



IDUG

2023 NA **Db2** Tech Conference

Machine Learning Optimization for production use in the next version of Db2

Calisto Zuzarte, IBM

Session Code: C08

Philadelphia

Agenda

- Motivation
- Tech Preview
- Plan for Productization



Motivation

Motivation – Challenges, Benefits and Goals (1|4)

- **Optimizer Challenges**

- Stability
- Tuning effort
- Development effort

- **Benefits of Machine Learning**

- Adaptability to specific customer **data**, **workload** and **environment** characteristics
- Feedback for Retraining

- **ML Optimizer Goals**

- Automate everything
- Simplify optimizer development
- Infuse ML gradually

Motivation – Infuse AI Gradually (2|4)

- Cardinality estimation for common **Local Predicates**
 - Equality, Range, BETWEEN, IN, OR
- Join cardinality estimation
 - Pairwise Joins Equality, Multiple Join Predicates, Multiple Joins
- Enhance cardinality estimation
 - Commonly used expressions, Parameter Markers, Group By etc.,
- Join planning
- Other aspects

Motivation - Why Start With Cardinality Estimation? (3 | 4)

- Cardinality Estimation is the number of rows input to or output from an operator
- Critical for cost-based optimizers
- Primary source of query performance problem issues from customers

Motivation - Tuning For Good Cardinality Estimates (4 | 4)

Actual : 10,113,972

1000X off !!!



20X off !!!




Close Estimate and Better Plan



1136.92

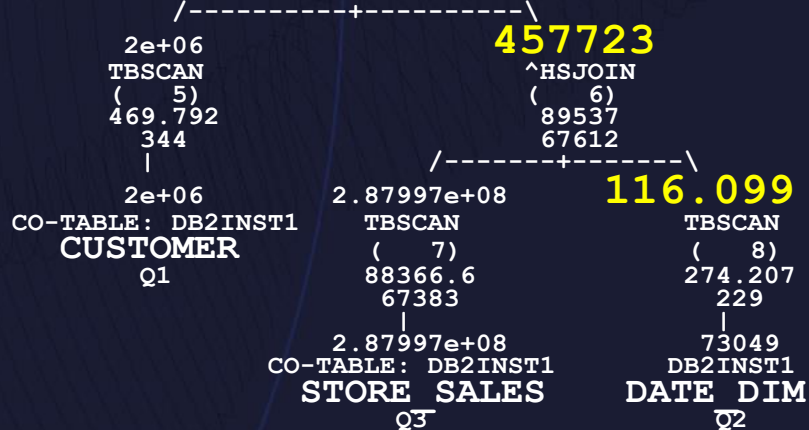
HSJOIN
(4)
90013
67956



Default Statistics

457723

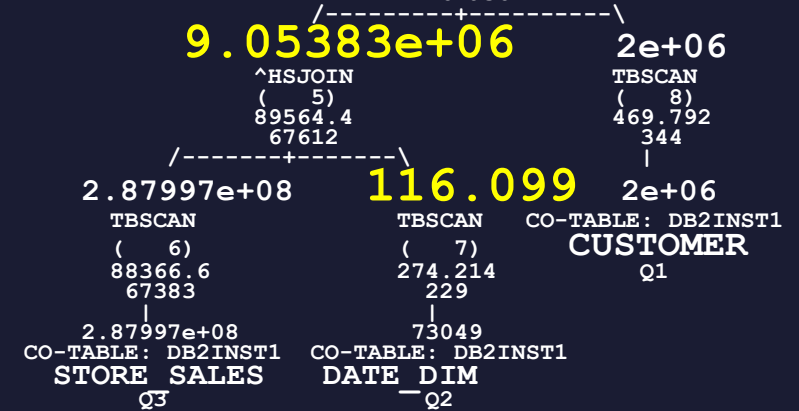
HSJOIN
(4)
90013
67956



With additional Column Group Statistics

9.05383e+06

^HSJOIN
(4)
90083.4
67956



With additional Statistical Views

The image is a dark blue background with a subtle grid pattern. On the left side, there is a circular inset containing a faint, light blue graphic of a server rack or data center. The text "Tech Preview" is centered in the lower half of the image in a white, sans-serif font.

Tech Preview

Tech Preview – Documentation (1 | 7)

- <https://www.ibm.com/support/pages/machine-learning-optimizer-technology-preview-db2-1156>
- Send questions and feedback to calisto@ca.ibm.com

Tech Preview – Try this on a Test System (2|7)

- **Enabling the ML Optimizer**
 - `db2set DB2_ML_OPT="ENABLE:ON"`
 - `db2 -tf MLOptimizerCreateTables.ddl`
 - Needs Auto-RUNSTATS
- **Toggle to look at how the traditional Optimizer does:**
 - `db2set -im DB2_SELECTIVITY="ML_PRED_SEL OFF"`
- **Disabling the ML Optimizer**
 - `db2set DB2_ML_OPT="ENABLE:OFF"`
 - drops all the models

Tech Preview – Try this on a Test System (3|7)

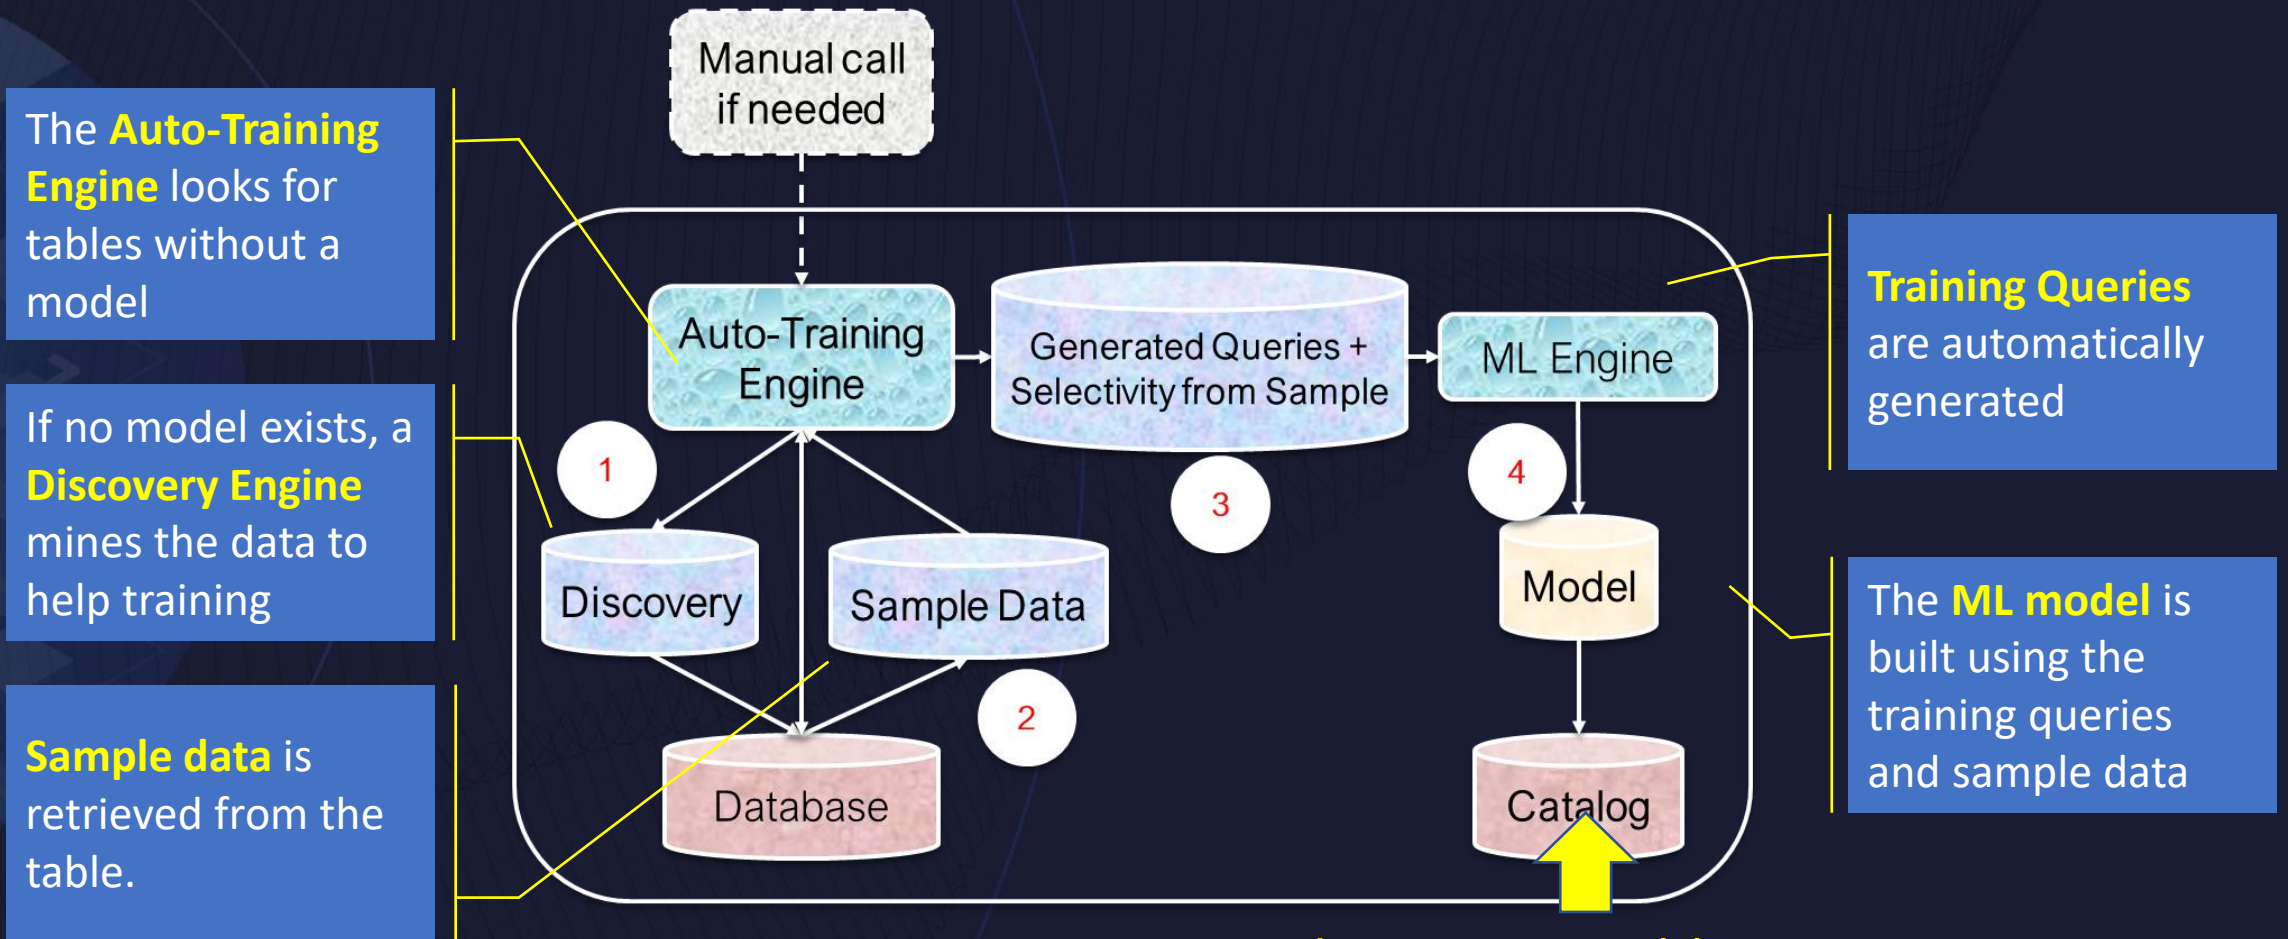
```
SELECT * FROM T1, T2
WHERE
  T1.C1 = 'abc' AND
  T1.C6 IN (5, 3, 205) AND
  T1.C2 BETWEEN 5 AND 10 AND
  T2.C3 <= 120 AND
  ((T1.C4 > 5 AND T1.C5 < 20) OR
   (T1.C4 < 2 AND T1.C5 = 100)) AND
  T1.C0 = T2.C0 AND
  T1.C3 = ? AND
  MOD(T1.C4, 10) = 1;
```

✓
Local Predicates with Equality,
Range, Between , IN, OR

✓
Pair-Wise Join Predicates

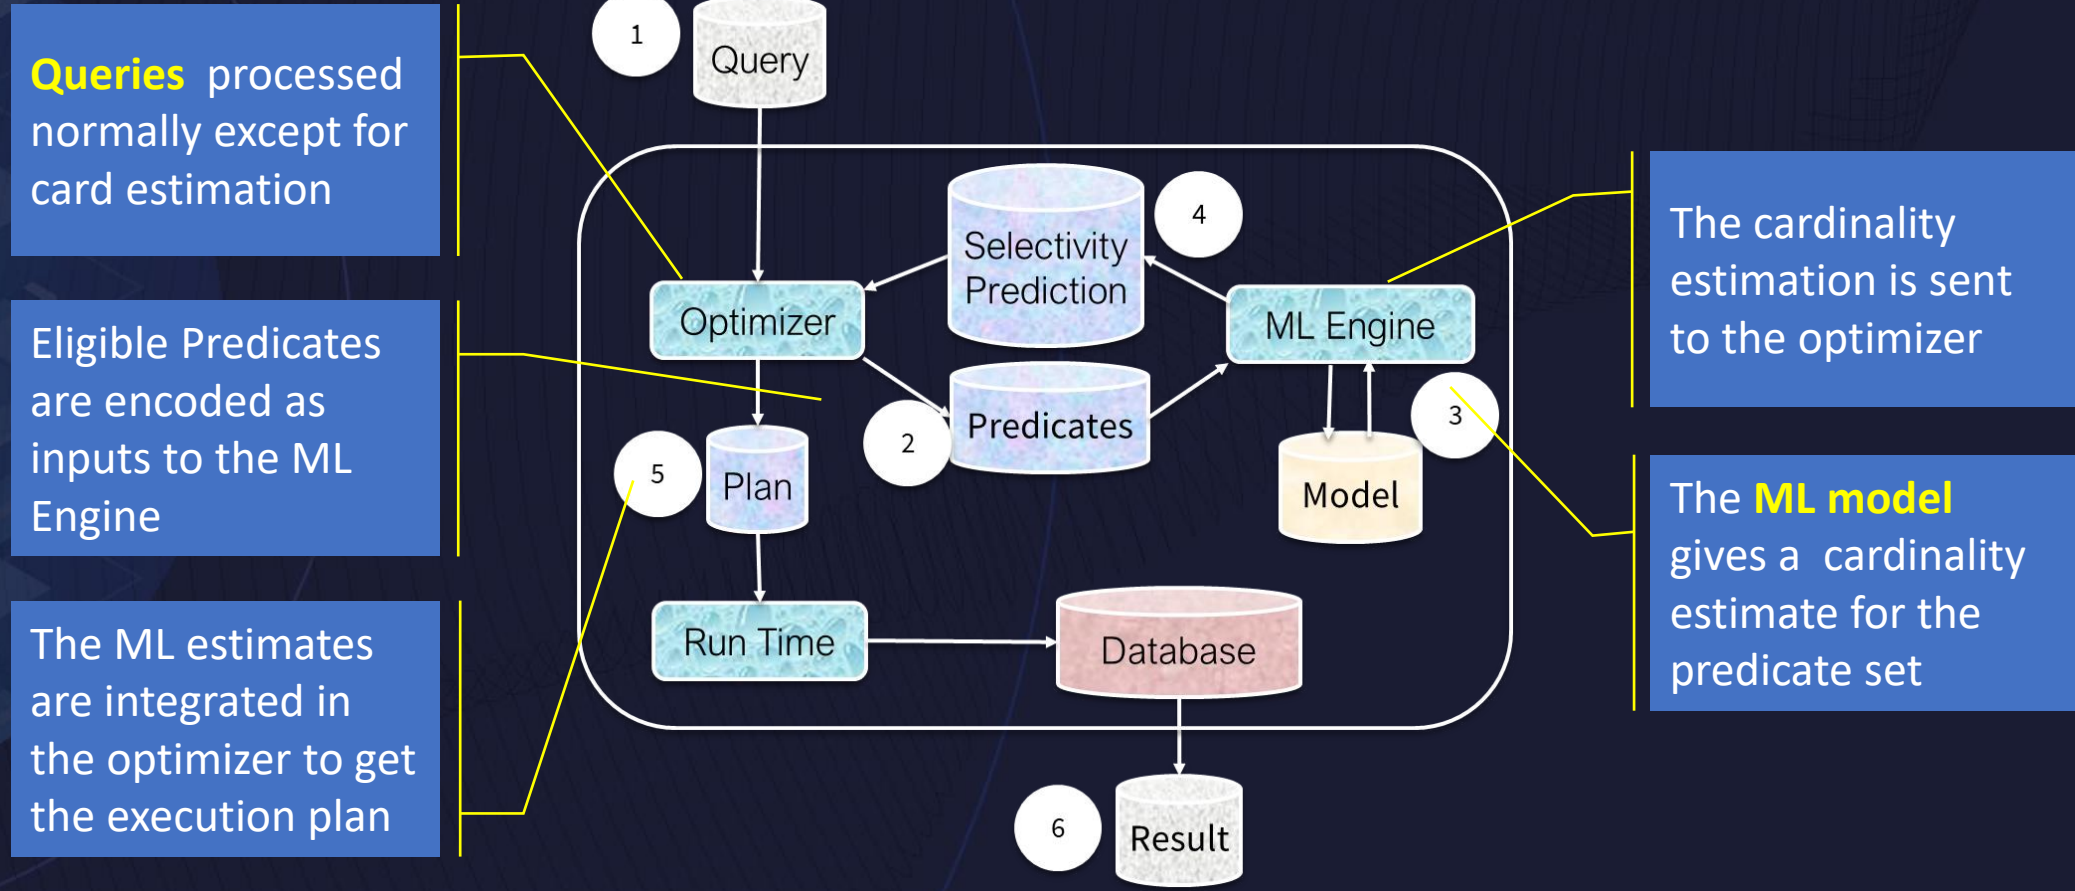
✗
Predicates With Parameter Markers
Predicates With Expressions
• These will be processed by the
traditional optimizer

Tech Preview – Automatic Training (4|7)



Tech Preview: Table in SYSTOOLS

Tech Preview – Cardinality Estimation Using the Model (5 | 7)



Tech Preview – Automatic Feedback (6|7)

- Like Auto-RUNSTATS, table data change counters are used to trigger retraining
- No optimizer or run time feedback in the Tech Preview

Tech Preview – Interesting Scenarios (7|7)

- ```
SELECT GUEST_LAST_NAME, ARRIVAL_DATE, DEPARTURE_DATE
FROM HOTEL_DB
WHERE (ARRIVAL_DATE <= '2019-12-25' and DEPARTURE_DATE >= '2019-12-25') OR
 (ARRIVAL_DATE <= '2018-12-25' and DEPARTURE_DATE >= '2018-12-25') OR
 (ARRIVAL_DATE <= '2017-12-25' and DEPARTURE_DATE >= '2017-12-25')
```
- ```
SELECT GUEST_LAST_NAME, ARRIVAL_DATE, DEPARTURE_DATE
FROM HOTEL_DB
WHERE DATE_COL BETWEEN '2019-08-01' and '2019-08-31') AND
      COMPANY = 'IBM'
```

The background is a dark blue gradient with abstract, overlapping circular and wavy shapes. On the left side, there is a faint, semi-transparent circular area containing a complex industrial or manufacturing diagram with various geometric shapes and lines.

Plan For Production

Plan for Productization – Infrastructure Enhancements (1 | 4)

- New system catalog table to store system AI models
 - SYSAIMODELS
 - Appropriate SYSCAT view for individual model type
- Security / access control
- Audit

Plan for Productization - Usability Enhancements (2 | 4)

- Explain support
- Activity logging
- Better model management
- Better configuration management
- Appropriate error messages
- DDL to manage models
- Appropriate dependency management

Plan for Productization – Model Enhancements (3|4)

- Improved ML model size : **~ 20 Kb to 30 Kb per table**
- Increased number of columns (**up to 20 instead of 10**) strongly correlated columns
- Improved Training Time : **~ 1 to 2 Minutes**
 - Dependent on number and characteristics of the columns included
 - Not so dependent on the table size

Plan for Productization – Configuration (4|4)

Automatic maintenance

Automatic database backup

Automatic table maintenance

Automatic runstats

Real-time statistics

Statistical views

Automatic sampling

Automatic column group statistics

Automatic reorganization

Automatic AI maintenance

Machine Learning Optimizer

Automatic Model Discovery

(AUTO_MAINT) = ON

(AUTO_DB_BACKUP) = OFF

(AUTO_TBL_MAINT) = ON

(AUTO_RUNSTATS) = ON

(AUTO_STMT_STATS) = ON

(AUTO_STATS_VIEWS) = OFF

(AUTO_SAMPLING) = ON

(AUTO_CG_STATS) = OFF

(AUTO_REORG) = OFF

(AUTO_AI_MAINT) = ON

(AUTO_ML_OPTIMIZER) = ON

(AUTO_ML_DISCOVER) = ON

Summary

- Infusing AI in the Db2 Optimizer is strategic
- Please try out the Tech Preview on your test system
 - <https://www.ibm.com/support/pages/machine-learning-optimizer-technology-preview-db2-1156>
 - Send questions and feedback to calisto@ca.ibm.com
- ML cardinality estimation plan for GA (vNext)
 - Initially with the local predicate model only and
 - Appropriate infrastructure for all future AI models for use within Db2

Thank You

Speaker: Calisto Zuzarte

Company: IBM

Email Address: calisto@ca.ibm.com

Session Code: C08

Please fill out your session evaluation before leaving!