

# Build your own Database

Week 7

# Agenda

---

- How was Sprint 3?
- Null Values
- Logistics
- How to work with the Code Base
- Topic Assignments
- First Meeting for Group Project

# NULL values in SQL

---

- NULL is used to represent absent values
- It is a state/marker not a value
- DBMS dependent behavior in many cases
  - Arithmetic operations involving NULL, usually result in NULL:  $\text{NULL} * 17 \rightarrow \text{NULL}$ , but what about  $\text{NULL} / 0$ ?
  - String concatenations involving NULL, result in NULL

# NULL values in SQL

- `SELECT 17 = NULL` → ?
- SQL offers three logical results: True, False, Unknown → Three-valued Logic (3VL)

a	b	a AND b	a OR b	a = b
True	True	True	True	True
True	False	False	True	False
True	Unknown	Unknown	True	Unknown
False	True	False	True	False
False	False	False	False	True
False	Unknown	False	Unknown	Unknown
Unknown	True	Unknown	True	Unknown
Unknown	False	False	Unknown	Unknown
Unknown	Unknown	Unknown	Unknown	Unknown

# NULL values in SQL

## Customer

c_id	firstname	lastname
0	NULL	Zayer
1	Alex	Geier
2	Frank	Meier

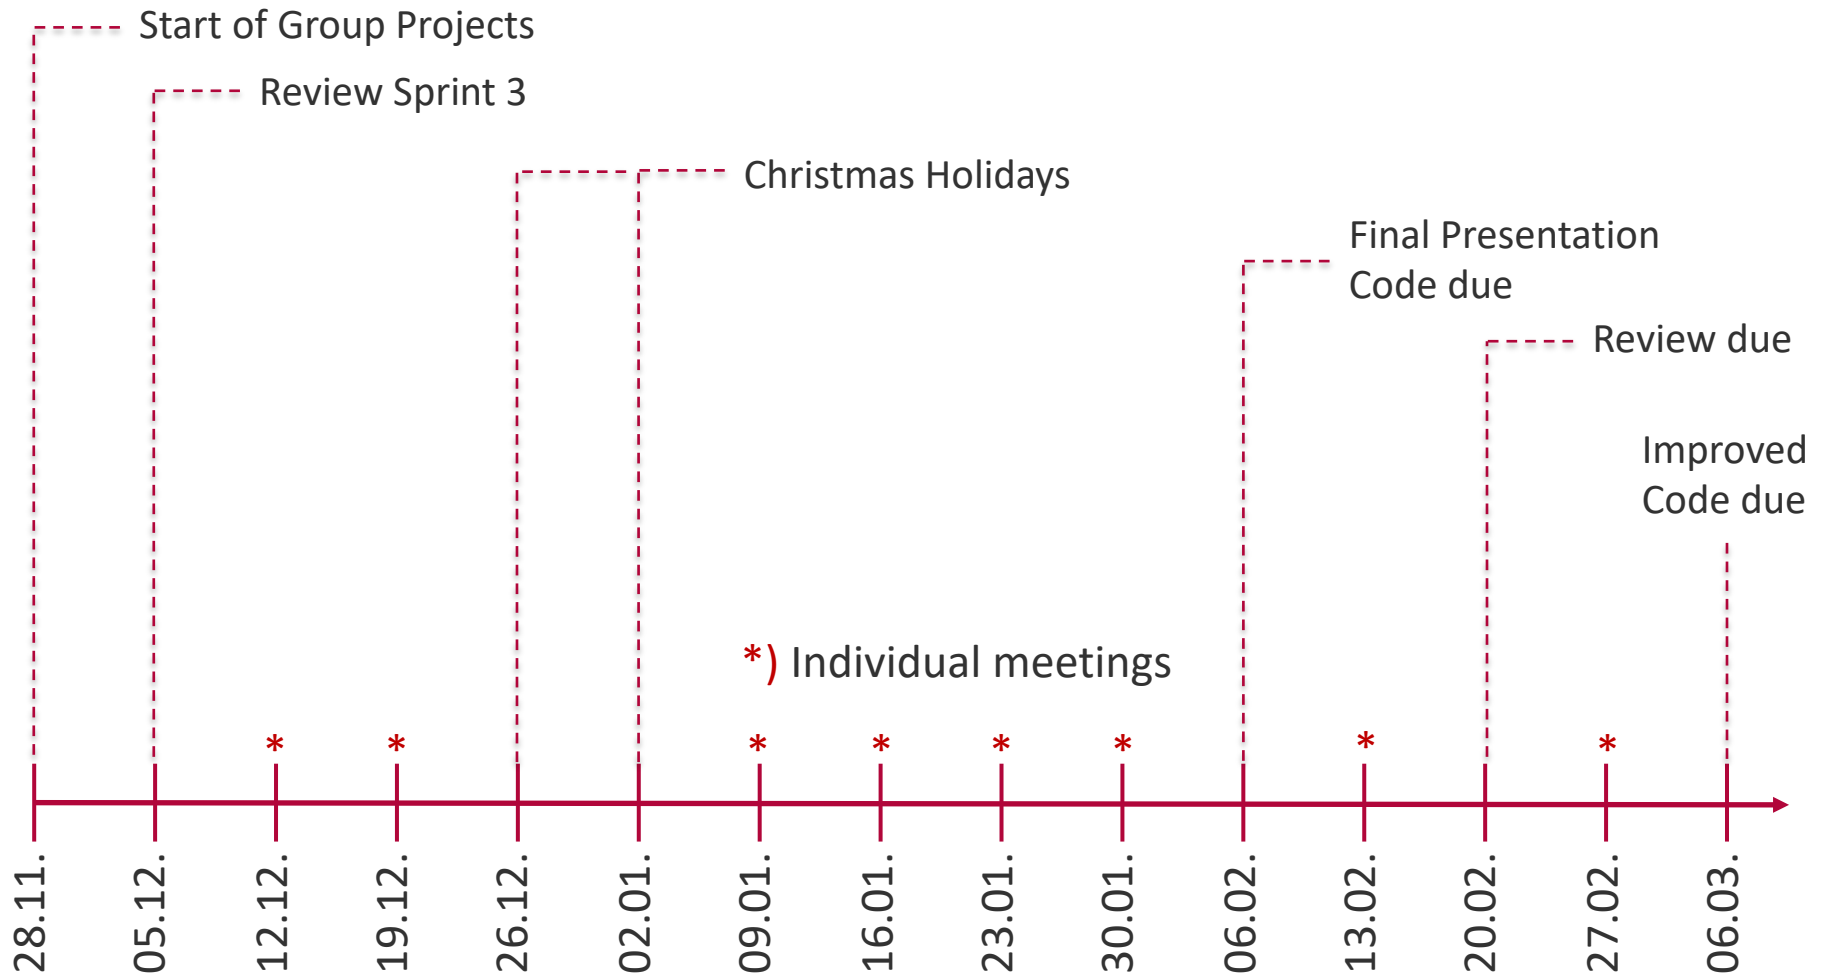
- `SELECT * FROM customer WHERE firstname = NULL; → ?`
- `SELECT * FROM customer WHERE firstname <> Alex; → ?`
- Rows for which predicates evaluate to Unknown are treated like rows that evaluate to False
- All standard comparison operators return Unknown when comparing NULL
- Thus, SQL offers `IS NULL` and `IS NOT NULL`

# NULL values in Hyrise

We need a representation for NULL for Value, Reference and Dictionary columns:

```
// Value Segments  
  
static const auto NULL_VALUE = AllTypeVariant{};  
  
// Reference Segments  
  
constexpr ChunkOffset INVALID_CHUNK_OFFSET{std::numeric_limits<ChunkOffset>::max()};  
  
const RowID NULL_ROW_ID = RowID{ChunkID{0u}, INVALID_CHUNK_OFFSET};  
  
// NULL ValueIDs for Dictionary Segments  
  
constexpr ValueID NULL_VALUE_ID{std::numeric_limits<ValueID::base_type>::max()};  
  
segment.unique_values_count()
```

# Logistics



\*) Individual meetings

# Date of Final Presentation

---

- Final day of instruction is on 06.02.
- We would like to have the final presentation on that day, but ~11 minutes per group is not enough
- Three options
  - Super Wednesday (06.02., 09:15 – 12:30)
  - Find a second slot in that week
  - Take the Wednesday before (our least favorite)



# Setting up Hyrise

---

1. **Fork** the Hyrise Repo
2. Run `./install.sh`
3. Setup a build folder
  - `mkdir build; cd build; cmake ..; make -jX`
  - Important cmake flags:
    - `-DCMAKE_CXX_COMPILER_LAUNCHER=ccache`
    - `-DCMAKE_BUILD_TYPE=Release` (or `RelWithDebInfo`)
    - `-DENABLE_JIT_SUPPORT=OFF`

# Hyrise Code Base

GitHub, Inc. (US) | <https://github.com/hyrise/hyrise> Search

Pull requests Issues Marketplace Explore

hyrise / hyrise Unwatch 33 Unstar 20 Fork 32

Code Issues 74 Pull requests 7 Wiki Insights Settings

Hyrise is a research in-memory database. <https://hpi.de/plattner/projects/hyri...> Edit

database in-memory-database cpp sql Manage topics

1,163 commits 75 branches 0 releases 33 contributors MIT

Branch: master New pull request Create new file Upload files Find file Clone or download

mrks Assert pos list size (#1327)

- cmake -Walmost-everything (#1148)
- scripts Fix lint.sh (#1248)
- src Assert pos list size (#1327)
- third\_party Update cxxopts, remove hack (#1315)

Clone with HTTPS Use SSH  
Use Git or checkout with SVN using the web URL.  
<https://github.com/hyrise/hyrise.git>  
Open in Desktop Download ZIP

# Demo

---

## Run Console

- `cd ..; ./build/hyriseConsole`
- `generate_tpch 0.1`
- `SELECT * FROM customer JOIN orders ON c_custkey = o_custkey WHERE o_orderpriority = '5-LOW' LIMIT 1`
- `visualize`

## Run TPCH Benchmark

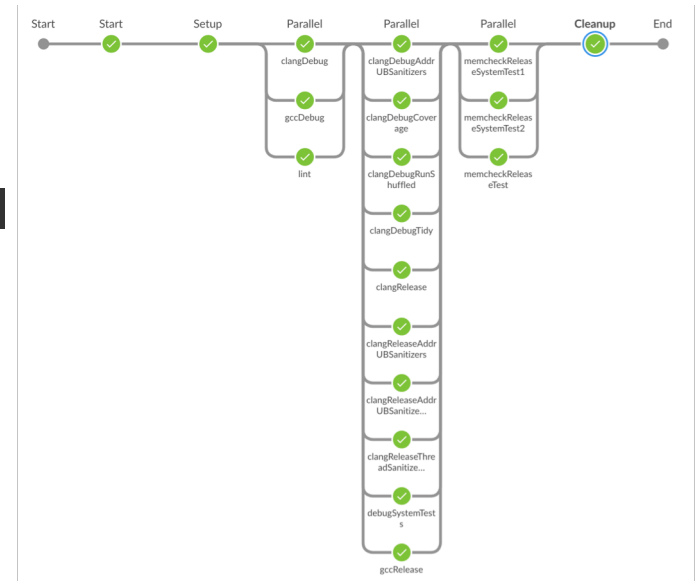
- `./build-release/hyriseBenchmarkTPCH -s 0.1 -v --visualize`

## Compare results

- `./scripts/compare_benchmarks.py old.json new.json`

# Continuous Integration

- <https://hyrise-ci.epic-hpi.de/>
- Once you open your first PR, we will assign Hyrise team membership
- Only then will the CI start to build your code
- Once your code is stable and reviewed, you can add the FullCI tag, which enables additional CI verification



# Pull Requests

---

- Please open PRs from early on
  - Early Feedback
  - Fewer Conflicts
  - Less Blocks
- The initial review will be done by *the other group*

# Fortnightly

---

- Every other Wednesday (starting next week), the Hyrise developers meet to discuss current issues and architecture changes
  - 13:30, Konfi
- This is a good chance of seeing what is happening in the other groups and what is happening with Hyrise in general

# Support

---

- If you have any questions, please don't wait for a week
- You can use Piazza for questions or email us
- If anything looks odd or is hard to understand, please let us know

# Hyrise Paper

- Our EDBT paper discusses some of the design decisions and the overall architecture
- Some parts have been covered in the lecture, others (e.g., JIT) should be new
- We will upload it to Piazza

## Hyrise Re-engineered: An Extensible Database System for Research in Relational In-Memory Data Management

Markus Dreseler, Jan Kossmann, Martin Boissier, Stefan Klauck,  
Matthias Uflacker, Hasso Plattner  
Hasso Plattner Institute  
Potsdam, Germany  
firstname.lastname@hpi.de

### ABSTRACT

Research in data management profits when the performance evaluation is based not only on single components in isolation, but uses an actual DBMS end-to-end. Facilitating the integration and benchmarking of new concepts within a DBMS requires a simple setup process, well-documented code, and the possibility to execute both standard and custom benchmarks without tedious preparation. Fulfilling these requirements also makes it easy to reproduce the results later on.

The relational open-source database Hyrise (VLDB, 2010) was presented to make the case for hybrid row- and column-format data storage. Since then, it has evolved from a single-purpose research DBMS towards a platform for various projects, including research in the areas of indexing, data partitioning, and non-volatile memory. With a growing diversity of topics, we have found that the original code base grew to a point where new experimentation was made unnecessarily difficult. Over the last two years, we have rewritten Hyrise from scratch, focusing on building an extensible multi-purpose research DBMS that can serve as an easy-to-extend platform for a variety of experiments and prototyping in database research.

In this paper, we discuss how our learnings from the previous version of Hyrise have influenced our rewrite. We describe the new architecture of Hyrise and highlight the main components. We then show how our extensible plugin architecture facilitates research on diverse DBMS-related aspects without compromising the architectural tidiness of the code. In a first performance evaluation, we show that the execution time of most TPC-H queries is competitive to that of other research databases.

### 1 INTRODUCTION

Hyrise was first presented in 2010 [19] to introduce the concept of hybrid row- and column-based data layouts for in-memory databases. Since then, several other research efforts have used Hyrise as a basis for orthogonal research topics. This includes work on data tiering [7], secondary indexes [16], multi-version concurrency control [42], different replication schemes [43], and non-volatile memories for instant database recovery [44]. Over the years, the uncontrolled growth of code and functionality became an impediment for future experiments. We have identified four major factors leading to this situation:

- Data layout abstractions were resolved at runtime and incurred costs that sometimes had a disproportional overhead.
- Prototypical components have been implemented to work in isolation, but did not interact well with other components.

- The lack of SQL support required query plans to be written by hand and made executing standard benchmarks tedious.
- Accumulated technical debt made it difficult to understand the code base and to integrate new features.

For these reasons, we have rewritten Hyrise from scratch and incorporated the lessons learned. We redesigned the architecture to provide a stable and easy to use basis for holistic evaluations of new data management concepts. Hyrise now allows researchers to embed new concepts in a proper DBMS and evaluate performance end to end, instead of implementing and benchmarking them in isolation. At the same time, we allow most components to be selectively enabled or disabled. This way, researchers can exclude unrelated components and perform isolated measurements. For example, when developing a new join implementation, they can bypass the network layer or to disable concurrency control.

In this paper, we describe the new architecture of Hyrise and how our prior learnings have led to a maintainable and comprehensible database for researching concepts in relational in-memory data management (Section 2). Furthermore, we present a plugin concept that allows to test-drive different optimizations without having to modify the core DBMS (Section 3). We compare Hyrise to other database engines, show which approaches are similar, and highlight key differences (Section 4). Finally, we evaluate the new version and show that its performance is competitive (Section 5).

### 1.1 Motivation and Lessons Learned

The redesign of Hyrise reflects our past experiences in developing, maintaining, and using a DBMS for research purposes. We motivate three important design decisions.

*Decoupling of Operators and Storage Layouts.* The previous version of Hyrise was designed with a high level of flexibility in the storage layout model: each table could consist of an arbitrary number of containers, which could either hold data (in uncompressed or compressed, mutable or immutable forms) or other containers with varying horizontal and vertical spans. In consequence, each operator had to be implemented in a way where it could deal with all possible combinations of storage containers. This made the process of adding new operators cumbersome and led to a system where some operators made undocumented assumptions about the data layout (e.g., that all partitions used the same encoding type). Instead of relying on operators to properly process data structures with varying memory layouts, Hyrise now follows an iterator-based approach. By accessing data through iterators, the implementation of new operators is decoupled from the implementation of new data storage concepts without compromising the flexibility. Operators can implement custom specializations for specific iterators, but execution falls back to the default iterator if no implementation exists. The iterator abstraction is explained in Section 2.3.

© 2019 Copyright held by the owner/author(s). Published in Proceedings of the 22nd International Conference on Extending Database Technology (EDBT), March 26-29, 2019, ISBN XXX-XXXXXX-XXX-X on OpenProceedings.org. Distribution of this paper is permitted under the terms of the Creative Commons license CC-by-nc-nd 4.0.



# Topic Assignments

---

- Gruppe 1 (EK, AL, AR): Aggregate (JK) here
- Gruppe 2 (CF, HR, JS): Sorted Segments (SK) Glaskasten
- Gruppe 3 (MK, TL, SS): Between (JK) here
- Gruppe 4 (PF, JJ, DS): Subselect-to-Join (MB) Konfi
- Gruppe 5 (MF, KO, MW): Multi-Predicate Joins (SK) Glaskasten
- Gruppe 6 (PB, MH, NT): Primary Keys (MD) V-1.16
- Gruppe 7 (FE, DJ, JM): Physical MVCC Delete (MD) V-1.16
- Gruppe 8 (JE, LI, LL): Heavy Compression (MB) Konfi