Experiences from Building the Open Database Performance Ranking with benchANT

Daniel Seybold
daniel.seybold@benchant.com
benchANT, Ulm, Germany

Jörg Domaschka
joerg.domaschka@benchant.com
benchANT, Ulm, Germany

Abstract

Benchmarking is an important method to advance database management systems (DBMS) from the industry and research perspective. Ensuring transparent and reproducible results is a key requirement to ensure the acceptance and credibility of benchmarking. To advance the research towards transparent and reproducible benchmark data, we report on building an open DBMS performance ranking with 130 benchmark configurations and ensuring comparability, transparency and reproducibility. We derive the required data on cloud, resource, DBMS and benchmark level to enable transparency and reproducibility and demonstrate the generation of such data sets with benchANT. Building upon such data, we outline future research directions for DBMS performance modelling, DBMS auto-tuning and decision support.

1 Introduction

DBMS benchmarking clearly helps to advance the DBMS industry and research community [9]. However, the acceptance and utilization of existing results is often limited by the fact that results are not fully transparent and not reproducible due to missing technical details or the software artefacts to reproduce them [5]. Leznik et al. even show that only a very limited number of performance related research results release their benchmark results as open data sets [11].

In order to address these challenges, we report on our experiences in building the first open database performance ranking (ODPR) focusing on transparent and reproducible benchmark results. Starting from previous research [7] together with industry requirements form DBMS and cloud providers, we report on how we structured data sets to ensure comparability, transparency and reproducibility. We envision that the extent and level of detail of the resulting data sets eases the process of generating data sets for future DBMS and cloud benchmark studies.

The following Section 2 presents the technical background on benchANT[1] [8] and the ODPR. Section 3 describes the structure of the data set. Section 4 outlines future research directions and Section 5 concludes.

[1]`https://benchant.com`

2 Open DBMS Performance Ranking

Performance benchmarking of cloud resources, DBMS and distributed systems in general is a continuously carried out task in research as shown by Leznik et al. [11]. Regarding DBMS, there are many performance studies available [11] but each of them only covers a dedicated area of the DBMS and cloud landscape and none provides a global performance overview. In order to provide a global performance overview, we build the open DBMS performance ranking (ODPR)[2].

To establish a global DBMS comparability[3], we select a basic workload that only includes simple write and read operations which can be applied to any data model. Further workloads covering specific data models are implemented in data model specific rankings. In consequence, the resulting workload does not reflect any real world workload and the results are only a very first performance indicator.

2.1 Implementation

We are able to provide this novel ranking by building upon the DBMS benchmark automation framework Mowgli[4] [2] and its enterprise version benchANT[8]. benchANT enables the definition of comparable DBMS benchmark configurations in a multi-cloud context and ensures the deterministic and reproducible benchmark execution.

Benchmark scenarios cover domain-specific properties such as the DBMS itself, cluster size, replication factor, cloud provider, resource capacity, and workload type are defined in a JSON-based model and submitted to the Mowgli API. Internally, Mowgli maps the model to the respective technical implementations and applies a cloud orchestration engine for the cloud resource allocation, DBMS and workload deployment. In addition, monitoring and data processing components are applied to collect and process benchmark and system metrics.

And the ODPR will be continuously extended and in its current state the ODPR provides performance data for seven DBMS operated on IaaS resources, three DBaaS providers, four cloud providers and three

[2]`https://benchant.com/ranking/database-ranking`
[3]relational DBMS, NoSQL DBMS and NewSQL DBMS
[4]`https://research.spec.org/tools/overview/mowgli/`

benchmark suites, resulting in 130 benchmark results. For all benchmark scenarios, the benchmark driver is deployed on a separate VM with 16 cores within the same region and network to ensure that the benchmark instance is not a bottleneck. Currently, each benchmark is executed only once which makes the results susceptible to cloud performance fluctuations but upcoming releases will consist of multiple iterations and a cloud volatility metric.

2.2 Performance Insights

The ODPR enables the high level performance and performance per $ comparison of DBMS, DBaaS and IaaS cloud resources. Moreover, it provides first insights into vertical and horizontal scalability of different relational, NoSQL and NewSQL DBMS as the ODPR includes benchmark configurations with increasing DBMS cluster sizes, cloud resource sizes and workload intensities. These DBMS performance results are available for different workloads ranging from simple write-read, transactional to time-series workloads. In the following paragraphs, we present an excerpt of interesting performance related findings.

DBMS Performance For single node deployments the relational DBMS PostgreSQL achieves better performance than NoSQL and NewSQL DBMS while NoSQL and NewSQL DBMS provide good horizontal scalability results.

Cloud Performance VM types that are comparable based on their compute, storage and cost dimensions can provide significantly different performance, *e.g.* 8622 ops/s on `Azure Standard_D2s_v4` vs. 19447 ops/s on `AWS EC2 m5.large` for PostgreSQL.

DBaaS performance DBaaS instances are not tuned by default since comparable instance types deployed on IaaS achieve similar or even better performance, e.g. a three node cluster of MongoDB on AWS EC2 achieves 12799 ops/s while MongoDB Atlas on AWS achieves 11814 ops/s.

3 ODPR Data Set

The fair comparison and reproducibility of DBMS benchmarks are major challenges as pointed out by Abadi et al. [9]. With the ODPR we address these challenges by building reference data sets that not only cover the benchmark results but all required details to ensure comparability and enable to reproduce the results with benchANT or any other manual or semi-automated approach.

For the generation comprehensive data sets, we build upon the capabilities of the Mowgli framework [2] which are further extend in the level of detail by requirements coming from benchANT [8] customers such as DBMS and cloud providers or enterprises with performance engineering teams. The transparency and reproducibility of the results are verified by three DBMS providers[5].

In the following subsections, we describe the required types of data to ensure comparability and reproducibility and discuss how we addressed this in the ODPR data set. The full data set of the ODPR is available on GitHub[6]

3.1 Benchmark Specification

The benchmark specification needs to provide a model of all configurable parameters for the benchmark execution, covering the workload but also the supported cloud and DBMS configurations.

In the ODPR data set we apply the concept of JSON-based evaluation scenarios introduced by Mowgli [2]. An evaluation scenario contains all configurable parameters to execute the benchmark via an automation framework or manually. Commonly modified parameters are the VM instance type, DBMS type, cluster size, workload intensity and workload request distribution.

3.2 Cloud & DBMS Data

Besides the static benchmark specification, it is also required to collect the dynamic runtime metadata on cloud provider, resource and DBMS level in order to enable in-depth and also time-bound analysis of the results. Deviations on the cloud and DBMS level can occur even with an identical benchmark specification, depending on the benchmark execution date, *e.g.* updated kernel version or unreachable DBMS nodes in large-scale clusters.

For this purpose, benchANT extends the metadata collection of Mowgli by additionally collecting cloud api responses, operating system details, DBMS configuration files and DBMS cluster states.

3.3 Performance Data

Performance data needs to be provided as time-series for each performance metric per benchmark execution and configuration for in-depth analysis and aggregated to enable higher level comparisons.

The ODPR data sets provides the raw performance metrics of each applied benchmark suite as well as aggregated metrics as shown in the ranking.

3.4 Monitoring Data

In order to enable the identification of potential benchmark bottlenecks such as an overload benchmark instance or insufficient network bandwidth, it is mandatory to monitor the system utilization of all components during the benchmark execution and include this data into the benchmark results.

Mowgli already providing monitoring of the involved components, benchANT extends these capabilities with respect to data export and processing.

[5]The DBMS providers can not be named directly due to a non disclosure agreement

[6]`https://github.com/benchANT/database-ranking`

3.5 Execution Data

For reproducible results, it is required to provide a timed execution log of the individual tasks that have carried out to execute the benchmark, such as allocate the cloud resources, deploy the DBMS or execute the workload phases.

For the orchestration of these tasks, benchANT relies internally on Apache Airflow[7] and provides fine-grained log files to enable the traceability of the benchmark execution.

4 Future Research Directions

The following research directions that address the automated operation of DBMS in the cloud as well as the automated performance optimization can be further developed by building on such comprehensive data sets. In particular, we envision the following research directions as potential beneficiaries:

4.1 DBMS Performance Modelling

Building DBMS performance models for predicting DBMS performance is a complex challenge due to the various impact factors on (distributed) DBMS and resource level where the cloud adds another level of complexity. In consequence, large scale measurement series are required that are represented as comprehensive data sets. The Baloo framework [4] presents a first approach for building such performance models for distributed DBMS on a single cloud by building upon data sets generated by Mowgli [2]. The extended data sets by benchANT enable to advance the performance models in a multi-cloud context.

4.2 DBMS Auto-Tuning

Auto-tuning of configuration parameters for distributed systems and for DBMS in particular is an ongoing research directions with promising results for distributed systems [10] and relational DBMS [6]. In order to expand this research towards distributed NoSQL DBMS that are operated in the cloud, the outlined data set structure can be a valuable input to the auto-tuning alogirthms.

4.3 Cloud & DBMS Decision Support

The research direction of decision support systems guide users in finding the optimal cloud resources and cloud-hosted DBMS for their use case based on functional and non-functional features. Decision support systems such as MiPACE [1] or Hathi [3] apply multi-criteria decision making (MCDM) algorithms combine non-functional feature measurements such as performance or scalability with costs and additional functional features into a global score. In consequence, comprehensive performance data sets are a valuable input for building advanced decision support systems for performance focused cloud applications.

[7] `https://airflow.apache.org/`

5 Conclusion

DBMS benchmarking is a continuous and valuable method to drive the development of DBMS. But ensuring transparent and reproducible results is still a challenge due to the manifold impact factors. We report our experiences in building an open database performance ranking that ensure transparency and reproducibility by providing a comprehensive data set. Moreover, we outline how such data sets can be applied in different performance-related DBMS research directions from which we will apply the data sets to improve DBMS configuration auto-tuning.

References

[1] J. Araujo et al. "Decision making in cloud environments: an approach based on multiple-criteria decision analysis and stochastic models". In: *Journal of Cloud Computing* 7.1 (2018).

[2] D. Seybold et al. "Mowgli: Finding Your Way in the DBMS Jungle". In: *ACM/SPEC ICPE*. Mumbai, India: ACM, 2019, pp. 321–332.

[3] J. Domaschka, S. Volpert, and D. Seybold. "Hathi: An MCDM-based Approach to Capacity Planning for Cloud-hosted DBMS". In: *2020 IEEE/ACM 13th UCC*. 2020, pp. 143–154.

[4] J. Grohmann et al. "Baloo: Measuring and Modeling the Performance Configurations of Distributed DBMS". In: *MASCOTS*. 2020.

[5] A. V. Papadopoulos et al. "Methodological Principles for Reproducible Performance Evaluation in Cloud Computing". In: *IEEE Transactions on Software Engineering* 47.8 (2021).

[6] A. Pavlo et al. "Make Your Database System Dream of Electric Sheep: Towards Self-Driving Operation". In: *Proc. VLDB Endow.* 14.12 (2021).

[7] D. Seybold. "An automation-based approach for reproducible evaluations of distributed DBMS on elastic infrastructures". en. PhD thesis. 2021.

[8] D. Seybold and J. Domaschka. "Benchmarking-as-a-Service for Cloud-Hosted DBMS". In: *Middleware: Demos and Posters*. Middleware '22. Virtual Event, Canada: ACM, 2021, pp. 12–13.

[9] D. Abadi et al. "The Seattle Report on Database Research". In: *Commun. ACM* 65.8 (July 2022), pp. 72–79.

[10] M. S. Iqbal et al. "Unicorn: Reasoning about Configurable System Performance through the Lens of Causality". In: *EuroSys*. ACM, 2022.

[11] M. Leznik et al. "Same, Same, but Dissimilar: Exploring Measurements for Workload Time-Series Similarity". In: *ACM/SPEC ICPE*. ACM, 2022.