,00
3	PostgreSQL v13 OpenSource vanilla small	AWS small	**19.447**	3,0	8,8	95	205,10
4	PostgreSQL v13 OpenSource vanilla small	MS Azure small	**8.622**	1,7	14,8	87	99,50

benchANT

Insights: DBaaS Peformance (YCSB read-write workload)

RANK	DATABASE	CLOUD	*THROUGHPUT* [ops/s]▼	READ LATENCY [ms]	WRITE LATENCY [ms]	MONTHLY COSTS [$]	THROUGHPUT PER COST [ops/s/$]
1	MongoDB CE v5.0.0 Community vanilla medium	AWS medium	**12.799**	3,2	19,7	533	24,00
2	MongoDB Atlas v5.0 DBaaS vanilla medium	AWS medium	**11.814**	4,2	21,5	871	13,60

RANK	DATABASE	CLOUD	*THROUGHPUT* [ops/s]▼	READ LATENCY [ms]	WRITE LATENCY [ms]	MONTHLY COSTS [$]	THROUGHPUT PER COST [ops/s/$]
1	PostgreSQL v13 OpenSource vanilla small	AWS medium	**34.976**	2,7	12,1	178	196,90
2	PostgreSQL AWS RDS v13.6 DBaaS vanilla small	AWS medium	**23.489**	5,6	14,8	323	72,70

benchANT

Benchmarking Data Structure

benchANT

Data Set Structure

- **R1:** provide raw and aggregated performance data
  - performance data for the load and run phase is provided as time-series and aggregated

- **R2:** provide dynamic configurations for cloud, database and workload domain
  - configurable benchmark parameters are defined in a model (*evaluationScenario.json)*
  - cloud & VM & database configurations are collected

- **R3:** provide monitoring data for all involved components
  - system metrics for database and benchmark instances are collected

- **R4:** enable a performance audit -> which benchmark step X is executed at time T^x
  - a task execution log for all executed benchmark steps is provided (*airflowTaskInstanceDetails.json*)

benchANT

Data Set Structure

- all data is available on GitHub: https://github.com/benchANT/database-ranking
- reproducibility of the results is validated by multiple database providers
- validation was carried out by using the benchANT platform and by executing the benchmarks manually based on the publicly available data sets

benchANT

Conclusion

- database benchmarking is still a highly relevant task to advance database research while cloud computing adds another level of complexity
- database benchmarking needs to ensure reproducible and transparent data sets, currently only a limited number of benchmarking studies follows these requirements
- based on a global database performance ranking, we provide a reference data set structure for reproducible and transparent performance results
- reproducibility is validated by multiple database providers
- comprehensive performance data sets are the foundation for advanced database research such as configuration auto-tuning

benchANT

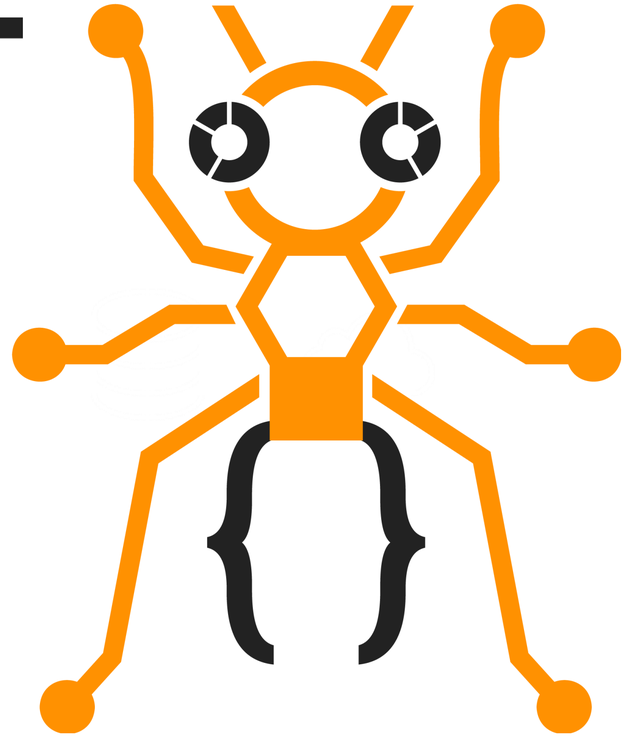

Thank you!

Daniel Seybold
benchANT | daniel.seybold@benchant.com