IBM Db2 AI for z/OS (Db2ZAI)

Machine Learning within Db2 for z/OS

RTP Db2 User Group

Jerome Gilbert, NA Analytics Technical Sales Leader

Dec. 5th, 2018

IBM | IBM **z Analytics**

Design Thinking - Problem Statement

The z platform service providers are always being asked to proactively manage the Total Cost of Ownership (TCO) on the mainframe

- This includes the DBA team, Application Development / Support team, Database Infrastructure team, etc.

Challenged by the complexity of the applications compared to the required skills necessary to tune these applications executing in the Db2z environment.

Goal: Run faster (elapsed), more efficiently (less work) and use less resources (MIPS, Memory, I/O) when using Db2 z/OS.

1. Identify good candidates for application tuning.
2. Bring the right skills to the challenge:
   - *Sysprog*: system tuning skills & environment knowledge
   - *DBA*: application knowledge / query tuning skills
3. Experience

Query Optimization Challenges

"25-30% of customer DBA time is spent fighting access path problems which cause performance degradation and service impact."

~John Campbell, IBM Db2 Distinguished Engineer

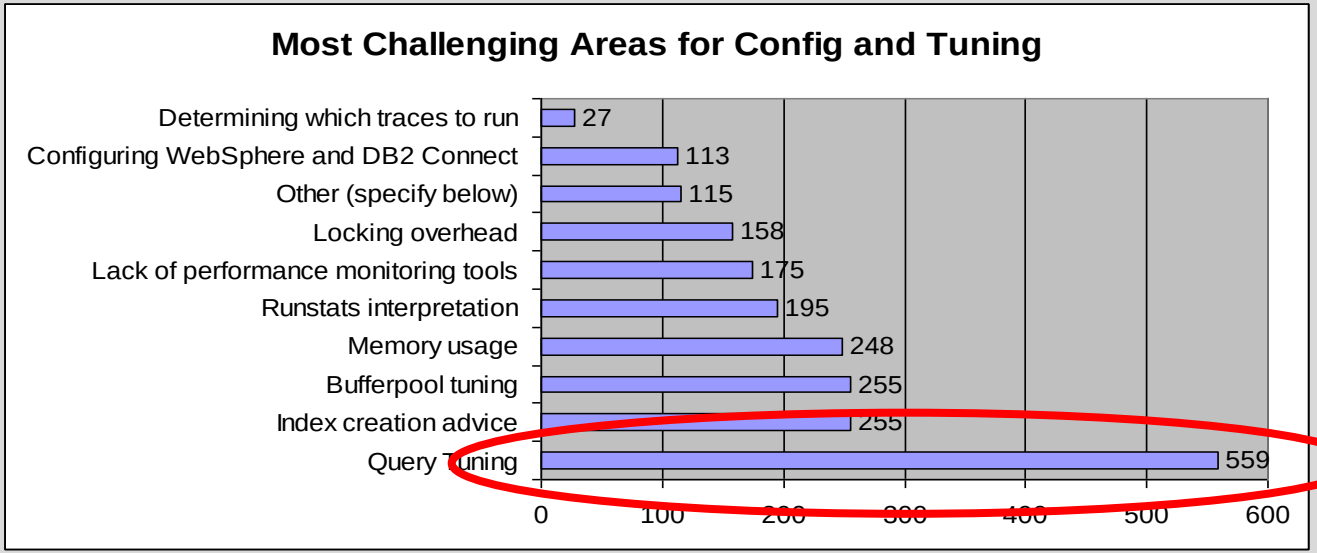

* Numbers are relative. Biggest pain point.

Query Optimization Challenges

User Challenge: Tuning just ONE query !

- Access path analysis can be complex and time consuming
  - Several factors: skill level, available tooling, etc

- IBM preferences:
  - Determine the appropriate RUNSTATS options for targeted Db2 objects
  - Indexing (New or Modified)
  - OPTIMIZE FOR n ROWS

Skill often missing

- Customer preferences:
  - BIND / REBIND with APREUSE or OPTHINTS
  - Query rewrite
  - “Fudging” catalog statistics

Organizations often revert to these

Db2ZAI – Query Optimization

- Solution
  - Targeting the most significant inputs to the Optimizer's access path choice – that Db2 must estimate today
    - Actual value of Host variable/parameter markers
    - Number of rows required by the application (OPTIMIZE FOR n ROWS)
- Learn from the actual customer workload
  - By collecting query execution history and feeding information into ML models
  - Available for the optimizer at next bind/rebind/prepare

Db2ZAI – Business Value

Db2z infused with ML

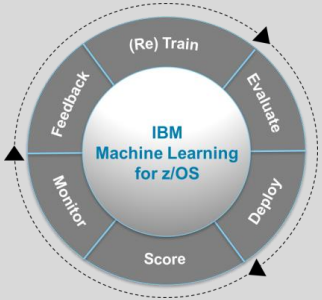

MLz

IBM Z

Up to 25% CPU Savings*

Value Proposition

Enable Db2 for z/OS optimizer to leverage Machine Learning for z/OS (“MLz”) services and IBM Z

- Reduce CPU consumption and IT cost through optimization for best query access paths
- Improve Db2 application performance
- Rapid model learning specific to the data/application behavior per subsystem without requiring data science skills

**CPU savings metric based on IBM internal benchmarks; actual savings will vary according to customer workloads and environment*

What is Machine Learning

- Computers apply **statistical learning techniques** to automatically identify patterns in data and make highly accurate predictions.
- Use example data or past experience to learn without being explicitly programmed.
- The OUTPUT is a function of the INPUT (Data -> Analysis -> Actionable Insight)
- Leading Open Source technologies & skills

What is Machine Learning (simplified):

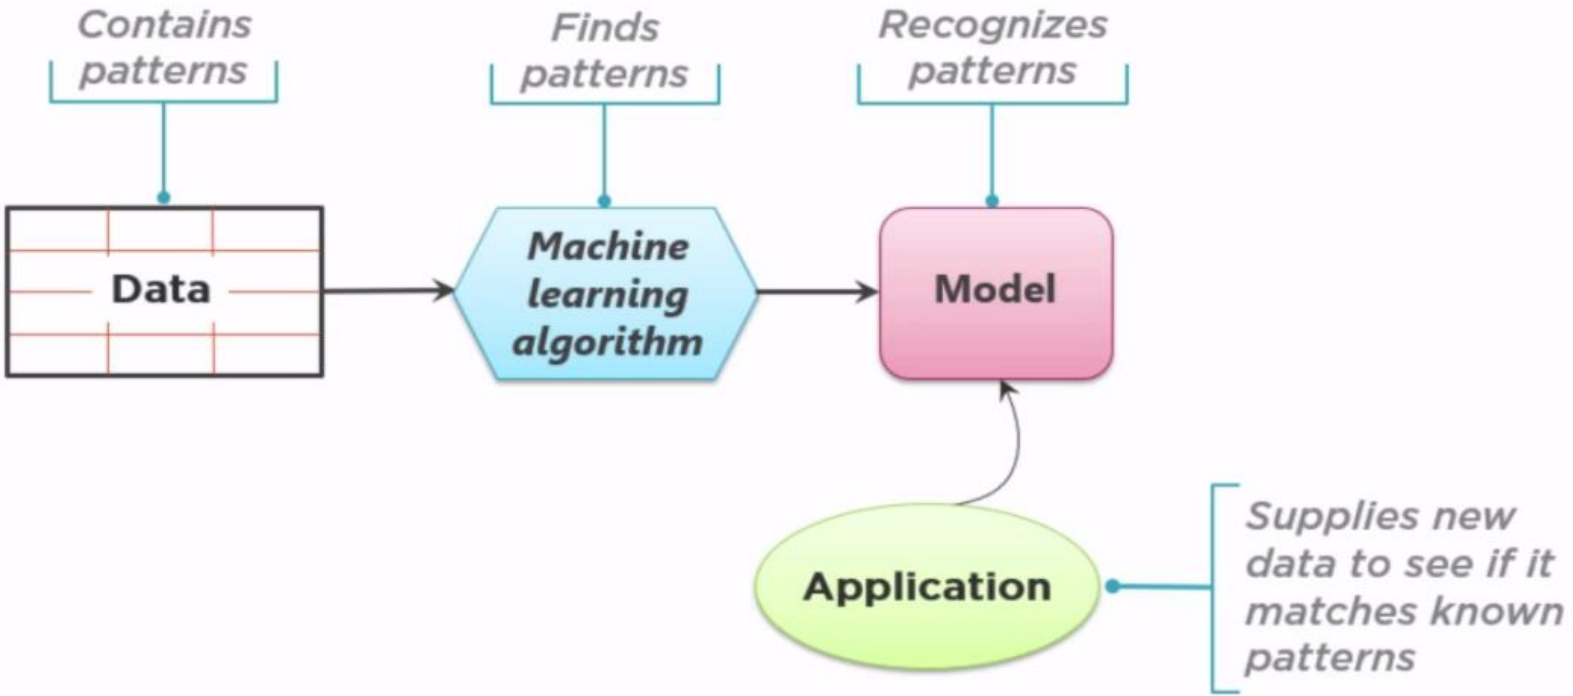

What is Machine Learning (drill down another level):

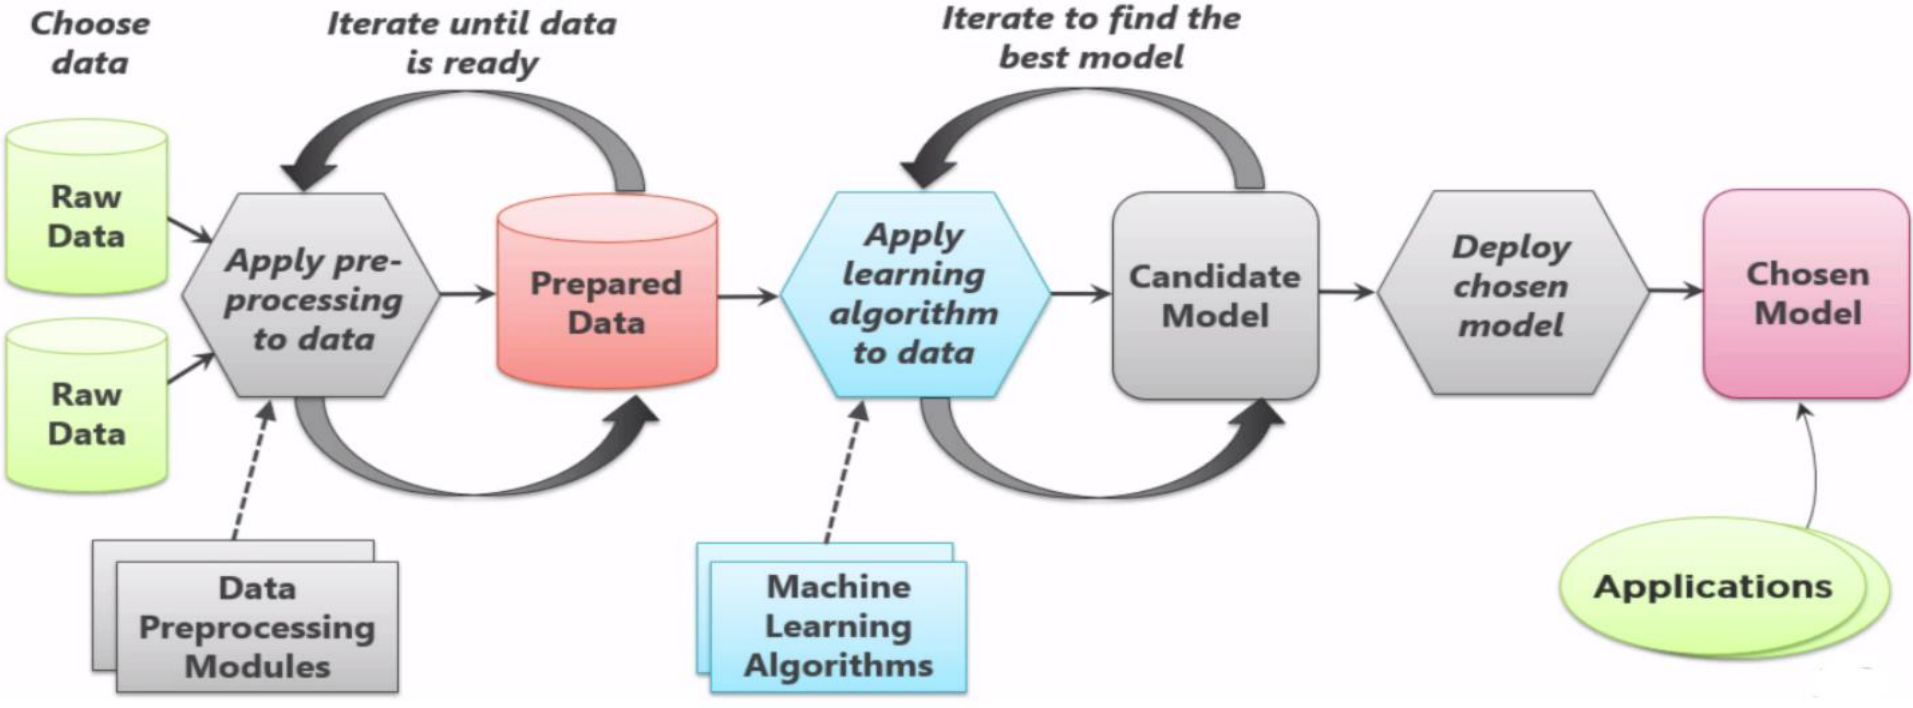

Db2ZAI – Query Optimization

- Solution
  - Targeting the most significant inputs to the Optimizer's access path choice – that Db2 must estimate today
    - Actual value of Host variable/parameter markers
    - Number of rows required by the application (OPTIMIZE FOR n ROWS)
  - Learn from the actual customer workload
    - By collecting query execution history and feeding information into ML models
    - Available for the optimizer at next bind/rebind/prepare

Key Db2ZAI Processing:

Data Scientist & DBA Resources Utilized: Minimum to None

1. Execution History in Db2:
    1. Execution information written to pseudo-catalog tables (per execution)
    2. No Traces; Internal Db2 data
    3. Db2 tables created as part of set-up
        1. Db2ZAI manages summarizations & purging
        2. User managed Reorgs & Runstats
2. Learning Tasks (Exploration)
    1. zIIP Eligible
    2. *Model creation/training/deployment*
3. Next BIND/REBIND/Prepare/Explain will use the model
4. User Interface in MLz
    1. Start / Stop ML per Db2 subsystem
    2. Display package and SQL performance metrics
    3. Display model statistics

Key Db2ZAI Externals

Externals all surfaced from the UI

- Identifying packages that are using and/or candidates for ML
- SQL Performance Metrics

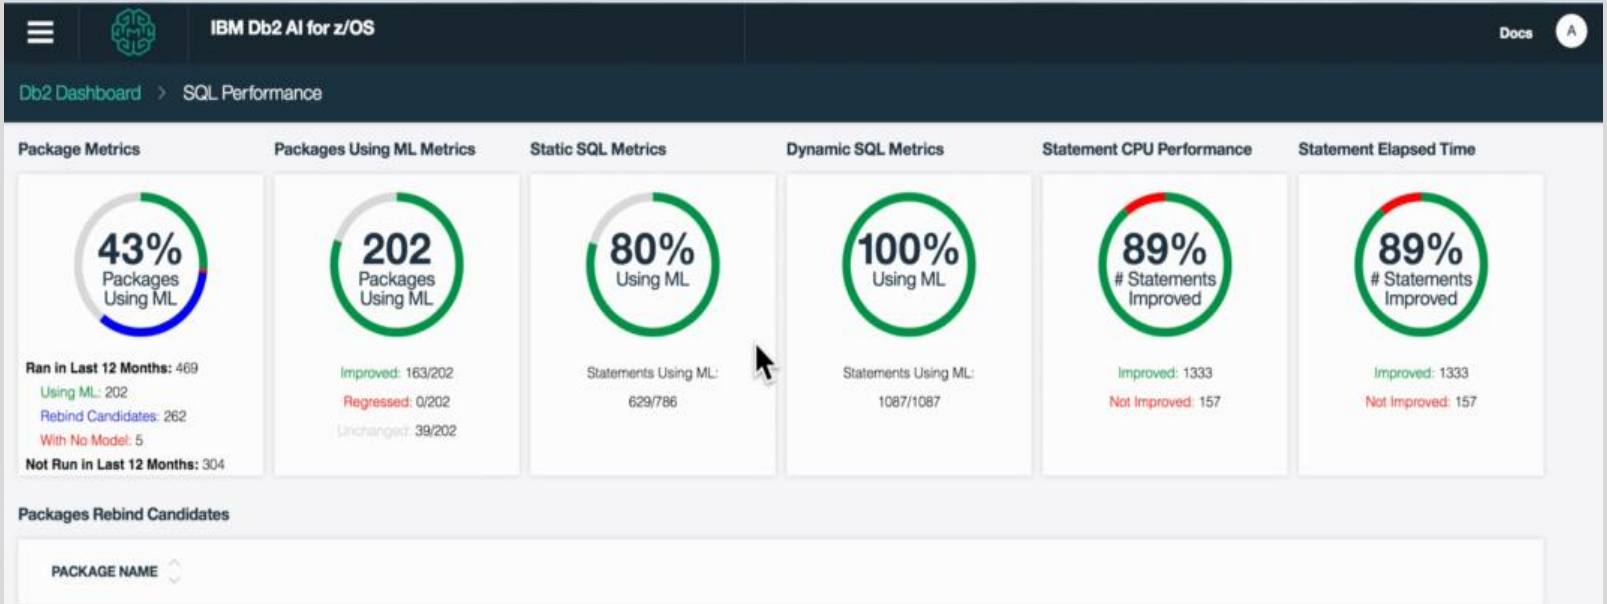

Db2ZAI - Architecture

= Db2ZAI

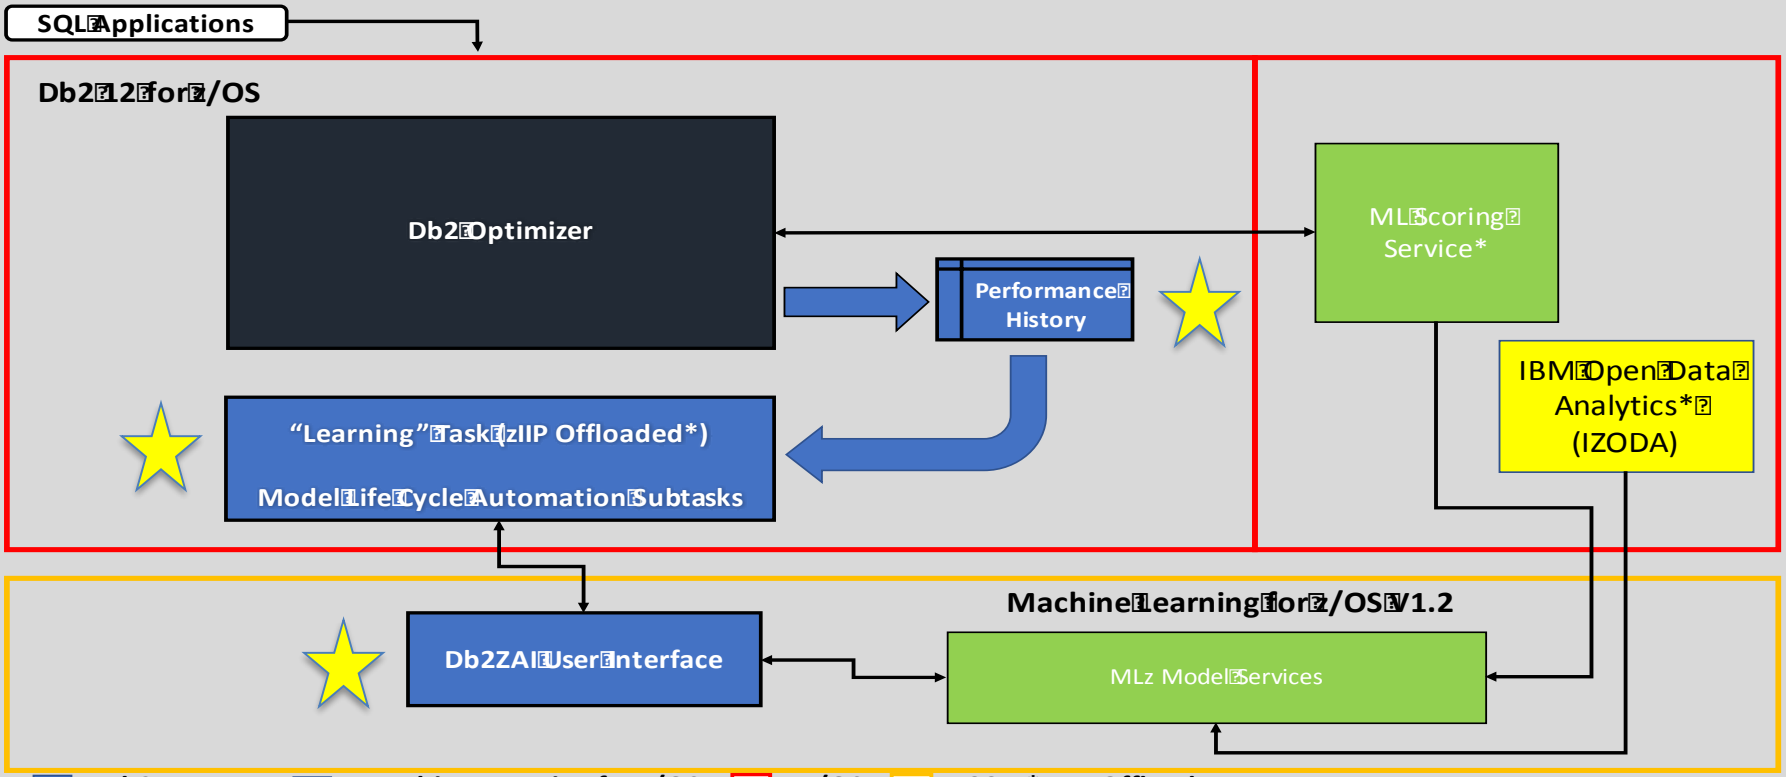

Db2ZAI – Software Pre-requisites

- IBM Db2 AI for z/OS, V1.1 requires:
  - z/OS 2.1 or later, with the Product Registration Services component.
  - **IBM Db2 12 for z/OS** (either 5650-DB2 or 5770-AF3),
    - plus Function Level 503 APAR PH00506. Note: FL503 *must be activated.*
  - **IBM Machine Learning for z/OS, V1.2** (5698-ML1)
    - plus APARs PI99400 and PI99401 and fix pack 1.2.0.1.
    - Incl IBM Open Data Analytics for z/OS

Value Proposition:

1. Automated ML Pipeline: Capture and retention of performance data (limited to no DBA activity).
2. Automated ML Pipeline: Feature extraction, transformation and exploration (no Data Engineer / Scientist)
3. Automated ML Pipeline: Model creation, training & deployment (no Data Scientist)
4. Optimizer Input: represents customer execution patterns (default filter factors and optimization based on the required portion of the result set)
5. MLz Foundation: Db2 as self tuning / self managing Subsystem (roadmap possibilities)
   a. Memory allocations: Query level & System Level
   b. PCTFREE / FREEPAGE allocations based on THE customer's Insert patterns
   c. Db2 Analytics Accelerator offload decisions
   d. Utilities: RUNSTATS / Optimizer integration;
   e. Applying maintenance

Thank You – Questions?

How do I know if my workload will benefit:

1. Db2ZAI Analyzer: script that can identify the number of
    1. SQL statements with the risk of sub-optimal access path due to
        1. Host variables / parameter markers
        2. WHERE clause predicates with data skew or range predicates
    2. Index design which exposes the optimizer to index choices that have risk
        1. Or multi-table queries where risk of choosing the wrong join sequence/method exists

The output of this script is intended to identify the number of SQL statements that have the statement structure which forces the optimizer to use estimates or defaults. It is these statements that Db2ZAI would attempt to improve the optimizer inputs for.

Db2ZAI compared to other Performance Tools

Db2ZAI

Improves the optimizer's intelligence by learning patterns from the collected data from workloads in customer's unique operating environment and determine the optimal paths for SQL statements entering Db2 for z/OS

Optim Query Workload Tuner

Reduces skill requirements for tuning Db2 queries and improving application performance. Advisors give actionable tuning recommendations.

Query Monitor

Low overhead SQL monitor. Can proactively identify SQL stmts which consume excessive resources and inadequate object structures. Can determine which tables and indexes are actually being used by queries for tuning.

IBM Db2 Analytics Accelerator

Optimizes the performance of longer running (data warehouse or analytics) queries