Releem

Workload-Based Tuning MySQL

How to use workload mapping for MySQL tuning

Kochetov Dmitrii — CTO, Releem

Releem

Introduction

High performance DBMS, which are a critical component of applications, is a difficult task.

DBMS performance targets:
- *Throughput*
- *Latency*

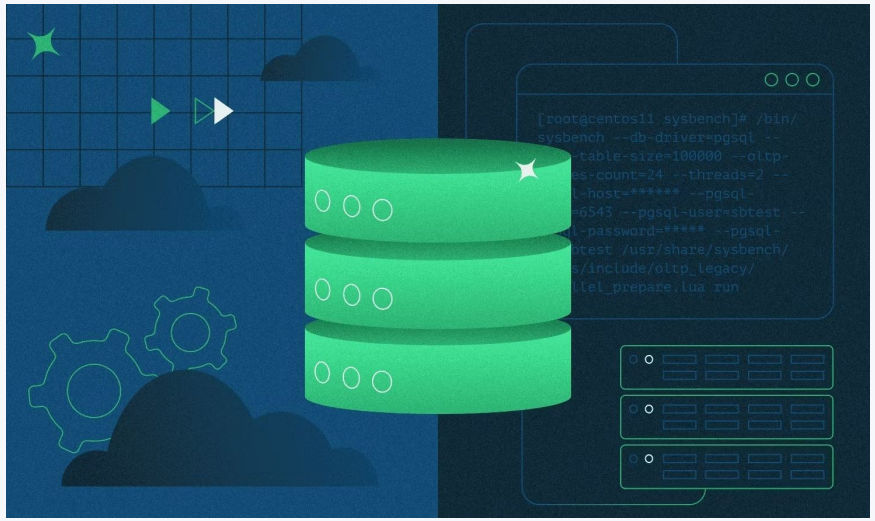

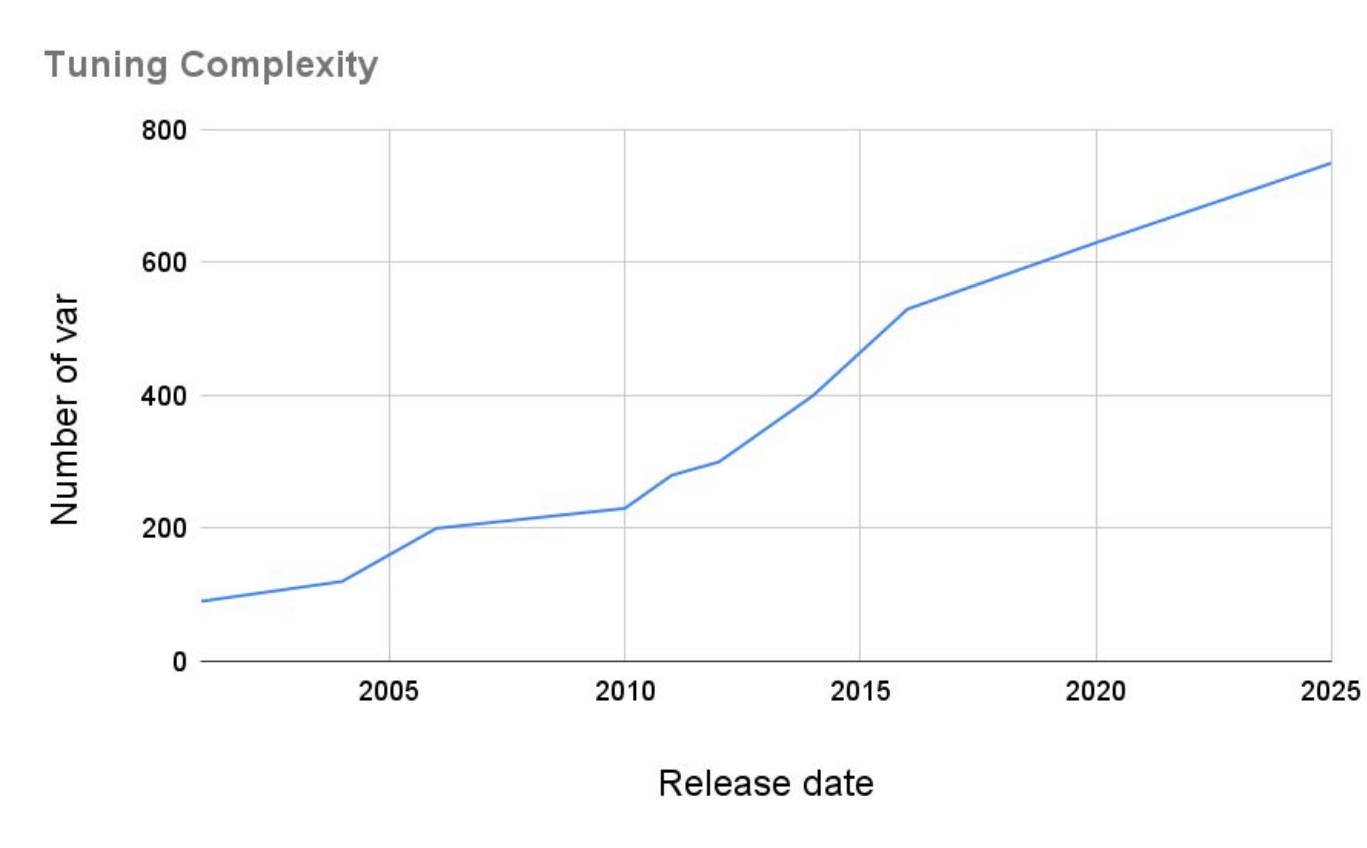

PERCONA

Releem

Problem

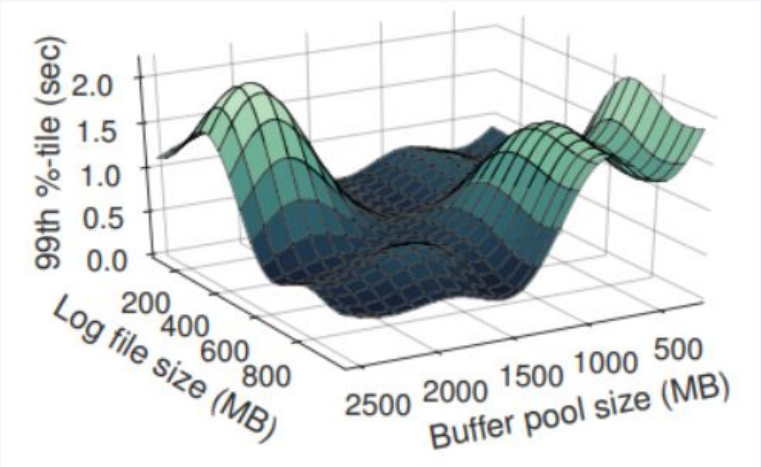

1 Dependencies

Requires Installation of Complicated Monitoring Solutions

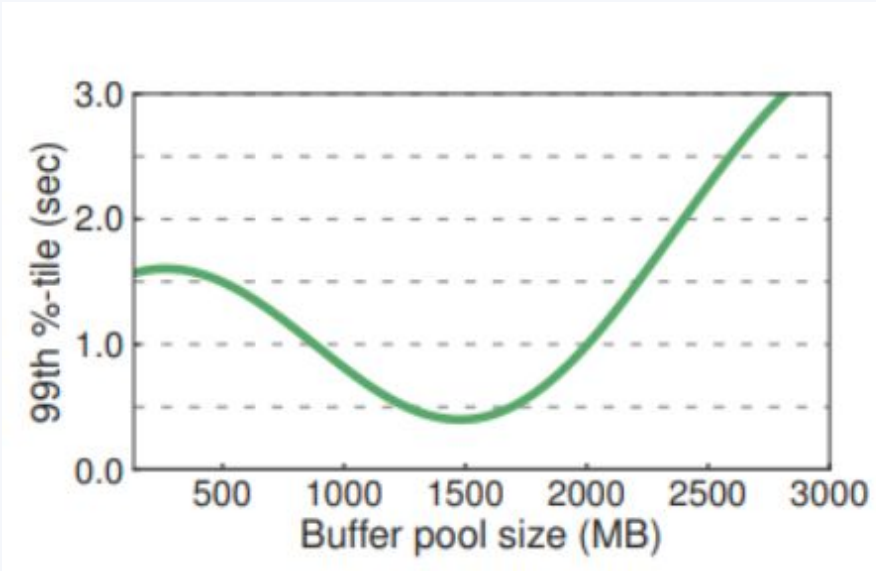

2 Continuous Settings

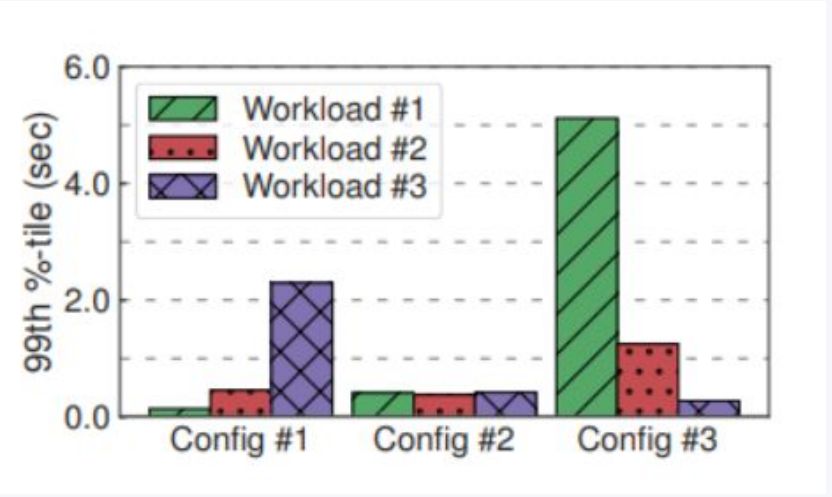

3 Non-Reusable Configurations

Needs Experimentation to Get Experience

PERCONA

Releem

Goal

How does MySQL Tuning impact Application Performance?

Application	Response Time (Latency)	CPU Utilization	Queries per Seconds
WordPress WooCommerce	-63%	-37%	+106%
Drupal Commerce Kickstart	-97%	-73%	+268%
Aimeos Laravel	-42%	-86%	+291%

Tuned MariaDB configuration compared to default MariaDB configuration

Benefits of MySQL Performance Tuning

- Enhanced Database Efficiency
- Improved Query Response Times
- Reduced Resource Usage
- Enhanced User Experience
- Reduced costs on servers

PERCONA

Releem

System Overview

- An overview of the components in the system.
- The Agent connects to the DBMS and collects information about the performance of the system.
- This information is then sent to the tuning platform where it is stored in its storage. It then builds models that are used to analyze, mapping and select an optimal configuration for the DBMS.

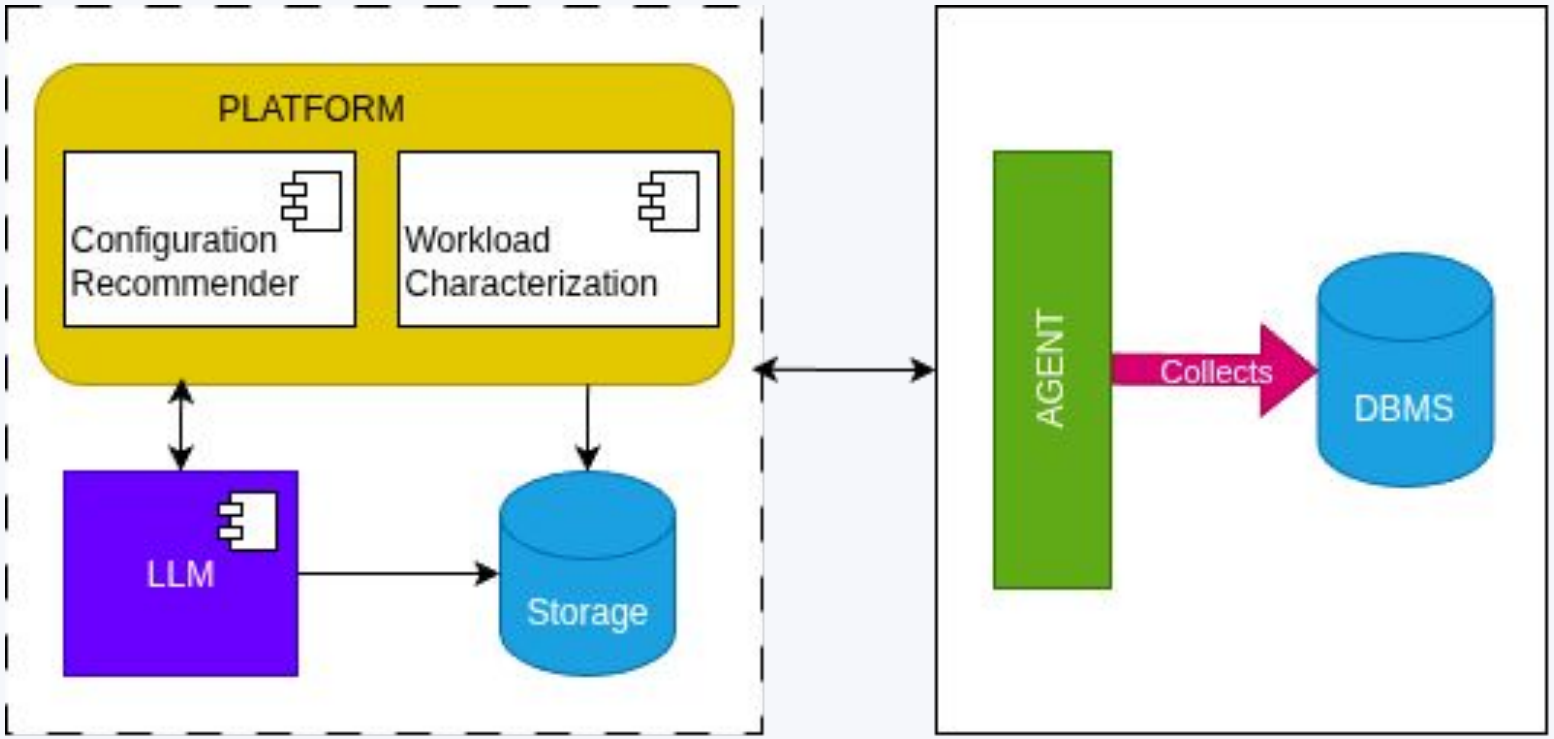

PERCONA

Releem

Workload Characterization

Discover for a model that best represents the distinctive aspects of the target workload

- The first is to analyze the target workload at the logical level.
- A second approach is to use the DBMS's internal runtime metrics to characterize how a workload behaves.

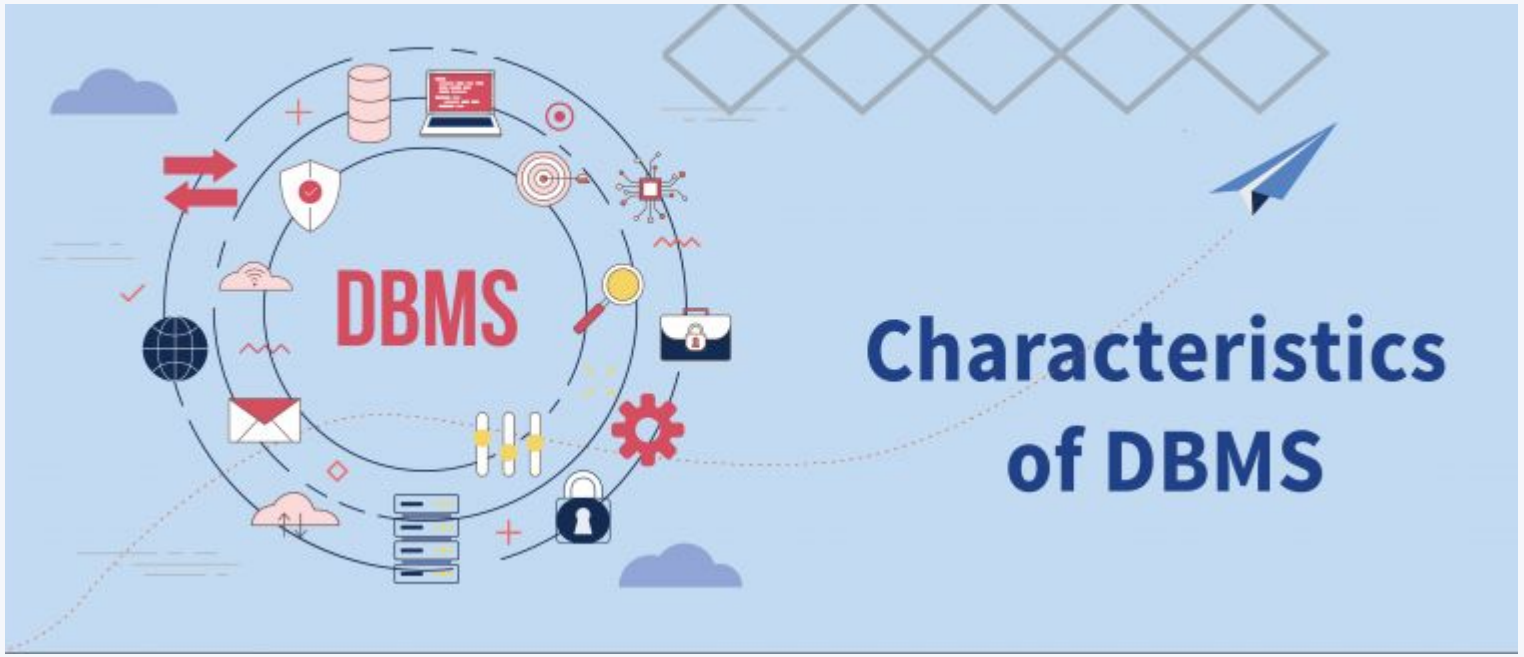

PERCONA

Releem

Selection of non-redundant metrics

Redundant DBMS metrics occur for two reasons:

- The first are ones that provide different granularities for the exact same metric in the system.
- The second type of redundant metrics are ones that represent independent components of the DBMS but whose values are strongly correlated.

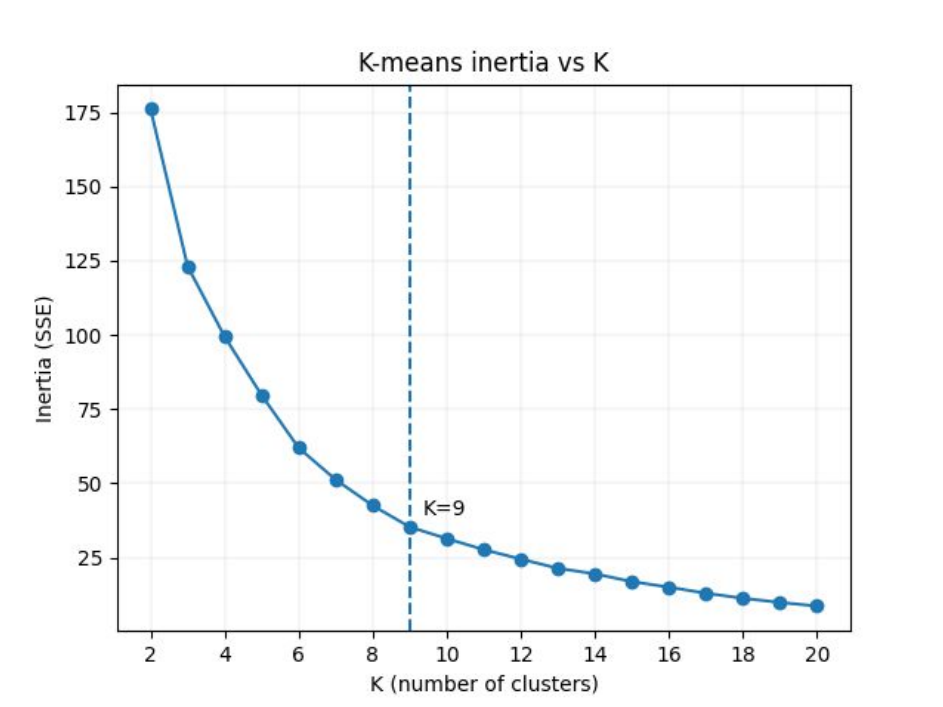

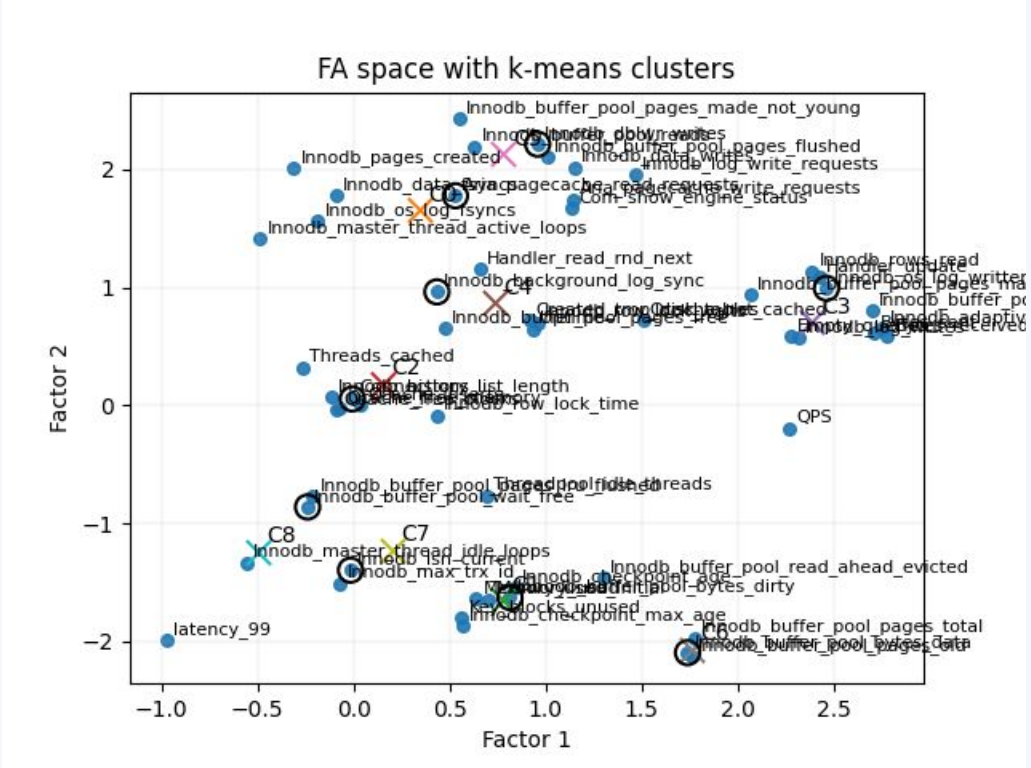

Releem

Workload Mapping

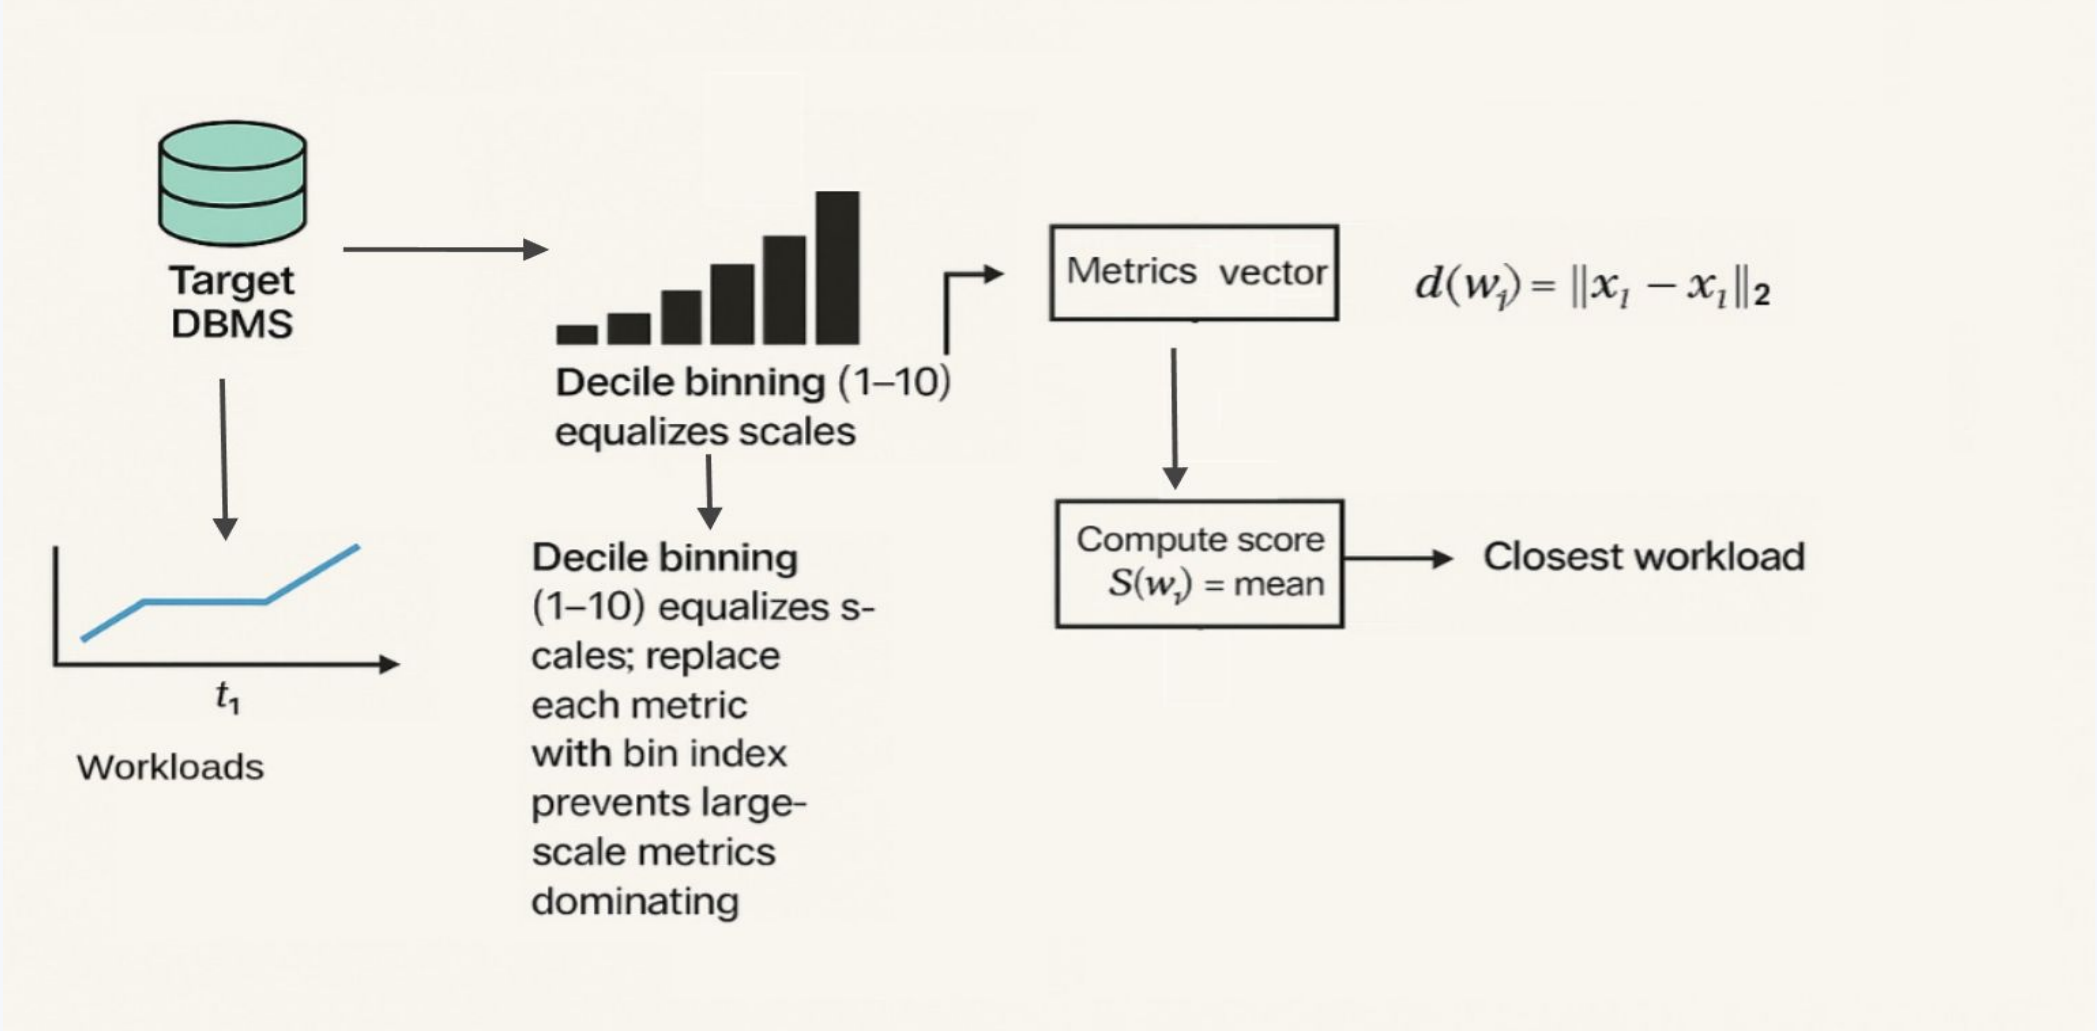

PERCONA

Releem

LLM Configuration Tuning
How it works?

- The Goal of optimization (Latency, Throughput, Resource usage)
- Reinforcement learning algorithms
- Workload Mapping

PERCONA

Releem

LLM Configuration Tuning. Identifying Important Variables

Algorithm LLM-based Variables Selection

Input: Var Set K; LLM F; DBMS D; Workload W; Tuning Lake L.
Output: Target Var Set T.
Configurable Var Set C ← $FILTER$(K);
// Filter out variables that are related to debugging, security and path-settin
System Level Selection: Cs ← F (C, D);
Workload Level Selection: Cw ← F (C,W);
Query Level Selection: Cq ← ∅;
for query qi in W **do**
    Ei ← $EXECUTE$(D, qi);
    // Get execution plan for query qi from D
    Cqi ← F (C, E);
    Cq ← Cq ∪ Cqi;
end
Variable Level Selection: Target Var Set T ← F (L, Cs ∪ Cw ∪ Cq);
return T;

PERCONA

Releem

LLM Configuration Tuning. Configuration Recommender.

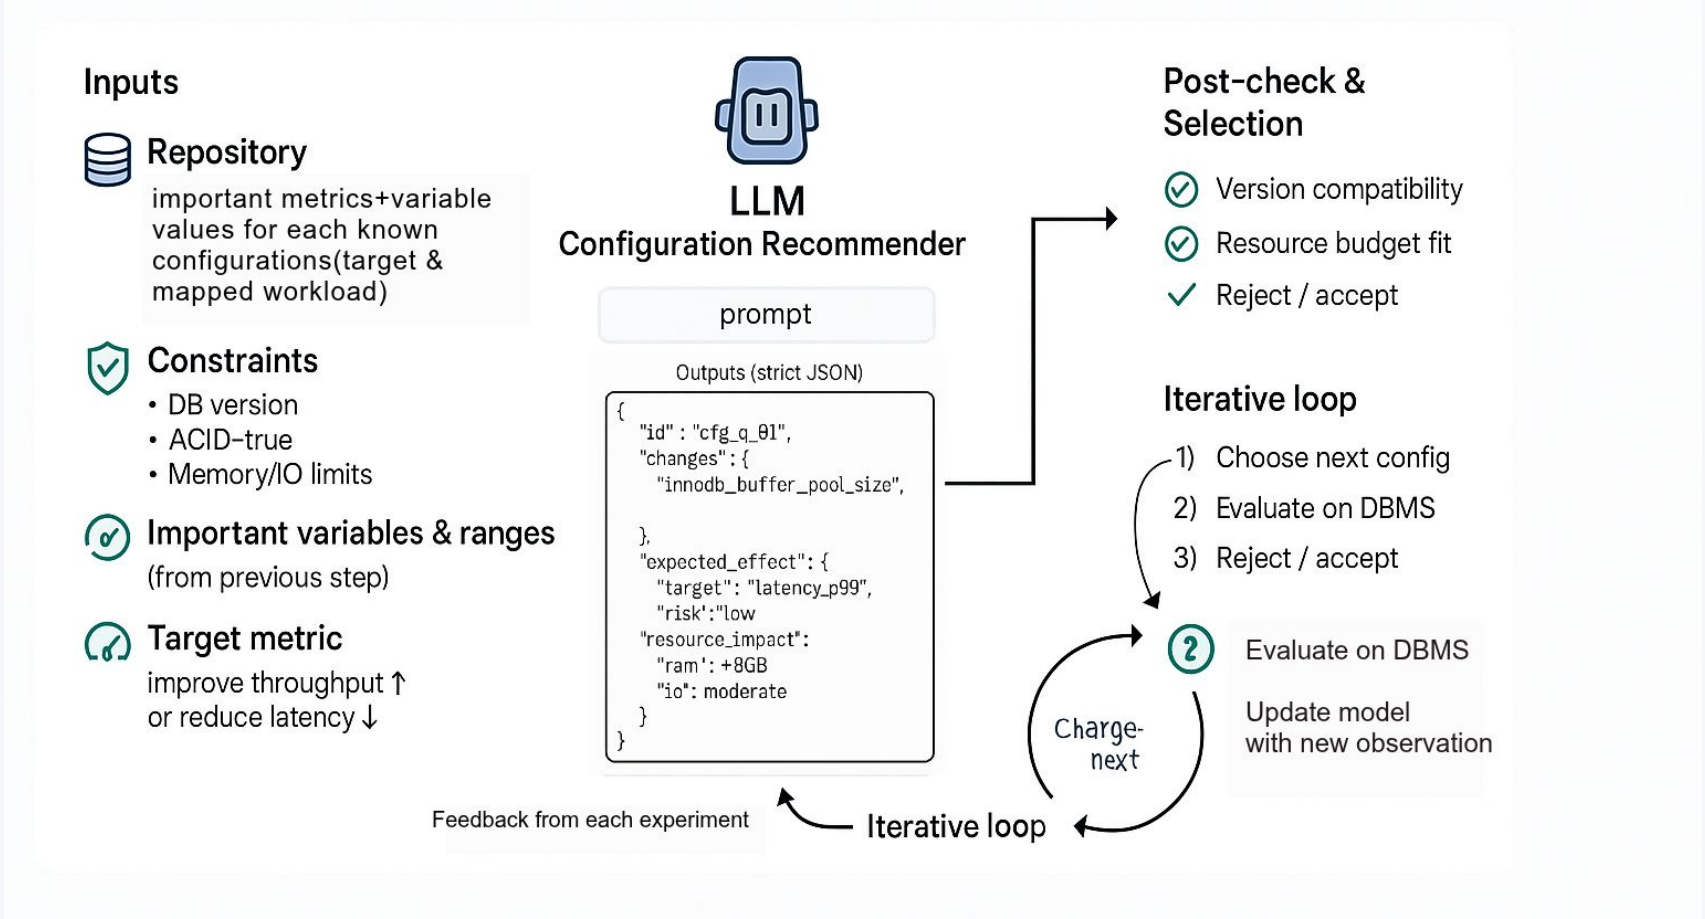

PERCONA

Releem

Training Data Collection

For these experiments, we use workloads from the BenchBase testbed that differ in complexity and system demands:

- YCSB: The Yahoo! Cloud Serving Benchmark (YCSB) is a collection of micro-benchmarks that represent data management applications whose workload is simple but requires high scalability
- TPC-C: This is the current industry standard for evaluating the performance of OLTP systems
- TPC-H: This benchmark illustrates decision support systems that examine large volumes of data, execute queries with a high degree of complexity, and give answers to critical business questions.

	ReadRecord	InsertRecord	ScanRecord	UpdateRecord	DeleteRecord	ReadModifyWriteRecord
YCSB Workload №1	100	0	0	0	0	0
YCSB Workload №2	50	0	0	50	0	0
YCSB Workload №3	70	10	10	10	0	0
YCSB Workload №4	10	20	10	30	10	20
.....						
YCSB Workload №15	33	34	0	33	0	0

Releem

Study results

PERCONA

Releem

Thank you for your attention.

Dmitrii Kochetov
CTO, ***Releem***

PERCONA