

Transparent Data Tiering for Mixed Workload IMDBs

Status Quo - Next Steps
2014-07-01

Martin Boissier
Carsten Meyer

Agenda

- Motivation / Context
- Dynamic Aging Rules
- Bachelorprojekt 2013
- Masterproject 2014
- Next Steps

Aging != Archiving

Data Aging:

- Performance and data volume improvements
- Data schema not changed
- „aged“ data must remain accessible (application transparent)
- Brand new!

Data Archiving:

- Legal and compliance requirements
- Transformation into archive format
- Explicit data access (data not visible to standard queries)
- > 20 years

A stone-old Challenge in a new DB-Era

*„Highly relevant (hot) data **should not be** handled the same way as irrelevant (cold) data*“*

- Buffering of disk-pages in main memory
- Pagelist management
- Page-optimization using indices

*e.g. DB2 requires 5% main memory for approx. 98% data accesses

Hot vs. Cold

HANA for real-time, mixed workload on transactional data

- Today, everything is (must be) in main memory
- Irrelevant (cold) data blocks precious resources
- No solutions for mixed workloads yet

Goal:

- A. Identify hot data and pin it in main memory.
- B. Avoid access on cold partition whenever possible to improve **memory utilization** and query **performance**

Mixed Workload Awareness

- `SELECT * FROM acc WHERE id = 8987789`
- `SELECT * FROM acc WHERE date >= today-30d AND status = 'open'`
- `SELECT sum(rev) FROM acc WHERE year >= 2009 AND customer = 149003`

The workload and data relevance is hard to anticipate at design-time in a mixed workload enabled system.

Requirements:


- Flexibility, ideally full application transparency
- Horizontal and vertical relevance classification
- Deterministic data access (hot or hot / cold) during query execution

Dynamic Aging Rules

1. **Sample** (trace) data -selection and -projection (**workload**)
2. Periodically analyze and rank these statistics > **derive rules**
3. Use rules **for classification** (aging) and **access** (query processing)


Aging Rules

- $\bigcup \pi(\sigma(R))$ - union of queries having different projections / selections, defining „hot“ data
- Rule1: SELECT rev, year, customer FROM acc WHERE year >= 2009 AND customer in (149003, 16765, ...)
- Rule2: SELECT * FROM acc WHERE status = 'open'



Hot & Cold in HYRISE

- Horizontal partitioning per column
- Different storage allocations
- Aging check before query exec.
- Hot-only operator (scan)



So far...

Idea: Workload-aware, dynamic aging rules, for:


- Data **classification** and
- Data **access** (avoid cold data access)

1. Trace of a real-world, productive SAP ERP system (100.000 tables, ~2.5 TB data in memory, ~**1.5 billion queries each day**)

2. Workload parsing and extraction of selection characteristics


3. Hot and Cold Partitions in HYRISE

BP Architecture




Demo

Explorative Workload Analyzer



Top 10 Query Templates


Template ID	Statement	Executions	Share	Type
95527	SELECT TOP 1 "TASKTYPE" FROM "CATSDB" WHERE "MANDT" = ? AND "PERNR" = ? AND "WORKDATE" = ? AND "TASKCOUNTER" = ?	596,788	48.80%	Partial Key Select
95589	SELECT * FROM "CATSDB" WHERE "MANDT" = ? AND "PERNR" = ? AND "ZZOBJNR" = ? AND NOT ("CATSQQUANT" = ?) AND ...	473,708	38.60%	Partial Key Select
95581	SELECT TOP 1 * FROM "CATSDB" WHERE "MANDT" = ? AND "COUNTER" = ?	325,308	26.40%	Key Select
95588	SELECT TOP 1 "TASKTYPE" FROM "CATSDB" WHERE "MANDT" = ? AND "COUNTER" = ?	126,130	10.40%	Key Select
95589	SELECT TOP 1 "ZZSUBTYPE", "TASKTYPE" FROM "CATSDB" WHERE "MANDT" = ? AND "TASKCOUNTER" = ? AND "PERNR" = ? ...	104,340	8.50%	Partial Key Select
95588	SELECT * FROM "CATSDB" WHERE "MANDT" = ? AND "PERNR" = ? AND "WORKDATE" = ? AND "WORKDATE" = ? AND "STATU" = ?	72,424	5.80%	Partial Key Select
95588	SELECT "COUNTER", "PERNR", "WORKDATE", "WABNR", "LGART", "LGORNR", "STATUS", "REFCOUNTER", "BELNR", "CATS..."	65,876	5.30%	Partial Key Select
95584	SELECT * FROM "CATSDB" WHERE ("MANDT" = ? AND "PERNR" = ? AND "STATUS" = ? AND "WORKDATE" BETWEEN ? AND ? A...	44,840	3.60%	Partial Key Select
95545	SELECT * FROM "CATSDB" WHERE "MANDT" = ? AND "COUNTER" IN (?) AND NOT ("ZZ_LGORTNR" = ?)	36,168	2.90%	Key Select
95588	SELECT * FROM "CATSDB" WHERE "MANDT" = ? AND "PERNR" = ? AND "WORKDATE" = ? AND "WORKDATE" = ? AND "STATU" = ?	22,482	1.80%	Partial Key Select




Workload Replay



Visualize **access** frequency and **pattern** for each table **record**

See **hot** and **cold** data







HPI Data Tiering - Status Quo / Next Steps




HPI Data Tiering - Status Quo / Next Steps




HPI Data Tiering - Status Quo / Next Steps



HPI Data Tiering - Status Quo / Next Steps



HPI Data Tiering - Status Quo / Next Steps



HPI Data Tiering - Status Quo / Next Steps

...The Master Project

- Use workload characteristics to derive optimized aging rules
- Implement a SAPphire-ready prototype* to validate dynamic aging for IMDBs
- Evaluate rules: For a real-world enterprise workload

Aspects

- Identify Aging Rules (relevant aging columns and column projections)
- Characterize Aging Rules (functional dependencies and time-correlation of columns)
- Optimized sampling for Aging Rules (reduce statistic overhead)
- Evaluate Aging Rules

For a productive Enterprise Workload