June, 2017

IBM DB2 Analytics Accelerator
for z/OS Update

Peter Bendel
IBM STSM

DB2 Analytics Accelerator
for z/OS on Cloud

DB2

DB2 Analytics
Accelerator
for z/OS

Disclaimer

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and at IBM's sole discretion.

- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.

- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.

- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

- Level-set:　　What is IDAA and the recent enhancements

- Big picture:　　Long-term strategy

- Work in progress:　　Top themes

- Level-set: What is IDAA and the recent enhancements

- Big picture: Long-term strategy

- Work in progress: Top themes

The hybrid computing platform on System z

➢ *Supports transaction processing and analytics workloads concurrently, efficiently and cost-effectively*

➢ *Delivers industry leading performance for mixed workloads*

➢ *The unique heterogeneous scale-out platform*

➢ *Superior availability, reliability and security*

Transaction Processing

Analytical Workload

DB2

DB2 Analytics Accelerator for z/OS

DB2 Accelerator personalities

➢ *Turbo charged access path with hardware assisted early filtering*

➢ *Full-width index*

➢ *Specialty engine*

➢ *Archive*

➢ *Tablespace*

➢ *ETL/ELT and in-database analytics acceleration*

➢ *Integration hub: fast federated joins across heterogeneous sources*

➢ *Hybrid cloud*

➢ *Accelerator-only tables* can benefit statistics and analytics tools that use temporary data for reports.
They enable acceleration of data transformations implemented via SQL statements. Storing interim results in accelerator-only tables enables subsequent queries or data transformations to process all relevant data on the accelerator with high speed

➢ *In-database analytics* capabilities enable acceleration of predictive analytics applications. This enables SPSS/Netezza Analytics data mining and in-database modeling to be processed within the Accelerator.

➢ *System Temporal and Bi-temporal tables support*
- Active and history table loaded to Accelerator
- Timestamp(12) data type now supported
  - Truncation to timestamp(6), new options for ZPARM QUERY_ACCEL_OPTIONS
- Queries using temporal SQL expressions can be routed to Accelerator

➢ *Transparent archive tables support (since V4PTF3)*
- Archive-enabled and archive table table loaded to Accelerator
- Archive table older partitions can be archived in IDAA

➢ *Encryption of data-in-motion*
- Encrypt network traffic between z Systems and the Accelerator using *Internet Protocol Security (IPsec).*
- Network traffic includes:
  - DRDA traffic (e.g. Queries, Table Loads)
  - Configuration Console traffic
  - Incremental Update traffic
- Requires configuration and enablement on z Systems and on the Accelerator

➢ *Incremental update operation enhancements*
- Ability to disable query acceleration automatically for suspended tables
- Continuous replication and re-load of replicated tables without taking a table lock on the source

➢ *Improved installation and maintenance process*
- Extended compatibility between stored procedure version level and accelerator version level
- Enhanced product packaging

➢ *More supported SQL functions*
- Date/Time arithmetic functions
- Inline scalar UDFs

➢ *IDAA on Cloud*

- Overwhelming general market trends and IBM focus
- Demonstrating z analytics vitality
- Fast deployment, quick time to value
- Basis for hybrid cloud offering
- Basis for change of accelerator engine (dashDB)
- Gaining necessary new skills

➢ *Mix and match*

- Simultaneous deployment of on-premises and cloud IDAA servers

➢ *R support in IDAA on prem*

- Support for data scientists most popular tool
- Enabling in-database analytics

Introducing Accelerator-only table type in DB2 for z/OS
Creation (DDL) and access remains through DB2 for z/OS in all cases

Non-accelerator DB2 table
- Data in DB2 only

Accelerator-shadow table
- Data in DB2 and the Accelerator

Accelerator-archived table / partition
- Empty read-only partition in DB2
- Partition data is in Accelerator only

Accelerator-only table (AOT)
- "Proxy table" in DB2
- Data is in Accelerator only
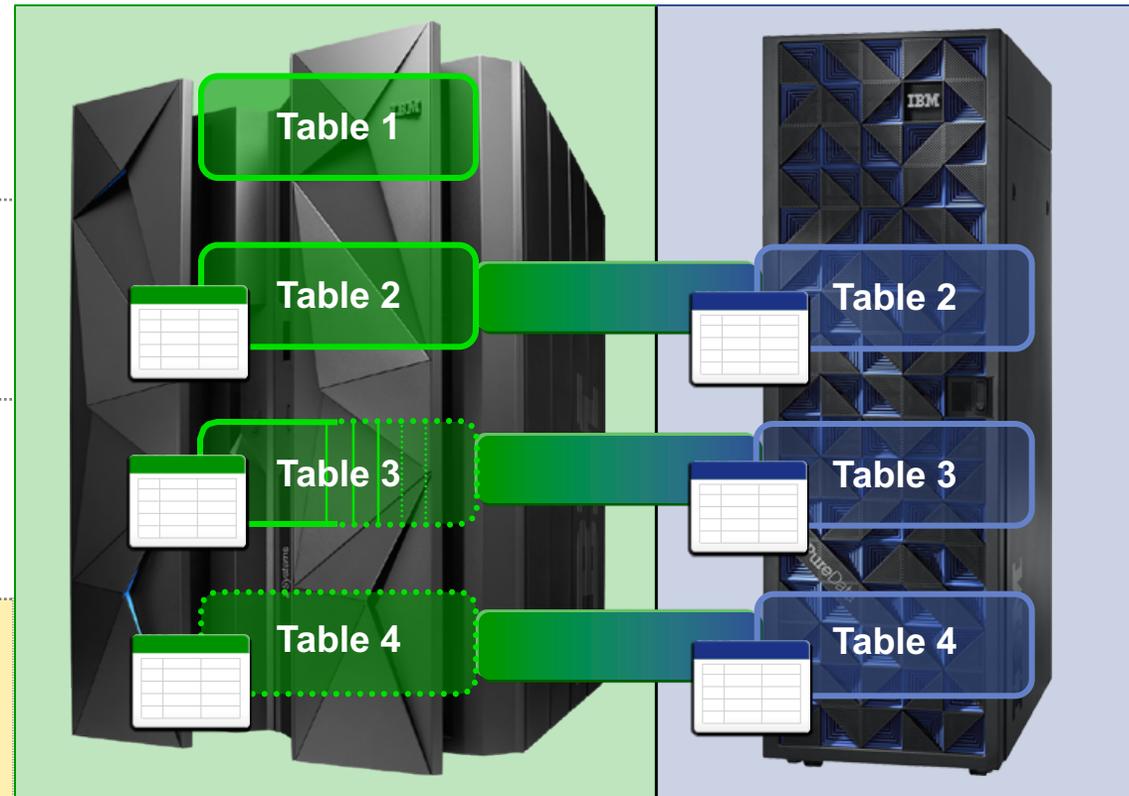
New use cases with Accelerator-only tables

- Adding application support for temporary objects
  - Multi-step Reporting
  - QMF
  - Campaign

- In-database transformation
  - Consolidation of ETL/ELT processing in DB2 for z/OS
  - Support for Data Stage Balanced Optimization

- Individual ad-hoc analysis
  - Data Scientist Work Area

- Integrate more data sources for analytics
  - Use DB2 Analytics Accelerator Loader to integrate with IMS data or data from other sources
  - Leverage High Performance Storage Saver

Application support for temporary objects

In-database transformation

Ad-Hoc analysis

Integrate more data sources for analytics

Accelerate existing workload

Reduce IT sprawl

Derive new business insight

Include external & historical data

1

2

3

4

IBM DB2 Analytics Accelerator Loader for z/OS v2.1
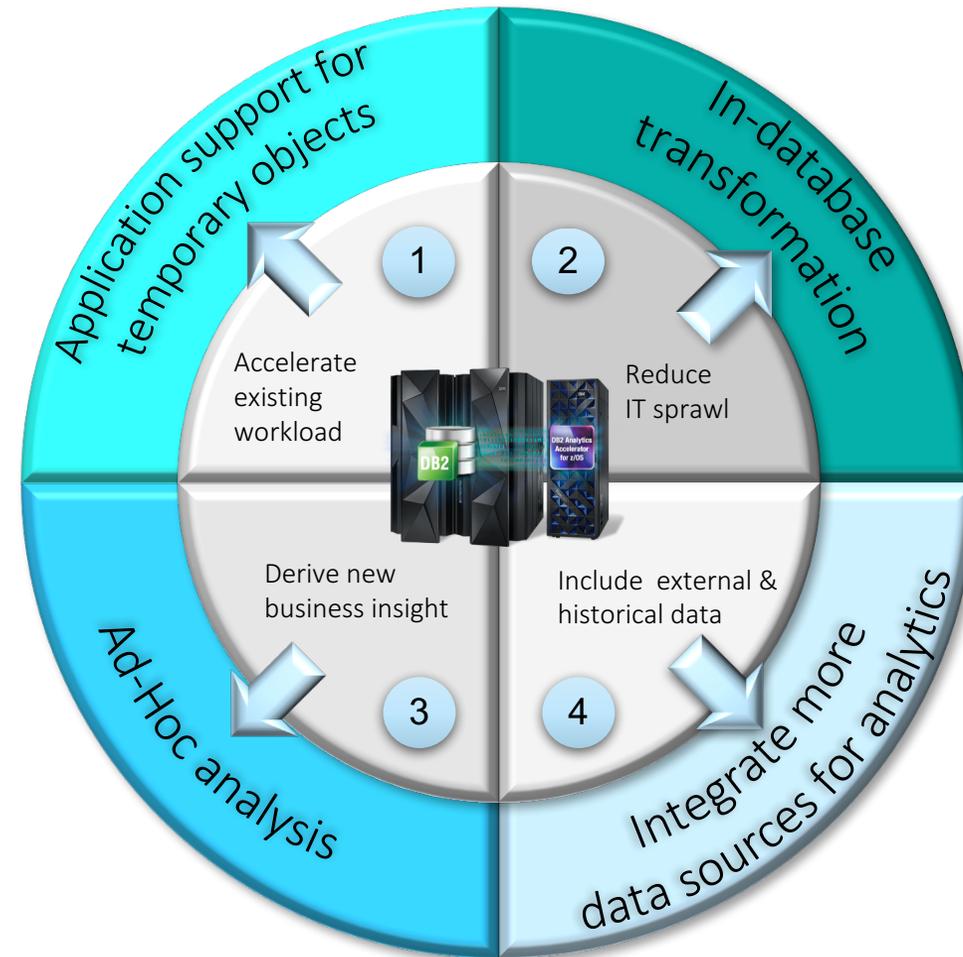
- *Extending *analytic capabilities* by bringing *non-DB2 data* to the accelerator and z Systems*
  - *Significant cost and time reduction* by eliminating manual ETL processes of non-DB2 data*
  - *Insight* into more data types such as *IMS, VSAM, sequential files, Oracle, Adabas, etc.*
  - *Accelerate *operational analytics* by supporting *SMF, RMF, Syslog, ... data*

- *Ensure High Availability* for critical analytics applications*
  - *Loading multiple accelerators in parallel: a single invocation and a single unload from DB2*
  - *Applies to non-DB2 sources as well*

- *Add data to existing table* to avoid reloading entire table*

- *z Systems synergy: *zIIP usage*

- IDAA V5 PTF 4 new features

  - Changing the IP address or the TCP/IP port of the Capture Agent

  - Restoring archived partitions from multiple accelerators

  - Directing accelerated queries to a specific accelerator

  - **Supporting multiple-row insert for Accelerator-only tables (AOTs)**

- DB2 Analytics Accelerator Loader 2.1 PTFs

  - **Backup and Recovery for accelerator tables including Accelerator-Only-Tables**

    - Available via APAR PI70981/PTF UI44867

    - see http://www.worldofdb2.com/profiles/blog/show?id=2837977%3ABlogPost%3A116258&xgs=1 &xg_source=msg_share_post

Directing accelerated queries to a specific accelerator

- For users with multiple accelerators: A new special register and bind option will allow users to direct accelerated queries to a specific accelerator
  - DB2 11 only
  - Delivered via DB2 PTFs only, no Accelerator update to PTF 4 required: PI72150, PI73563

- Currently, a user cannot specify a target accelerator for an accelerated query to be executed. If a user has multiple accelerators where the same set of DB2 tables are accelerated, the accelerator workload balance algorithm will distribute the queries depending on queue length on each accelerator.

- Use Cases:
  - User has an additional accelerator on a remote location (DR location) in addition to locally deployed accelerators. Current algorithm does not take latency into account caused by the long network distance. Therefore some high priority queries might be sent to the remote accelerators and incur increased elapsed time.
  - User wants to direct queries to different accelerators depending on the priorities of the workload, with high priority queries directed to the fastest, highest capacity accelerator and low priority queries directed to a slower accelerator.

Directing accelerated queries to a specific accelerator - Interfaces

- **CURRENT ACCELERATOR (special register)**
  - Specifies the name of preferred accelerator(s) to which DB2 should send accelerated dynamic queries.
  - If none of the accelerator servers named by CURRENT ACCELERATOR are available or eligible, DB2 will consider other available accelerator servers.
  - Syntax:
    - SET CURRENT ACCELERATOR = <accelerator_name>
    - <accelerator-name> is a single accelerator name or an accelerator logical name as recorded in SYSIBM.LOCATIONS (the logical name represents one or more accelerator server names).

- **ACCELERATOR BIND OPTION**
  - Set the target accelerator(s) to be used for static queries
    - single accelerator name or an accelerator logical name
  - Only affects runtime behavior of an accelerated query
    - If one of the specified accelerator(s) does not exist during Bind time, then DB2 issues a warning message and will not fail the Bind.

Multiple-row insert for Accelerator-only tables (AOTs)

- Support of multiple-row INSERTs for accelerator-only tables
  - Allows faster inserts into accelerator-only table through JDBC, ODBC or local SQL applications
- Note: This feature is still in Beta state
- DB2 11 only
- Required: PTF 4 on the Accelerator and DB2 PTF PI66573

- Performance numbers from test in Lab:
  - Multiple-row insert is ~66 times faster than single row insert
    - JDBC program with batch size of 32000 rows for multiple-row insert

- Use cases:
  - Populating AOTs using INSERT statements from external sources
  - ETL processing outside DB2, then insert result into AOT
  - Populating AOTs from analytics tools such as IBM Campaign or Data Stage

From DB2 INSERT statement:

```
multiple-row-insert:

>>-VALUES--+-+-expression---------+--------+-----
          | +-host-variable-array-+       |
          | +-NULL----------------+       |
          | '-DEFAULT------------'        |
          |    .-,-----------------------. |
          |    V                         | |
          '-(---+-expression---------+-+-)-'
                +-host-variable-array-+
                +-NULL----------------+
                '-DEFAULT------------'
```

Example for Multiple-row insert using JDBC

```
PreparedStatement prepStmt = conn.prepareStatement(
            "INSERT INTO SYSADM.MYEMP (EMPNO, LASTNAME, MIDINIT, FIRSTNME) VALUES(?, ?, ?, ?)");
System.out.println ("successful PREPARE of INSERT" );

prepStmt.setString(1,"111111");
prepStmt.setString(2,"Davis");
prepStmt.setString(3,"F");
prepStmt.setString(4,"Paul");
prepStmt.addBatch();

prepStmt.setString(1,"222222");
prepStmt.setString(2,"Spenser");
prepStmt.setString(3,"K");
prepStmt.setString(4,"Mary");
prepStmt.addBatch();

int [] numRowsUpdatedArray=prepStmt.executeBatch();
```

Agenda

- Level-set: What is IDAA and the recent enhancements

- **Big picture:** Long-term strategy

- Work in progress: Top themes

1. **Bring analytics to z data**

   ▪ DB2 for z/OS as the control point

     • Deeply integrated and transparent to typical personas that use database

     • Integration hub for z and non-z data

   ▪ Industry leading:

     • Complex query performance

     • In-database transformation acceleration

     • Archiving

2. **Hybrid platform architecture**

   ▪ Workload-optimized 'system of systems'

   ▪ Best of breed for each component

   ▪ Hybrid analytics/transactional processing (HTAP)

3. **Support hybrid cloud**

   ▪ Uniform experience, simultaneous use and easy transition between on-premises, private cloud and public cloud implementations

Planned major features delivered in upcoming V5 PTFs

- HTAP

  – Real-time processing on real-time data on "best-of-breed" technologies for transactional and analytical workload

- Schema change support

  – Adjust the schema on the IDAA side to the changed schema in DB2

- Federated access support

  – Enable read access to all data in one accelerator through a single DB2 system

HTAP – Real-time processing on real-time data

HTAP

HTAP Workload

Transactions

Analytics

DB2 Analytics Accelerator for z/OS
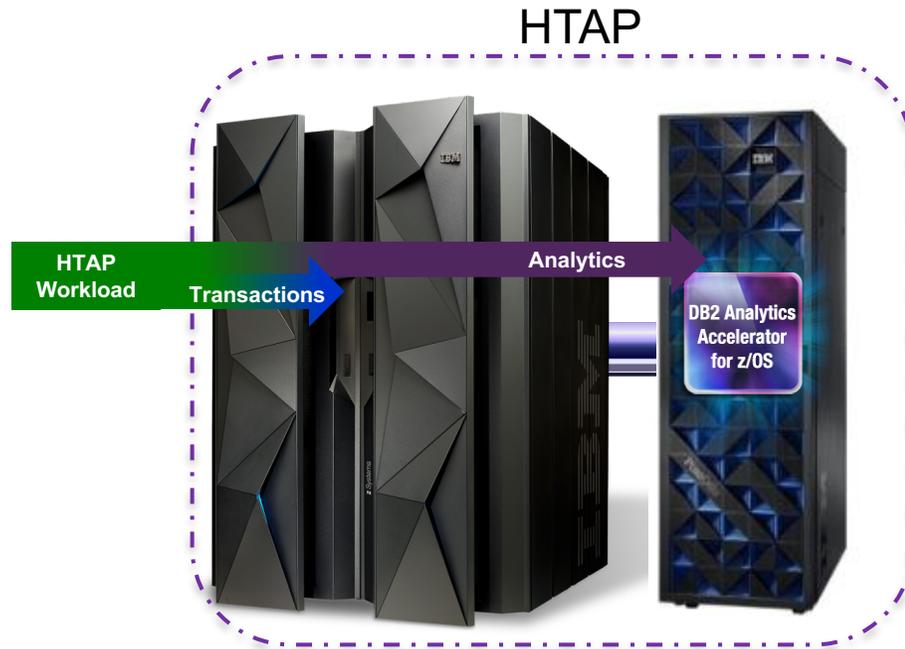
Top priority theme
Hybrid Transactional/Analytical Processing

- Real-time processing on real-time data on "best-of-breed" technologies for transactional and analytical workload

- IBM research team has pioneered this technology

- A patented *just-in-time* currency replication protocol guarantees transactional data coherency for analytical requests

- IBM looks at the ability to leverage DB2 z/OS data on the IBM DB2 Analytics Accelerator w/o recognizable latency with very high focus and views this as a high priority theme

- The result is a unique-in-industry, heterogeneous scale-out solution for enterprise-grade HTAP

Basic idea

Reading most recent committed data during asynchronous replication
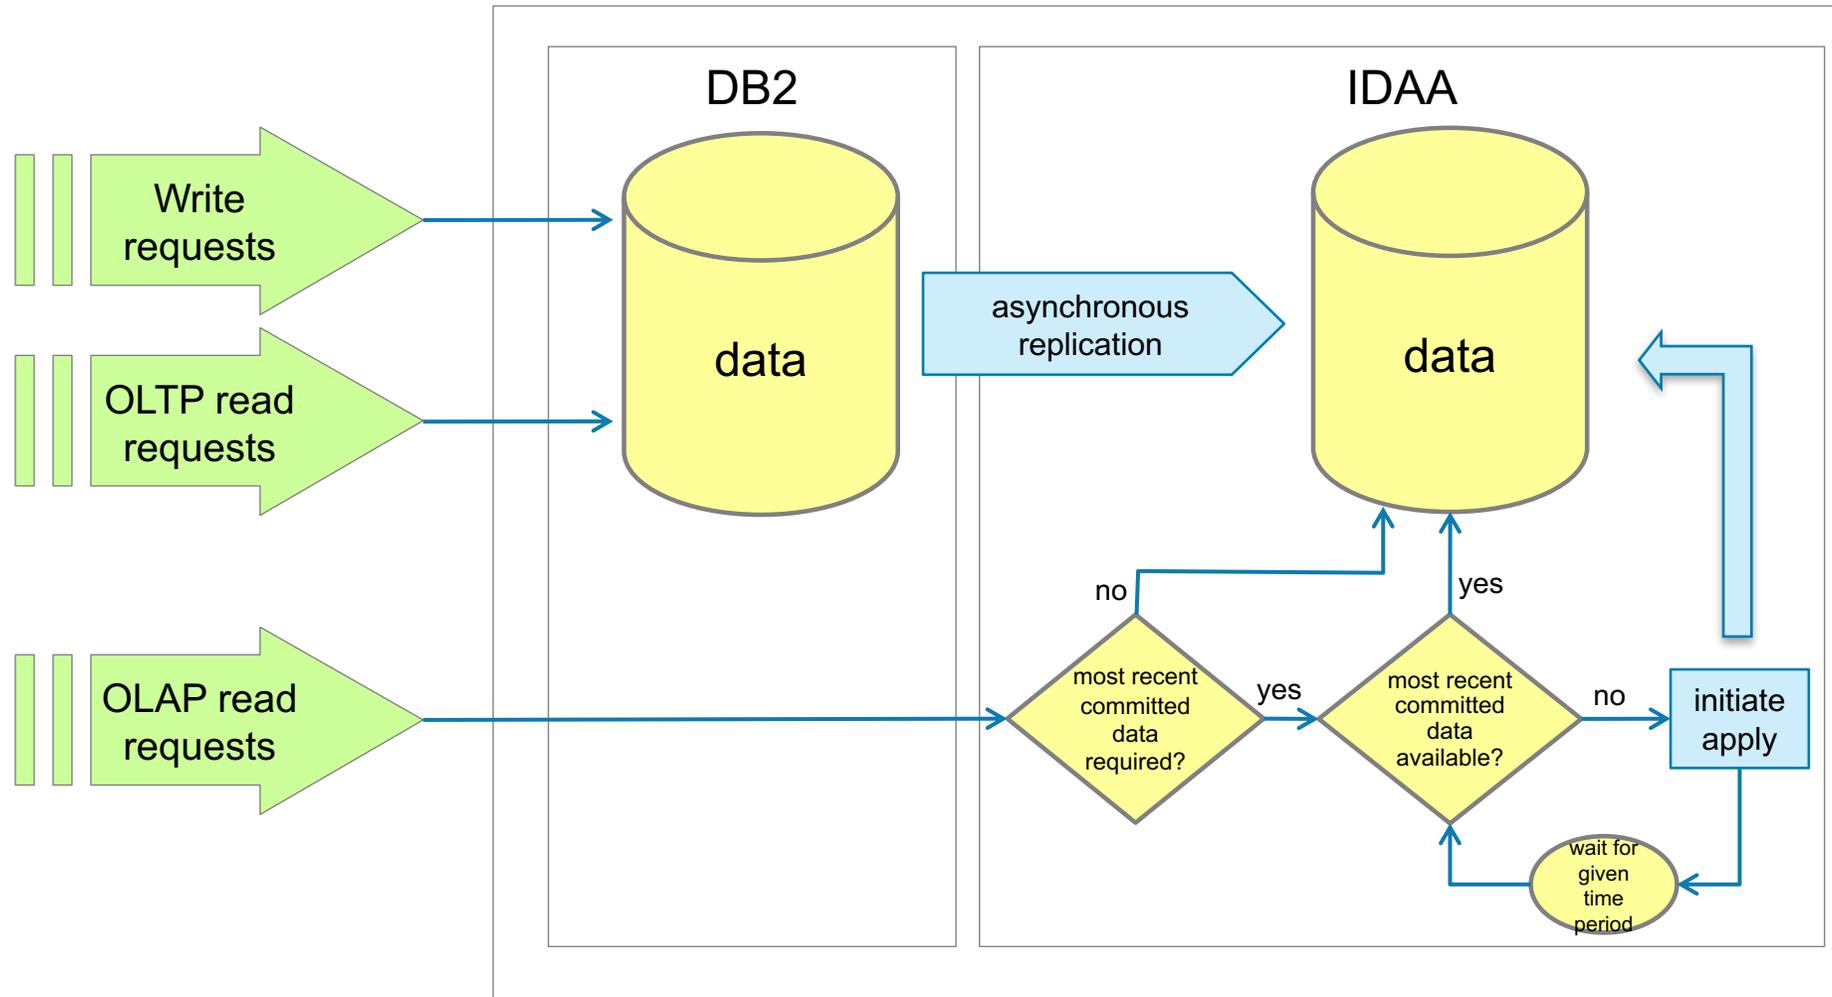
- Introducing a new zparm / bind option / special register
  - CURRENT QUERY ACCELERATION WAITFORDATA = n.m
  - 0.0 - 3600.0 seconds, 0.0 = no wait, default: 0.0

    0.0:
    - immediately execute in accelerator without waiting

    > 0.0:
    - wait for most recently committed changes before query to be applied by asynchronous replication and then run the query
    - or: fail the query if max wait time specified exceeded
    - or: execute query in DB2 if "WITH FAILBACK" is specified in CURRENT QUERY ACCELERATION special register

CURRENT QUERY ACCELERATION = ENABLE WITH FAILBACK
and query is routable to IDAA
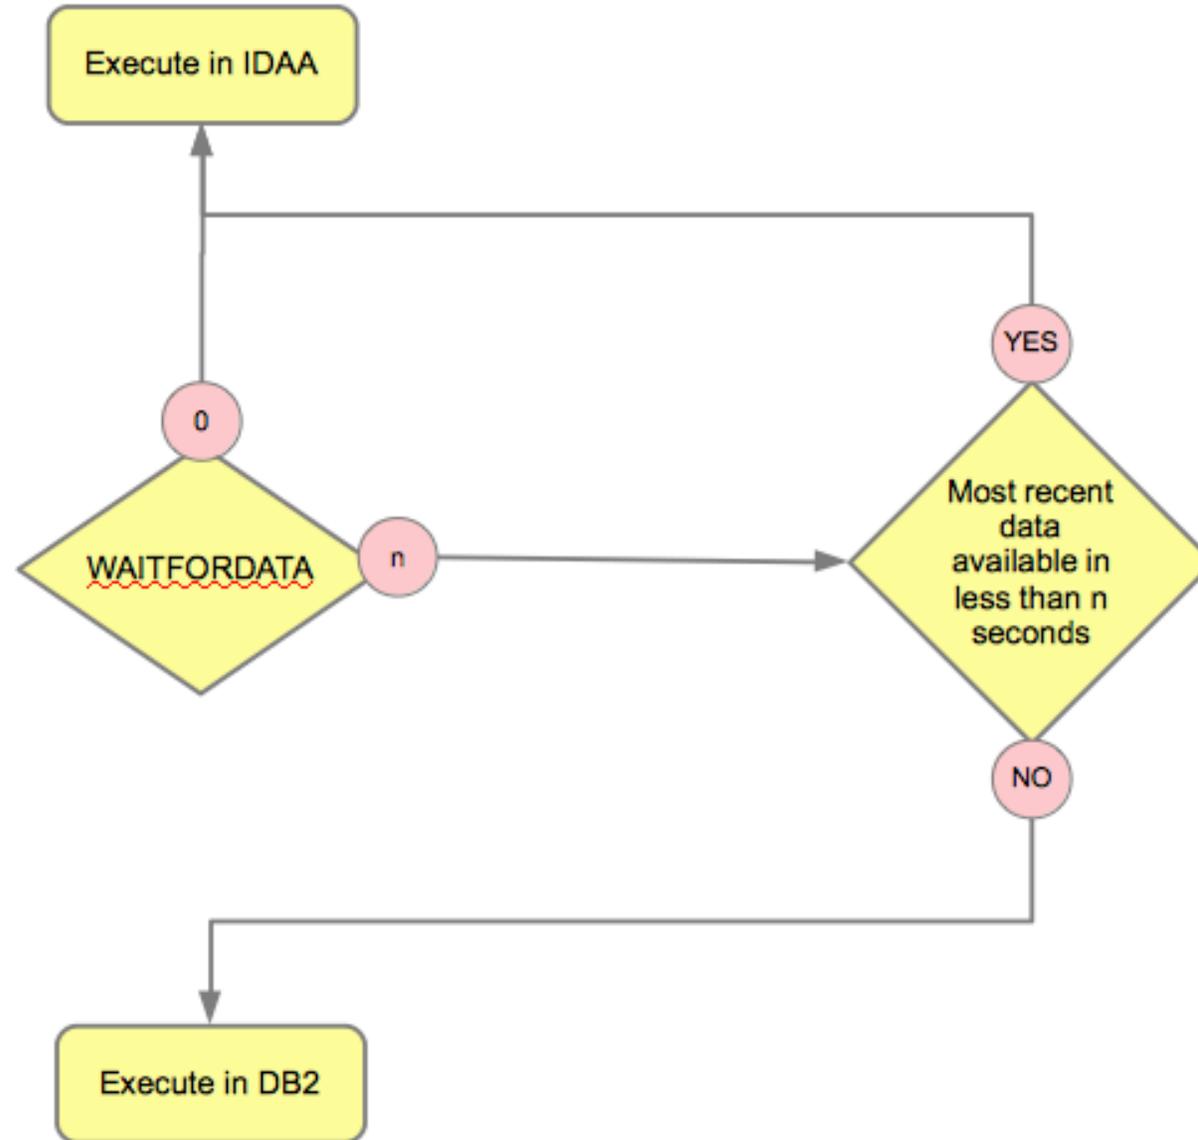
# Schema change support: Adding a column to IDAA

ID	Name

ID	Name

Adding a column to IDAA

ID	Name	Team

ID	Name

Queries referencing the new column can not be offloaded

Table in IDAA can not be loaded anymore

Data changes in DB2 does not get replicated anymore
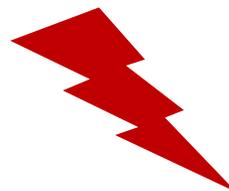
Workaround – as of today

ID	Name	Team

ID	Name

1. Remove the table from the accelerator

2. Re-add the table to the accelerator

3. Re-load the data of the table
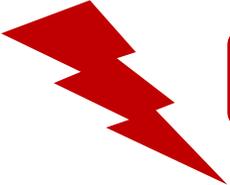
Table outage during the entire process

Workaround HPSS Schema Change

ID	Name	Team

ID	Name

1. Restore data in DB2

2. Remove the table from the accelerator

3. Re-add the table the accelerator

4. Re-the data of the table

5. Re-archive data

Table outage during the entire process

Adding a column to IDAA

Schema Change Support

ID	Name	Team

ID	Name	Team

Schema change support

- After a column has been added/altered to a DB2 table the column can be added/altered at the associated table on the Accelerator by executing a synchronization operation
  - No need any more to remove the table from the Accelerator, re-add it and re-load it again

- Supported table types:
  - Accelerator shadow tables (With or without Incremental Update enabled)
  - Acclelerator archive tables

- Supported Schema Changes
  - ALTER TABLE ADD COLUMN (adding a nullable column)
  - ALTER TABLE ADD COLUMN ... NOT NULL WITH DEFAULT (adding a non-nullable column with a default value)
    - Default values are supported as long as they are explicit. Default values like e.g. TIMESTAMP WITH DEFAULT; are not supported, as there is no way to determine the DB2 side value at the time of the ALTER statement in this case
  - ALTER TABLE ALTER COLUMN <name> SET DATA TYPE (limited to increase length of VARCHAR)

Planned interfaces for the schema change support

- New stored procedure ACCEL_SYNCHRONIZE_SCHEMA to synchronize the schema and alter the table on the Accelerator

- New Accelerator Studio controls:

■ Level-set: What is IDAA and the recent enhancements

■ Big picture: Long-term strategy

■ Work in progress: Top themes

  – Hybrid Transactional/Analytical Processing

 – Schema change support

  – Federated access support

… what works today
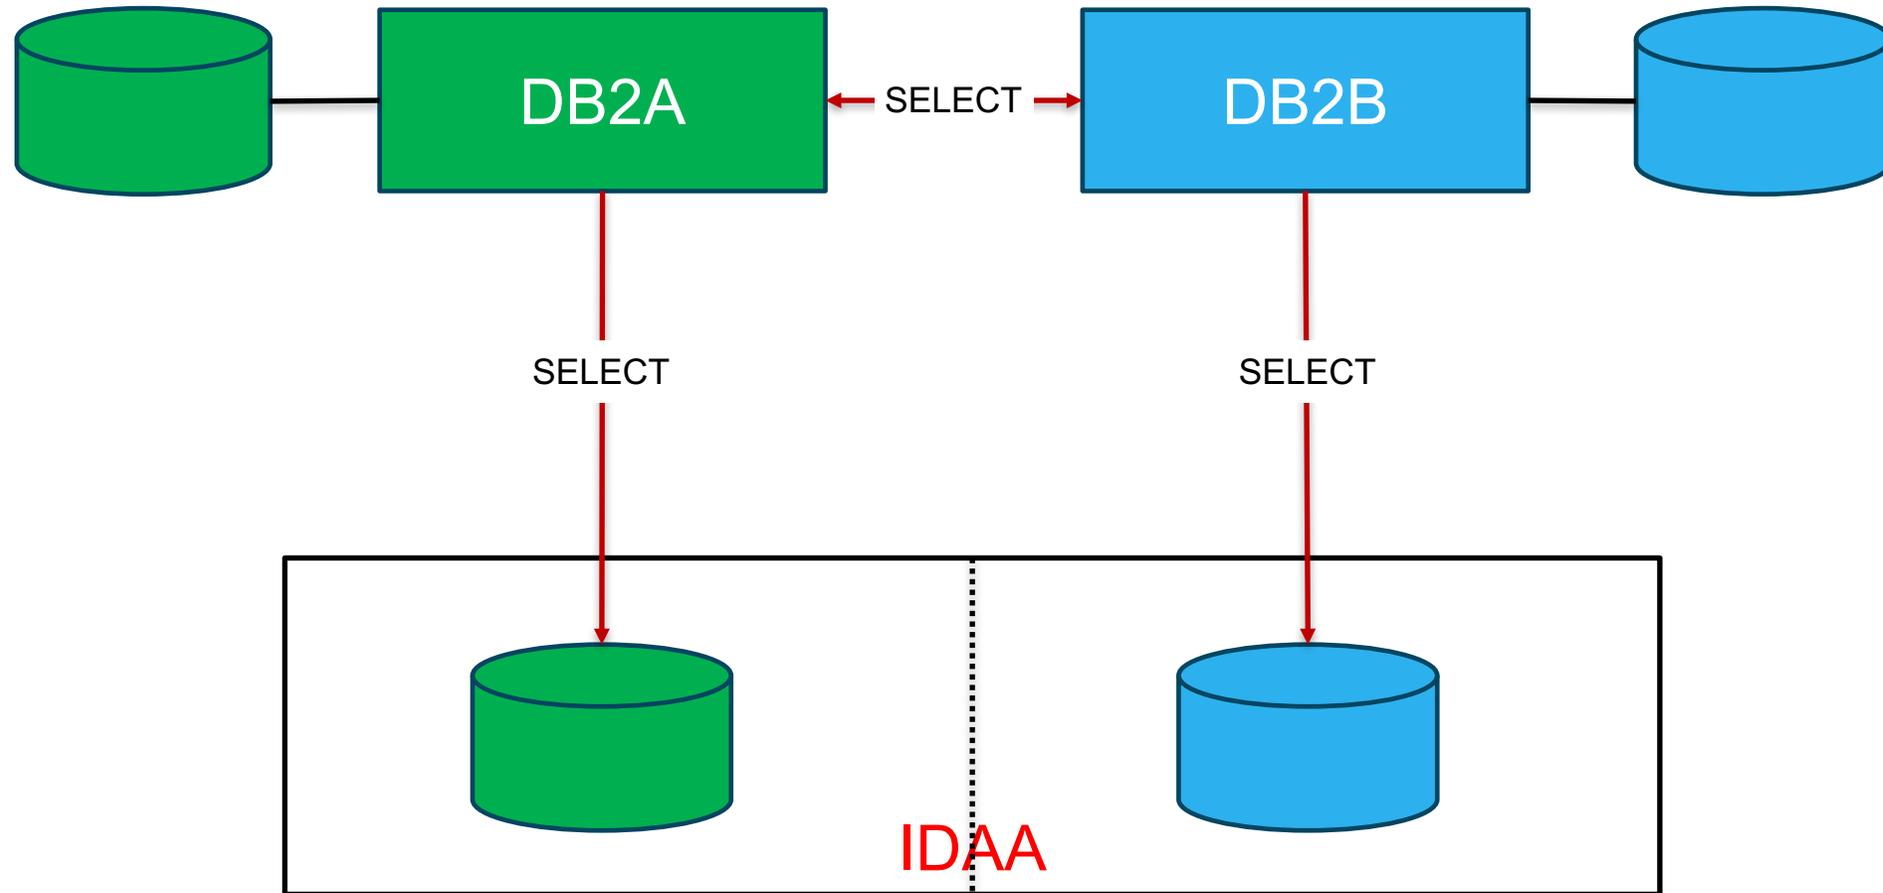
… what works today
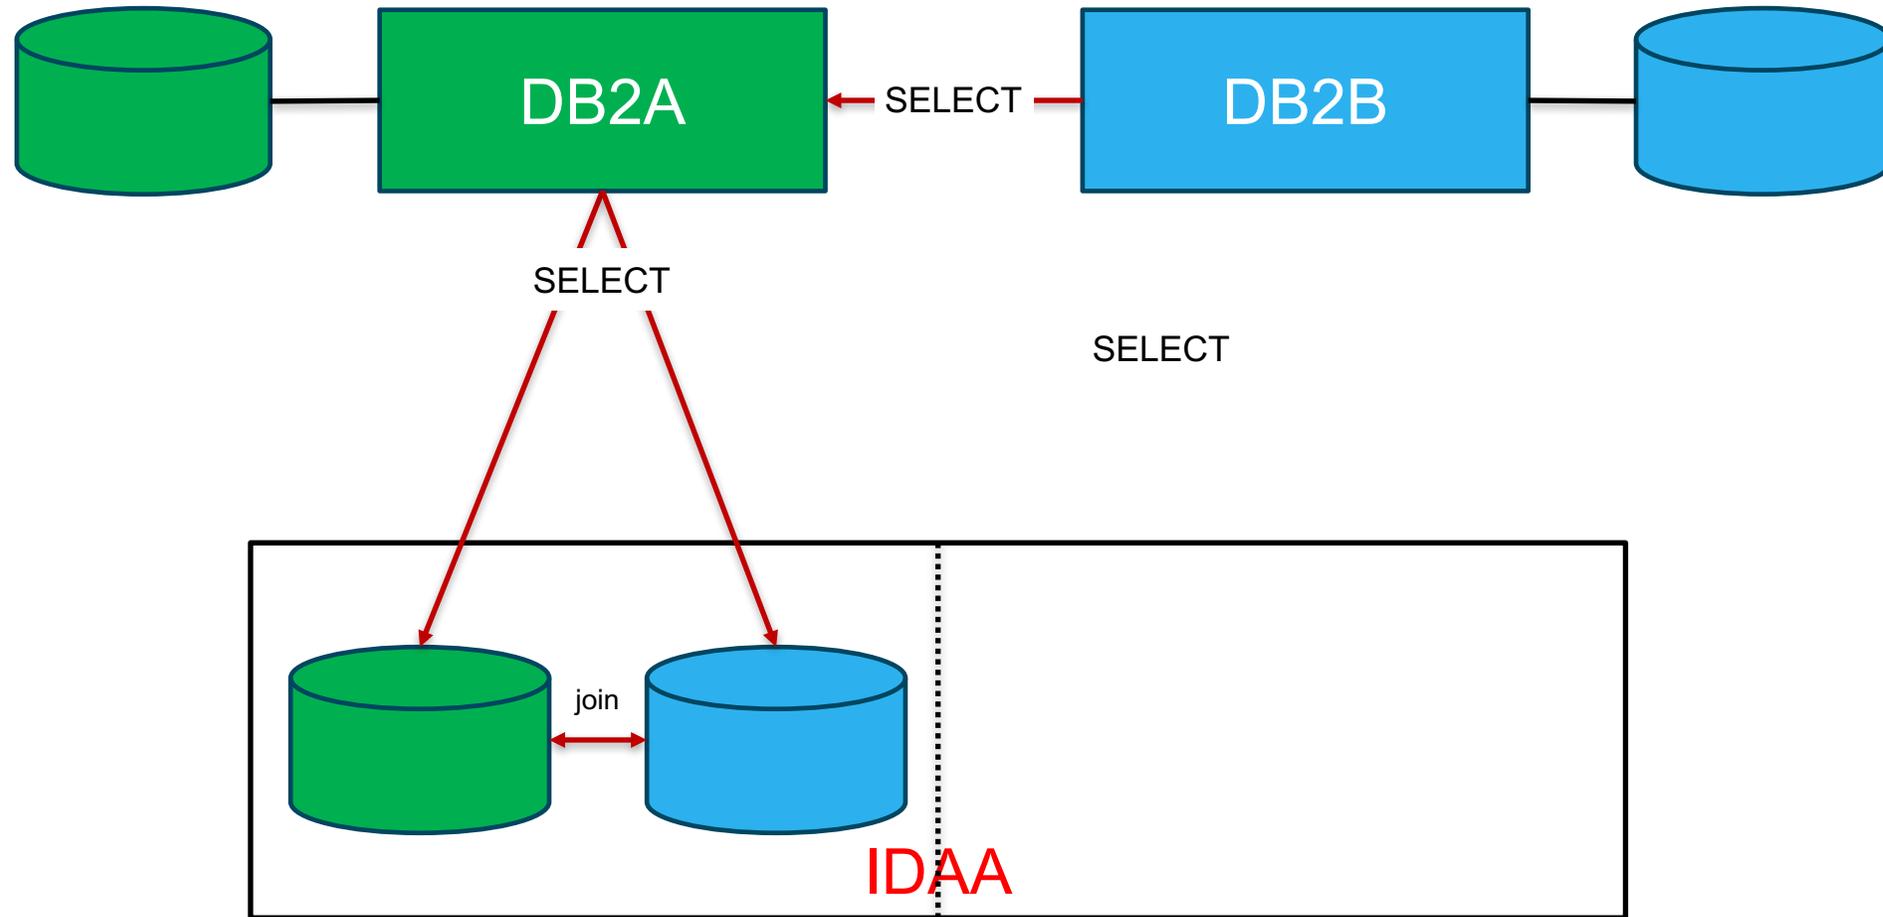
… work in progress
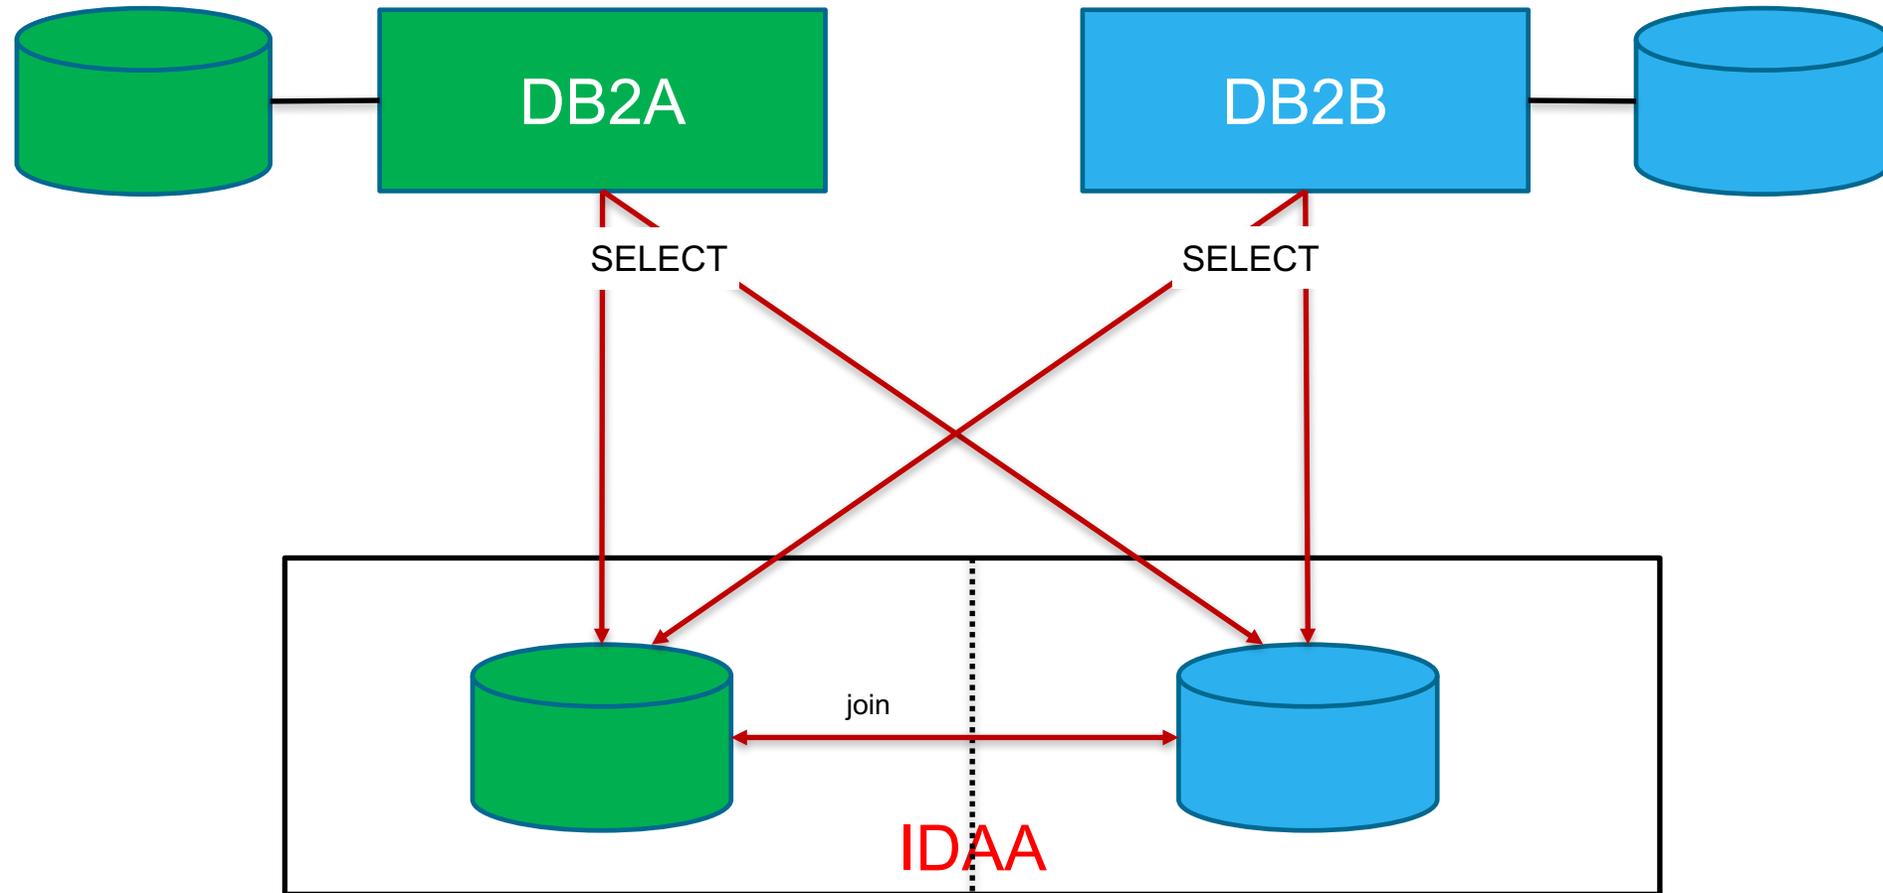
Federated access support – How is it going to work

- **For DB2 owning Subsystem DB2O**:

  – Invoke a new IDAA stored procedure to **GRANT** read access for tables to other referencing DB2 subsystems (not a DB2 GRANT)

    ACCEL_GRANT_TABLES_REFERENCE (ACCEL1, DB2R, Set< acceleratedtables>)

- **For DB2 referencing Subsystem DB2R:**

  – Invoke another new IDAA stored procedure to **CREATE** references to tables of other DB2 subsystems

    ACCEL_CREATE_REFERENCE_TABLES (ACCEL1, DB2O, Set<acceleratedtable-on-db2o>)

  – For each specified table, this stored procedure creates an AOT as a reference table and a synonym that points to the table of the owning subsystem on IDAA

    • No data duplication on IDAA, just references (reference AOT and synonym)

- Afterwards the DB2 referencing Subsystem DB2R can use the created reference AOT in queries and join it with other tables on IDAA (either ‚own tables' or other reference tables)

Thank You