

Leveraging IBM DB2 Analytics Accelerator to Analyze your IMS Data



Agenda

- DB2 Analytics Accelerator for z/OS Overview
- Challenges of Loading to the Analytics Accelerator
- Accelerator Loader – External ‘Dual’ Load
- Loading IMS data to the Accelerator
- General Accelerator Loader Details
- ISPF Interface Examples
- Summary

Analytics Accelerator Customer Usage

Administration ?
 What version ?

Pain Points ?

Do you current have an accelerator ?



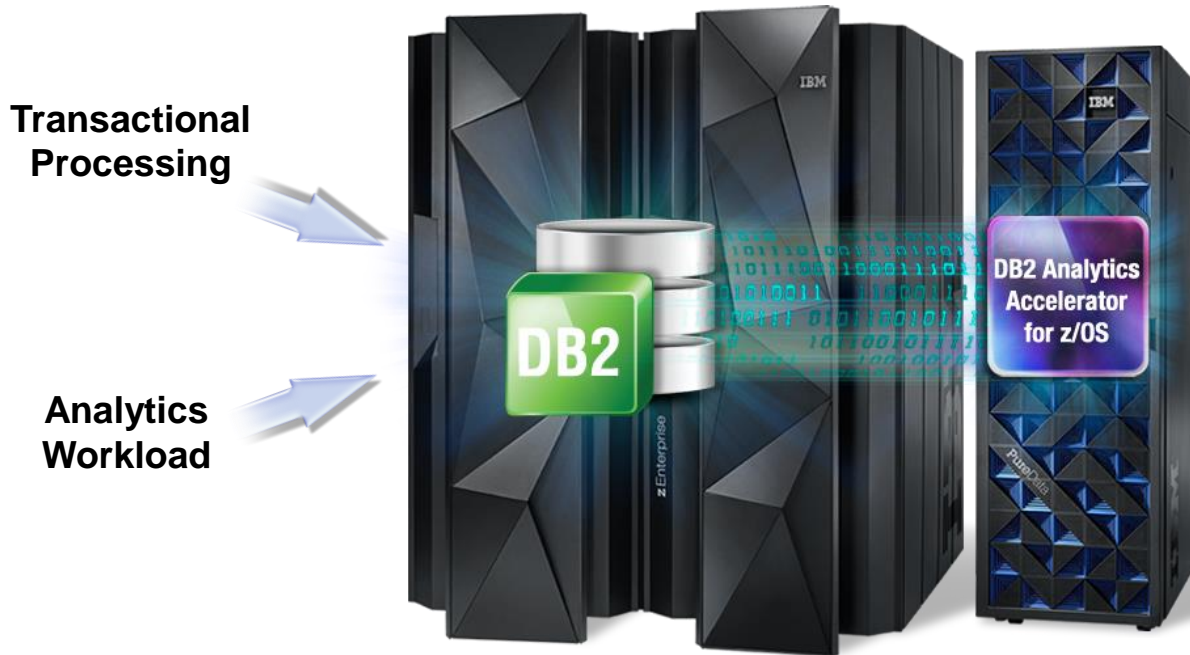
How many ?
 Replication ?

How are you using your accelerators ?

Future accelerator plans ?

IBM zEnterprise and Analytics Accelerator

Driving revolutionary change




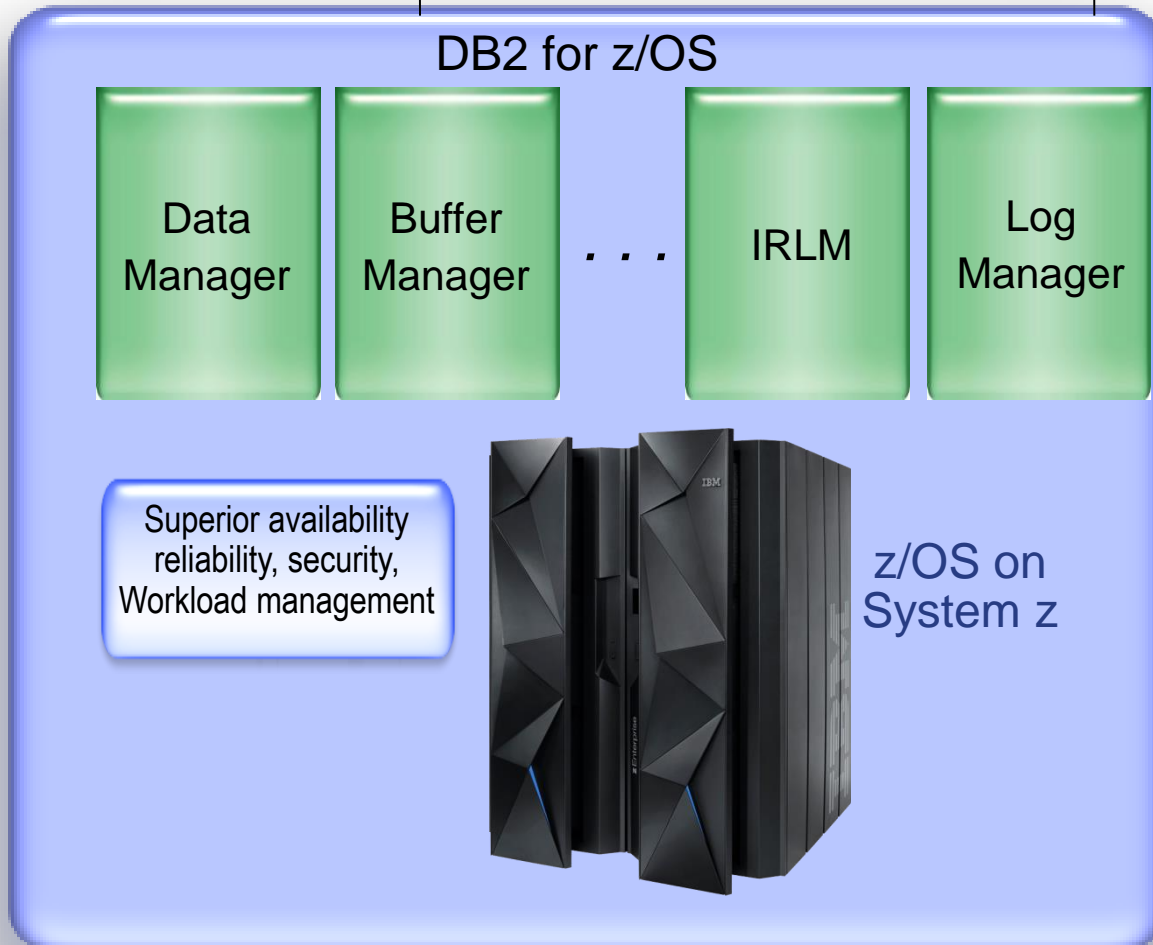

The hybrid computing platform on zEnterprise

- *Supports transaction processing and analytics workloads concurrently, efficiently and cost-effectively*
- *Delivers industry leading performance for mixed workloads*

DB2 Analytics Accelerator and DB2 for z/OS

A self-managing, hybrid workload-optimized database management system that runs query workloads in the most efficient way, so that queries are executed in the optimal environment for greatest performance and cost efficiency

Deep Integration within zEnterprise



DB2 Analytics Accelerator

Further extending the features


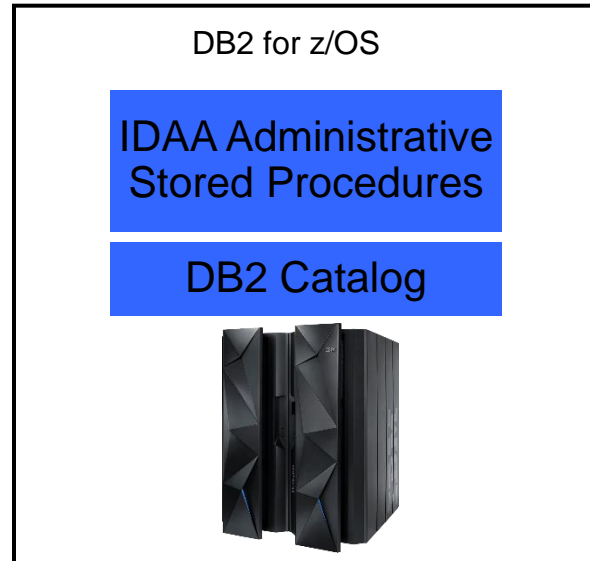


Blending System z and PureData technologies to deliver unparalleled, mixed workload performance for complex analytic business needs.

More insight from your data

- Unprecedented response times for “right-time” analysis
- Complex queries in seconds rather than hours
- Transparent to the application
- Inherits all System z DB2 attributes
- No need to create or maintain indices
- Eliminate query tuning
- Fast deployment and time-to-value

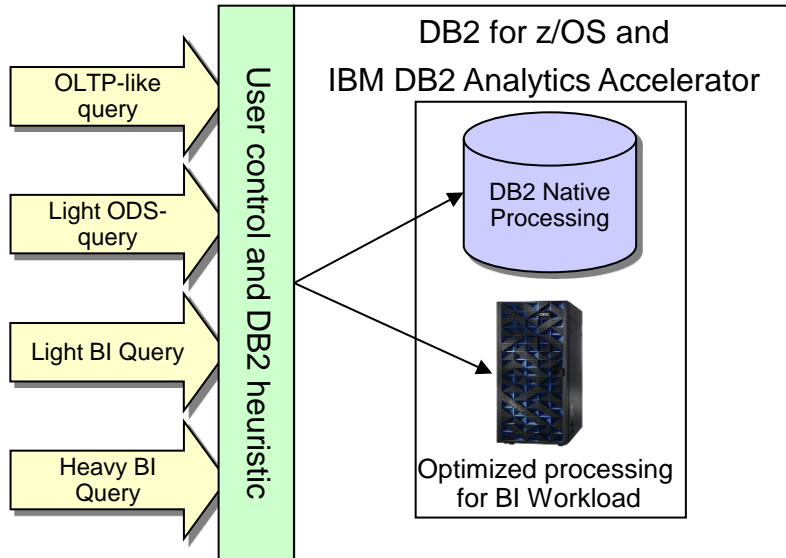
IBM DB2 Analytics Accelerator Table Definition and Deployment



- The tables need to be defined and deployed to IDAA before data is loaded and queries sent to it for processing
 - Definition: identifying tables for which queries need to be accelerated
 - Deployment: making tables known to DB2, i.e. storing table metadata in the DB2 and PureData catalog
- IDAA Studio guides you through the process of defining and deploying tables, as well as invoking other administrative tasks
- IDAA Stored Procedures implement and execute various administrative operations such as table deployment, load and update, and serve as the primary administrative interface to IDAA from the outside world including IDAA Studio

Query Routing Analysis

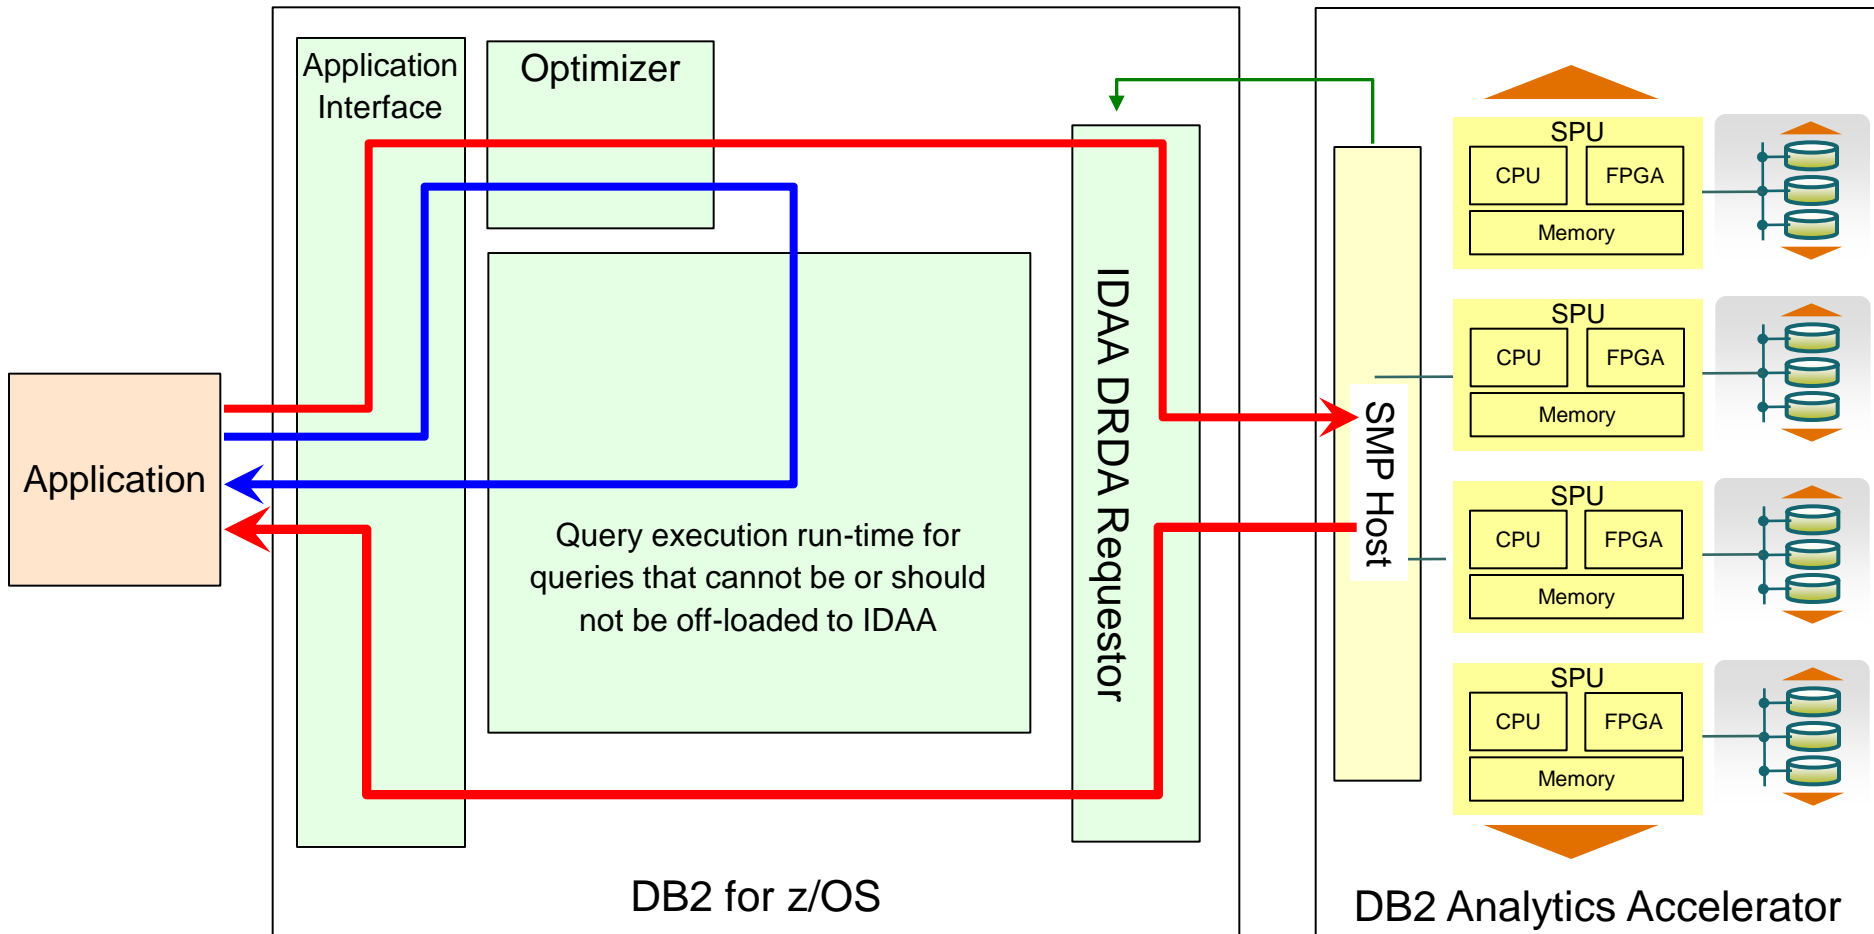
Values for CURRENT QUERY ACCELERATION



- Single and unique system for mixed query workloads
- Dynamic decision for most efficient execution platform
- New special register QUERY ACCELERATION
- New heuristic in DB2 optimizer

Value	Description
NONE	No query is routed to the accelerator
ENABLE	<p>A query is routed to the accelerator if it satisfies the acceleration criteria including the cost and heuristics criteria. Otherwise it is executed in DB2.</p> <p>If there is an accelerator failure while running the query, or the accelerator returns an error, DB2 will return a negative SQL Code to the application</p>
ENABLE WITH FAILBACK	<p>A query is routed to the accelerator if it satisfies the acceleration criteria including the cost and heuristics criteria. Otherwise it is executed in DB2.</p> <p>Under certain conditions the query will run on DB2 after it fails in the accelerator. In particular any negative SQL code will cause failback to DB2 during PREPARE or first OPEN. No failback is possible after a successful OPEN of a query</p>
ALL	A query is routed to the accelerator, if it cannot execute the query fails and a negative return code is passed back to the application
ELIGIBLE	A query is routed to the accelerator if it satisfies the acceleration criteria irrespective of the cost and heuristics criteria. Otherwise it is executed in DB2

Query Execution Process Flow



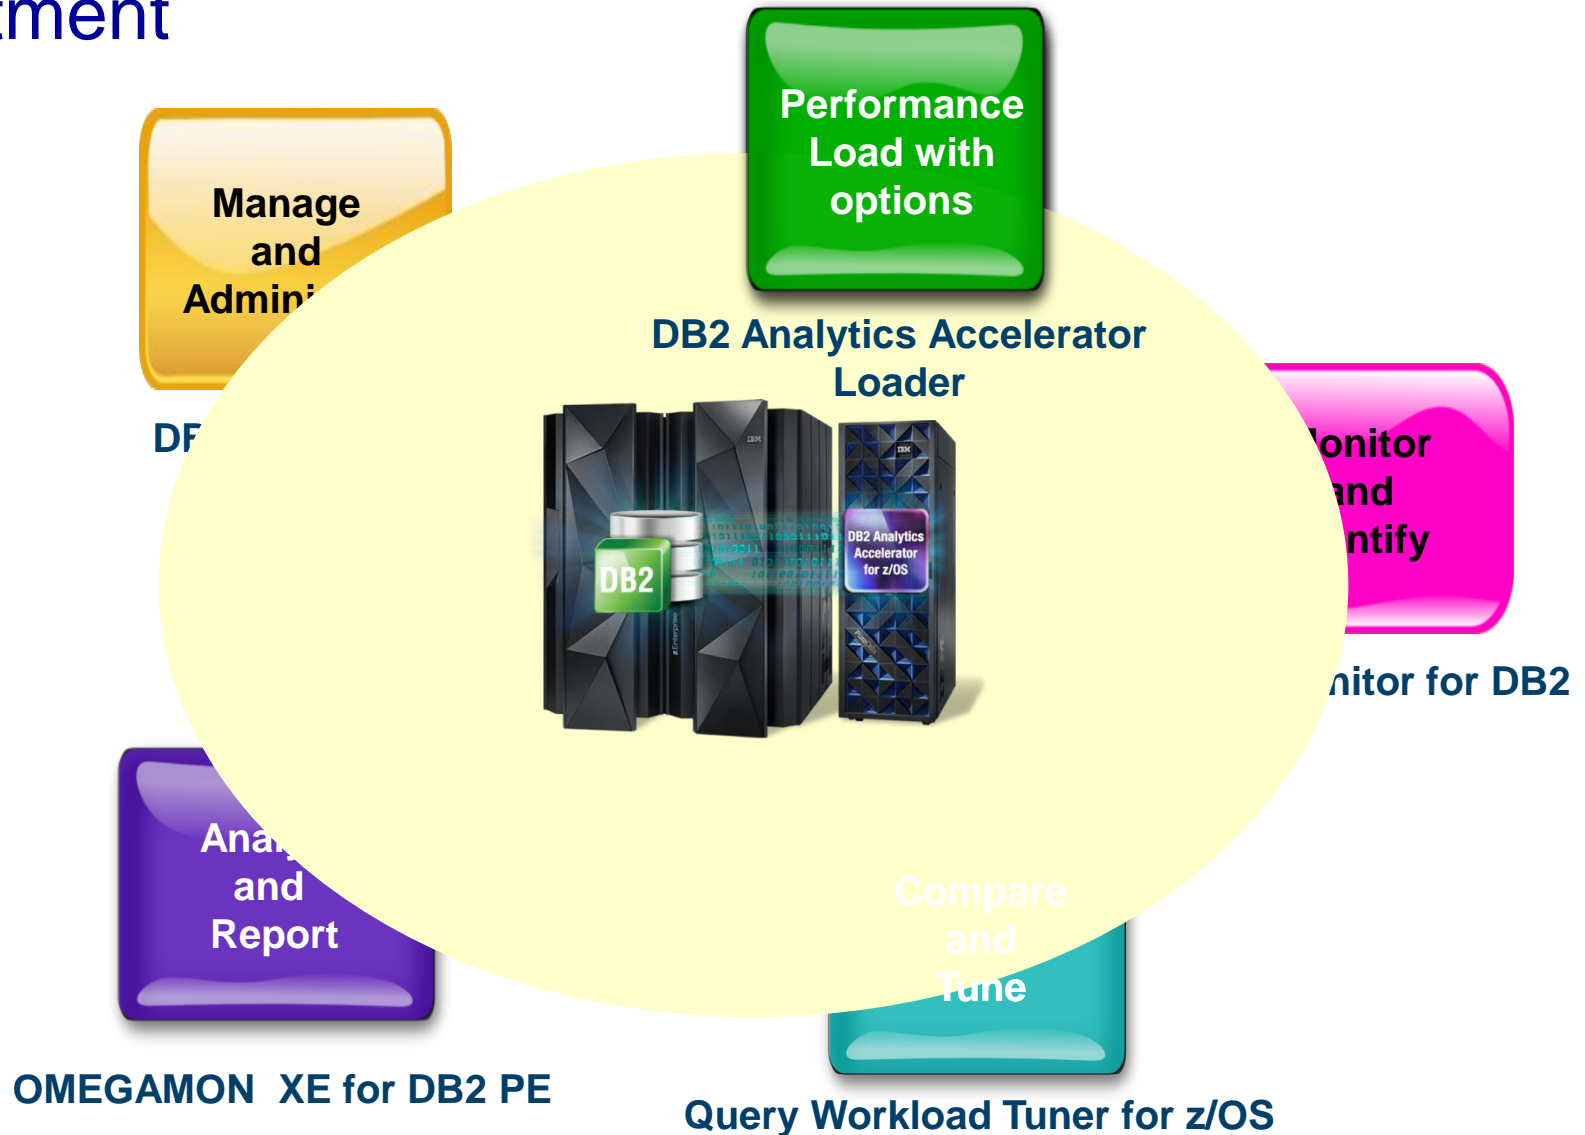
- Heartbeat (DB2 Analytics Accelerator availability and performance indicators)
- Queries executed without DB2 Analytics Accelerator
- Queries executed with DB2 Analytics Accelerator

How IBM Tools Can Maximize Accelerator Value

- Customers want to learn more about their investment in the Accelerator and maximize its use in their environment
 - Customer's are looking at creative ways to exploit the Accelerator....
 - IMS, VSAM, SMF Data, Non-z/OS data
 - Data Mining, IT Analytics, Reporting
- Three different areas where tools can provide value
 - Assessment
 - Do I have a workload that would benefit from the Accelerator?
 - Optimization
 - Can I optimize the workload to take advantage of the Accelerator?
 - Administration
 - Can I manage the Accelerator more effectively?



IBM Tools: Maximizing your Analytics Accelerator Investment



Loading Data into the Accelerator

Tremendous Possibilities

▪ How does data get loaded into the Accelerator today?

- The standard DB2 UNLOAD utility extracts data
- Places tables in Read-only mode until Accelerator load is done (when using option for transactional consistent data)
- Accelerator reflects DB2 data at point in time load was initiated
- Optionally use CDC for near-real-time replication to Accelerator

▪ Introducing the IBM DB2 Accelerator Loader V1.1

- Features
 - Group Consistent Load
 - External 'Dual' Load
- Built for performance & Usability
- Loader populates data in the Accelerator
 - Load DB2 & non-DB2 data
 - Load in parallel to avoid application downtime
 - Load to historical point-in-time

I want to maximize the power of DB2 and System z for business analytics. How do I bring IMS data to the Accelerator for query optimization?

I have to prepare a summary report from my business application as of last Thursday. How do I capture the data from a date that is other than current, that can be considered for query acceleration?

How do I improve the loading of objects into DB2 and/or Accelerator without impacting my DB2 production data?

IBM DB2 Analytics Accelerator Loader: What is External (Dual) Load

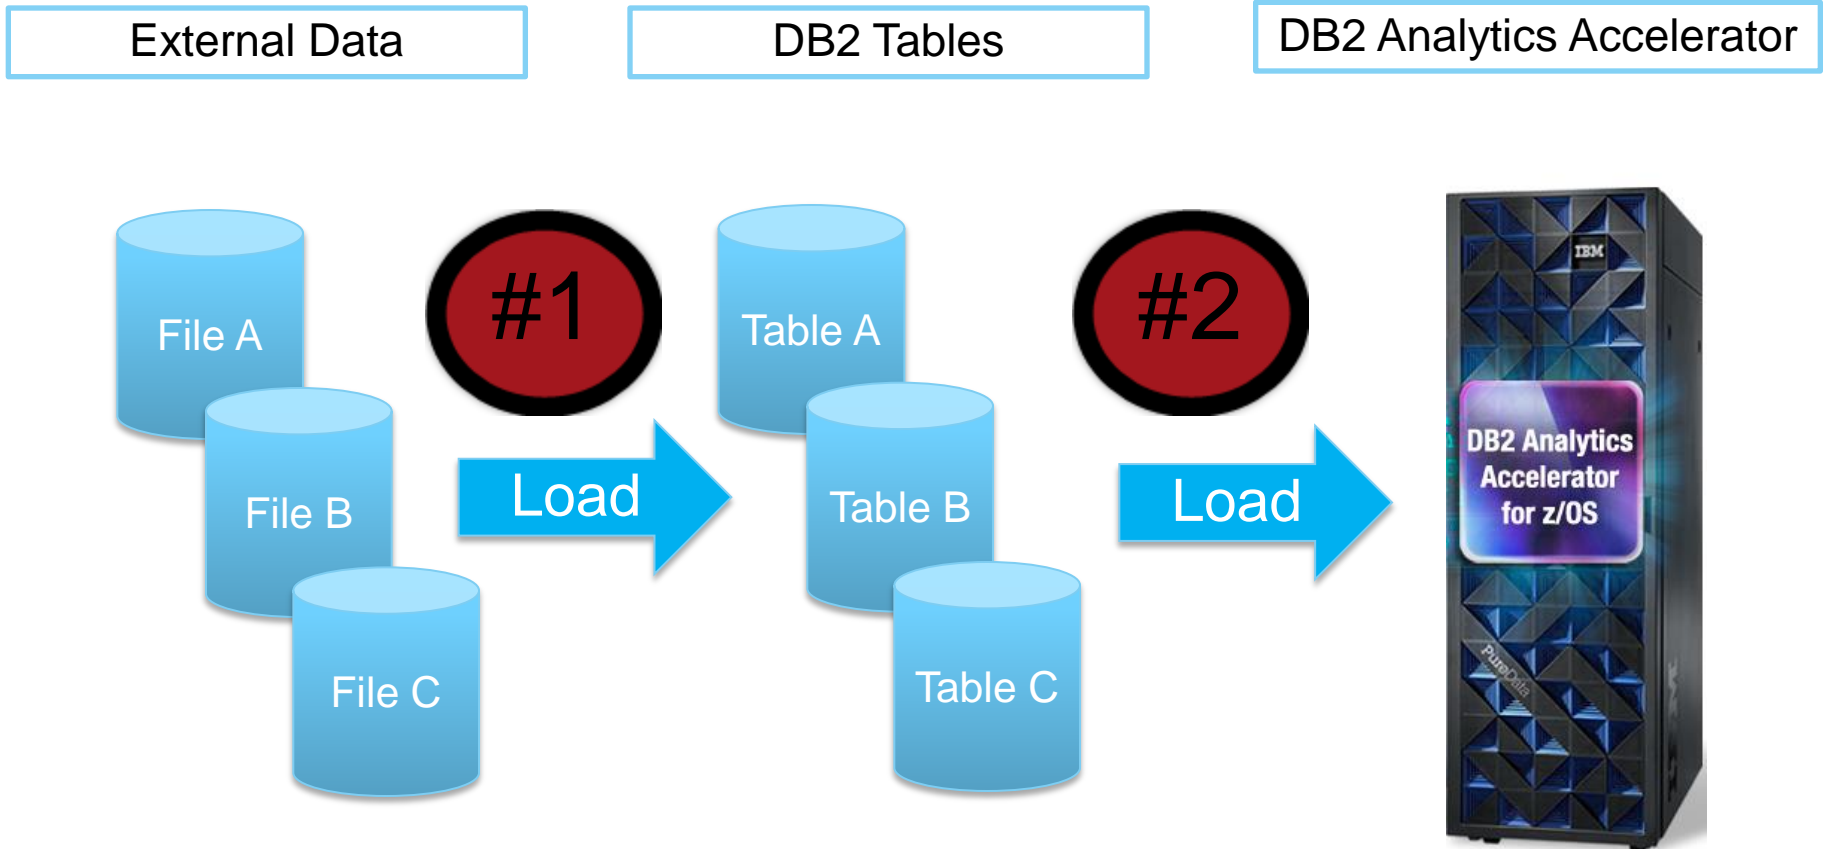
- Accelerator Loader can load data from a file in one of two methods:
 1. Dual External Load
 - Loads data into both DB2 and the Accelerator in parallel
 2. Accelerator Only
 - Accelerator Loader loads directly into Accelerator (no load in DB2)

- User is responsible for building the load file
 - Extracted data can come from various sources
 - IMS, VSAM, Oracle.....etc
 - File must be compatible for input into the DB2 LOAD utility
 - Field specification must describe input data format. This must be compatible with the DB2 LOAD utility.
 - Null-able columns or not null with default can be missing from input data
 - Defaults will be supplied by Accelerator Loader

IBM DB2 Analytics Accelerator Loader: External Load (Dual Load)

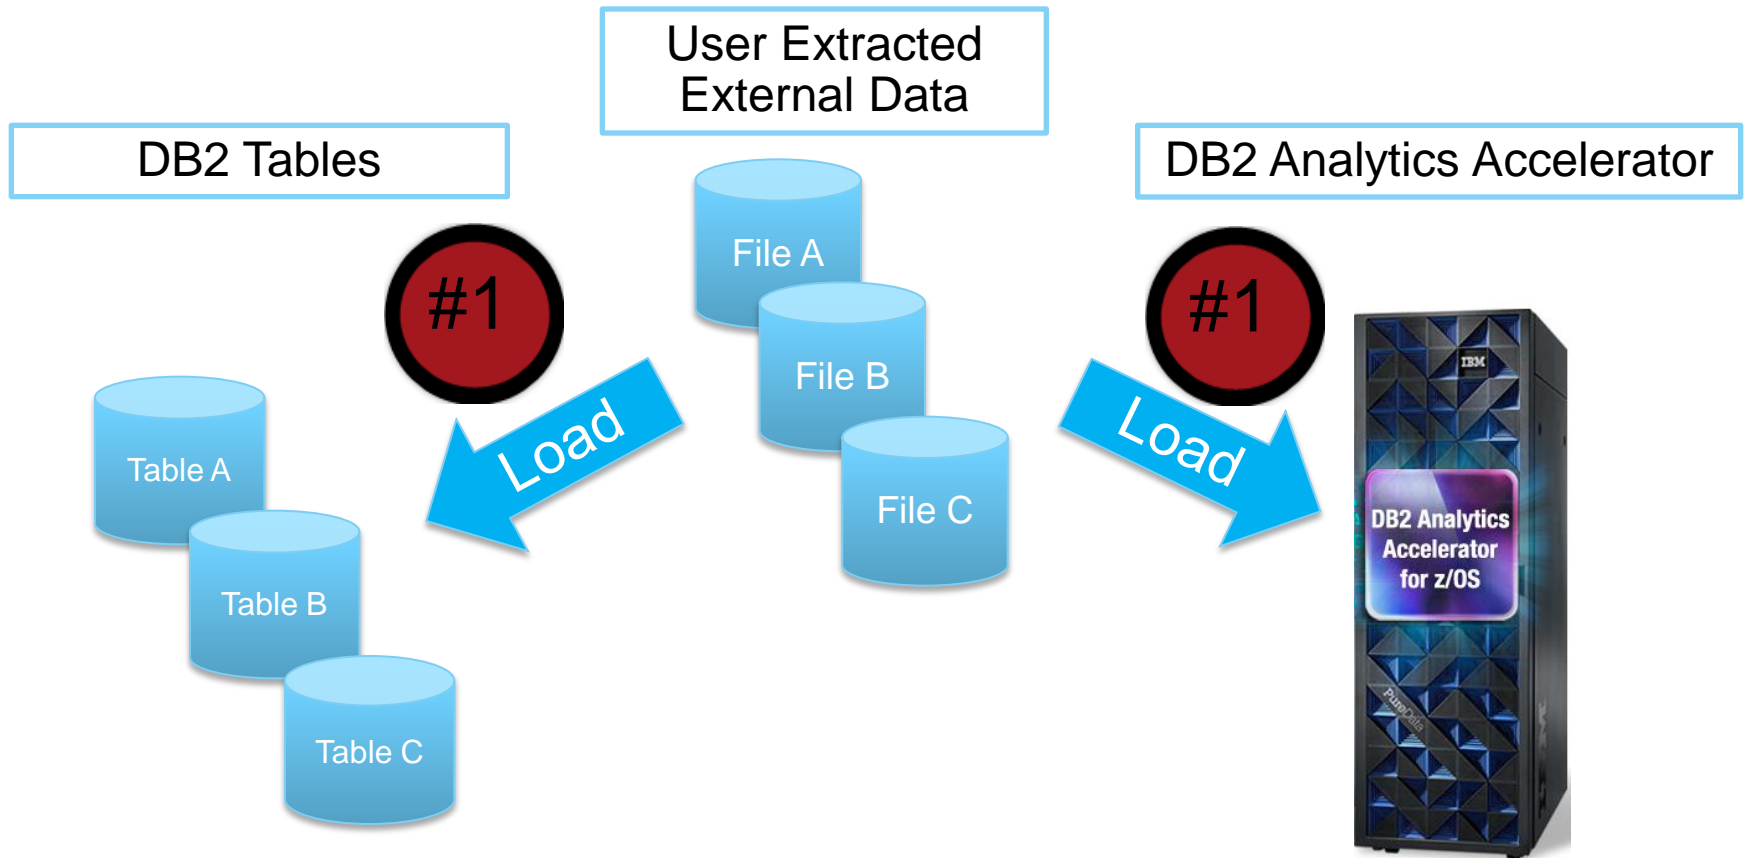
- Loading data external from DB2 into Accelerator
 - Pain Points
 - Data must first be loaded into DB2, then from DB2 into Accelerator
 - 2 Step Process
 - Loading data into DB2 can be CPU intensive
 - No way to load data directly to Accelerator
 - DB2 Accelerator Loader Business Value
 - Improved and simplified process - one step instead of two
 - Exploits zIIP processor to reduce cost of loading to Accelerator
 - Provides significant CPU and DASD savings on the source (DB2)
 - Provides TCO savings with CPU and Elapsed Time improvements
 - Provides path to load accelerator directly from external file

Building Data Warehouse on DB2 z/OS



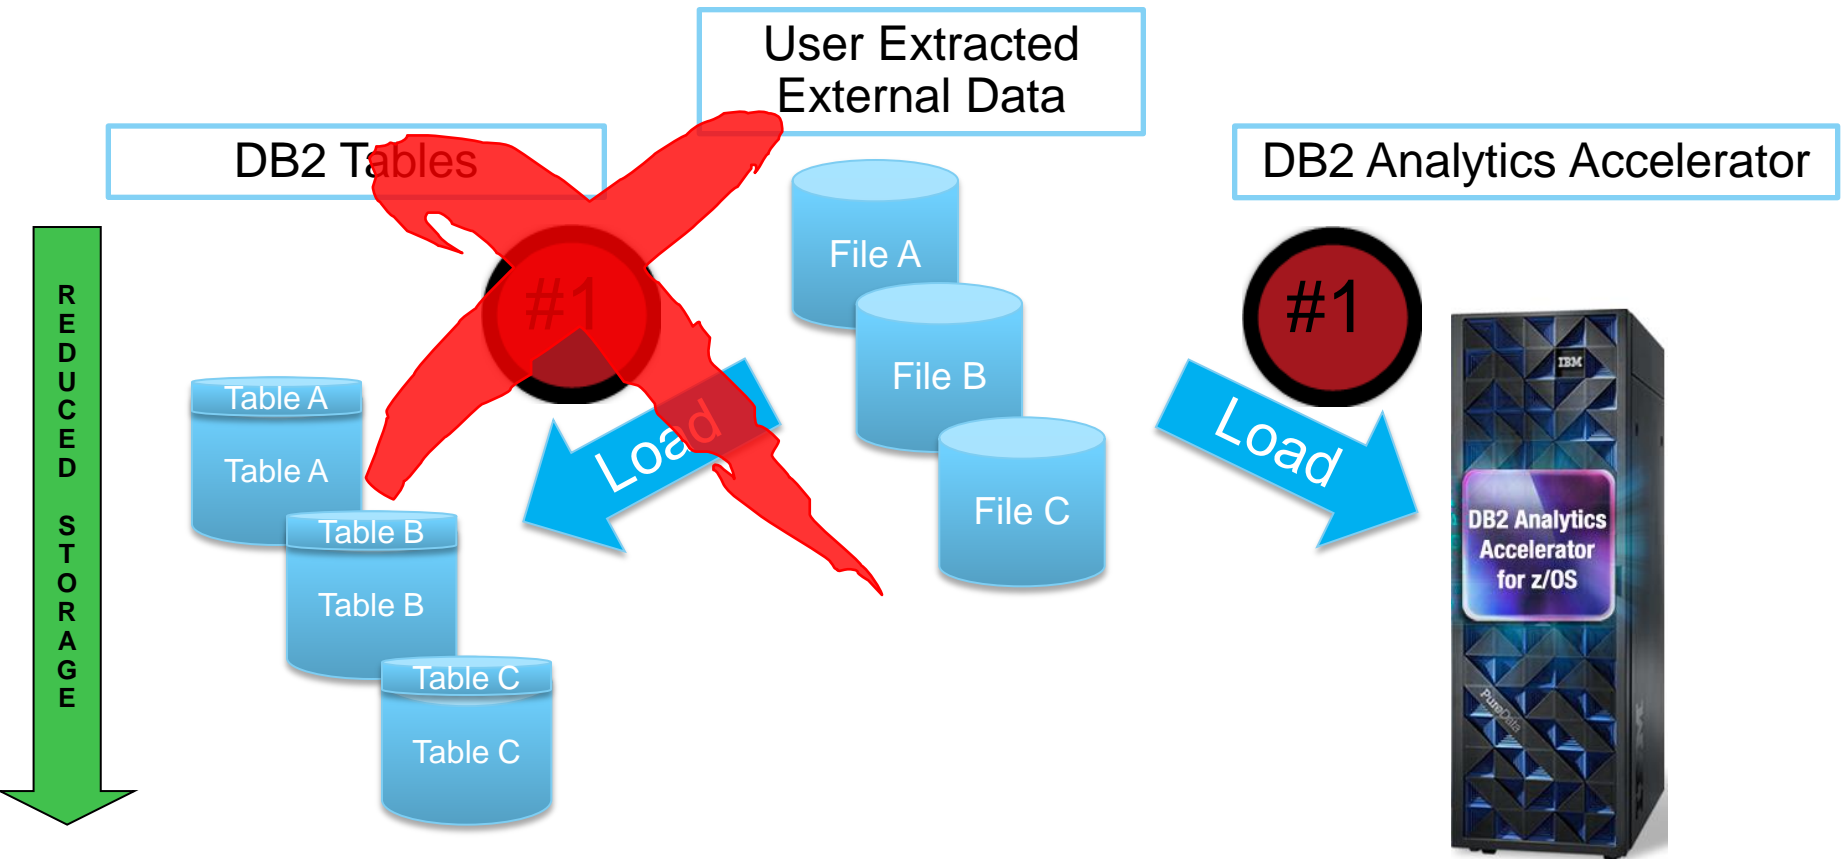
Two Step Load Process – Elongated Load Cycle - CPU Resource Intensive

DB2 Analytics Accelerator Loader: External Load (Dual Load Option)






Parallel Load into DB2 and Accelerator! – Faster Load Cycles! – Reduce Costs!

DB2 Analytics Accelerator Loader: External Load (DB2 Analytics Accelerator Only Option)



IMS to DB2 Analytics Accelerator: External Load Process – Use Case



IMS Segment – DB2 Table Mapping	
<pre> PARTNO 01 0000-0000 02 01-0000-0000 03 01-0000-0000 04 01-0000-0000 05 01-0000-0000 06 01-0000-0000 07 01-0000-0000 08 01-0000-0000 09 01-0000-0000 10 01-0000-0000 </pre>	<pre> Attribute Standard Info Process Code Char(1) Service code Char(1) Cust. Order Char(10) DB Part Number Char(15) </pre>

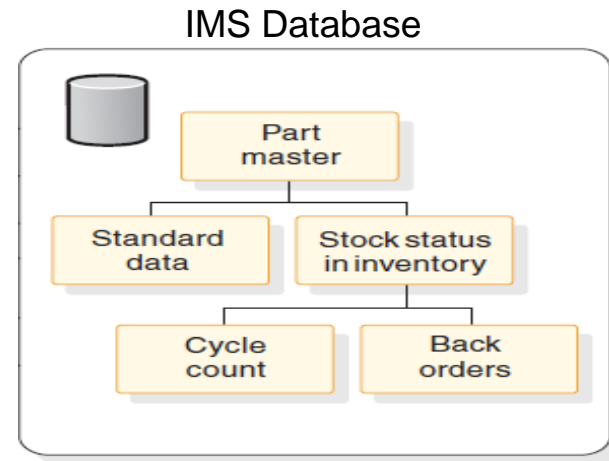
Mapping and Transforming Data

- Segment -> Table
 - Field -> Column

- Data type not required by IMS
 - Many times FIELD only defined for sequence fields
 - Data content not enforced by IMS

- Where are field descriptions defined?
 - IMS Catalog
 - Copy books
 - JAVA Classes

- Non-unique or non-keyed segments



DB2 Tables

&schema.Part_Master

Part_No	Part_Description
---------	------------------

&schema.Standard_Info

Process Code	Invoice Code	Cost Center	PM_Part_no
--------------	--------------	-------------	------------

&schema.Stock_Status


Area	Dept	Project	Division	PM_Part_no
------	------	---------	----------	------------

⋮


Flattening IMS Database Records

- Concatenated Keys
 - Concatenated key fields not stored with segment data
 - Key fields needed for each row to maintain referential integrity

IMS Segment – Part Master



IMS Segment – Standard Data



- OCCURS clauses
 - Multiple instances of a field in a single instance of a segment
 - Multiple 'rows' should be generated

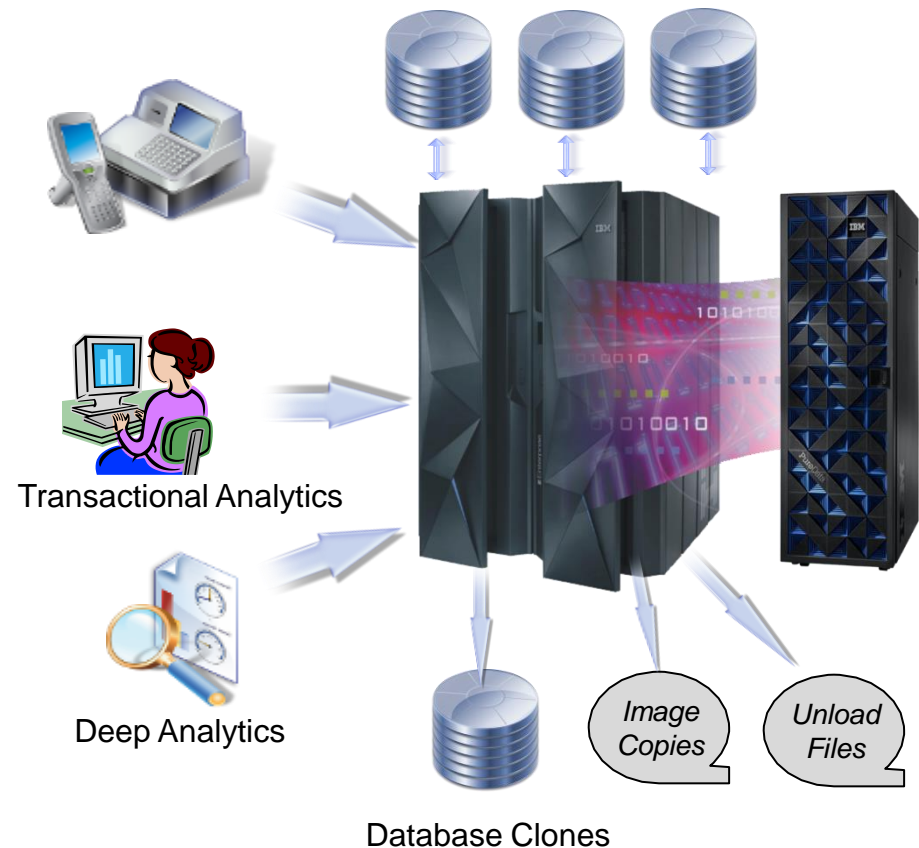
```

STANINFO (Standard data)
01 STAN-INFO.
 02 SI-PROC-CODE  PIC XX.  <-Key
 02 SI-INV-CODE   PIC X.
 02 SI-COST-CTRS  PIC X(30).
 02 SI-COST-CTRS-D REDEFINES SI-COST-CTRS.
 05 SI-COST-CTR-NO OCCURS 3 TIMES
                    PIC X(10).
          
```

Extraction Considerations and Methods

- Considerations
 - Availability requirements
 - Frequency
 - Impact to OLTP workload
 - What data is needed?
 - Entire database record?
 - Certain segments?
 - Data from multiple databases
 - Consistent point in time

- Extraction Methods
 - IMS Application
 - Additional online workload
 - Data can still be changing
 - Database Clone (IMS Cloning Tool)
 - Group of databases at a point in time
 - DLI applications
 - Image Copies/Unload Files
 - Additional knowledge of data set structure needed



Transforming Non-Relational Data

- Transformation may need to be done for each field of each segment
 - Data types
 - Flattening
 - OCCURs
- Transformation is typically most CPU intensive portion of ETL
 - Analysis and operation performed on smallest entity
- Accelerator Loader inputs
 - DB2 unload file format
 - All fields/columns are in native DB2 data type
 - Data types can be described to Loader
 - Loader will transform data to DB2 format during load

Accelerator Loader Performance Benefits

- Data Conversion
 - External to Internal Data conversion is zIIP enabled
 - Faster DB2 load utility (DSNUTILB) due to Internal Data
- Dual Load – “Double Load”
 - DB2 and DB2 Analytics Accelerator Loaded in Parallel
 - Input SYSREC is read once
- Performance Estimates
 - DB2 Load/Native DB2 Analytics Accelerator Load compared to Accelerator Loader - External Load
 - Up to 55% reduction in elapsed time
 - Up to 35% reduction in CPU
 - Mileage may vary
- DB2 Analytics Accelerator ONLY LOAD
 - IDAA_ONLY ON *accelerator_name*
 - Up to 60% general purpose CPU reduction
 - DB2 Storage savings

External Load: Internal Details

1. Users submit DB2 Load JCL with additional directive
 - Add IDAA_DUAL ON *accelerator_name*
 - Or IDAA_ONLY ON *accelerator_name*
 - Add ACCEL_LOAD_TASKS for partition parallelism
 - Add //HLODUMMY DD *
 - Can be added into existing JCL or created via ISPF profiles

2. Loader Intercepts DB2 Load Utility
 - Sees additional syntax (above)

3. Loader reads SYSREC (input file)
 - Converts data to DB2 Internal Format (Under zIIP)
 - Changes DB2 Load to 'Internal Format' (runs faster)

4. Loader sends internal format data to Accelerator and DB2 Load Utility or Accelerator Only

External Load Accelerator-Only Considerations

- When should user consider loading accelerator only?
 - Data is maintained and updated elsewhere
 - DB2 is not required for data backup and recovery
 - Data Validations
 - Load Utility is not run - No RI or Constraint Checking
 - All queries are qualified for accelerations
 - Not unsupported data types such as LOB or XML
 - If DB2 executes query locally
 - Initial Updates – Fail/SQL Code 100
 - 0 rows in DB2 table after load
 - Inserts – Will execute
 - Initial Deletes – Fail/SQL Code 100
 - 0 rows in DB2 table after load

External Load 'Accelerator-Only' Considerations

- Table must still exist in DB2 Catalog
 - Will be emptied upon load
- Users should define small tablespaces
 - Prevent large datasets sitting unused
- Access to accelerated table remains via DB2
 - All DB2 security is honored

- DB2 Optimizer must send queries to accelerator
 - Any query accessing DB2 table will have inconsistent results
- User Beware: Minor application change required
 - Special register directs all queries to Accelerator
 - **SET CURRENT QUERY ACCELERATION = ALL**
 - Can be set in JDBC/ODBC connections
 - Future change coming in DB2 for Accelerator Only tables

DB2 Analytics Accelerator Loader: Group Consistent Load

- What is Group Consistent Load?
 - Loads groups (or sets) of operational DB2 tables to the accelerator at a user-specified time. One time specified for all tables.
 - Uses DB2 Image Copies and DB2 Logs as input
 - Transaction Consistent: Uncommitted transactions at the specified time are not loaded to accelerator
 - Ex: Update Parent Table, Update Child Table, Commit
 - If Load is run after parent update but before child update, the update to parent table will not be loaded to accelerator
 - **No tables locked during consistent load process!**
 - Usability feature not a performance feature



DB2 Analytics Accelerator Loader: Group Consistent Load

Loading data to Accelerator from Operational DB2 Tables

- Possible pain points with existing load process:
 - Loading related tables requires taking DB2 tables offline for update
 - No method to load historical copy of tables into Accelerator

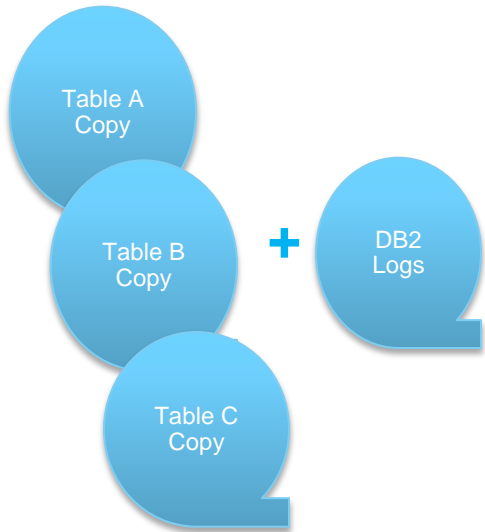
- DB2 Accelerator Loader business value:
 - Maintain availability of related tables during the load process
 - Load or refresh accelerator with zero impact to business critical data
 - Ability to load to accelerator with historical data
 - Improved availability and flexibility!

IBM DB2 Analytics Accelerator Loader

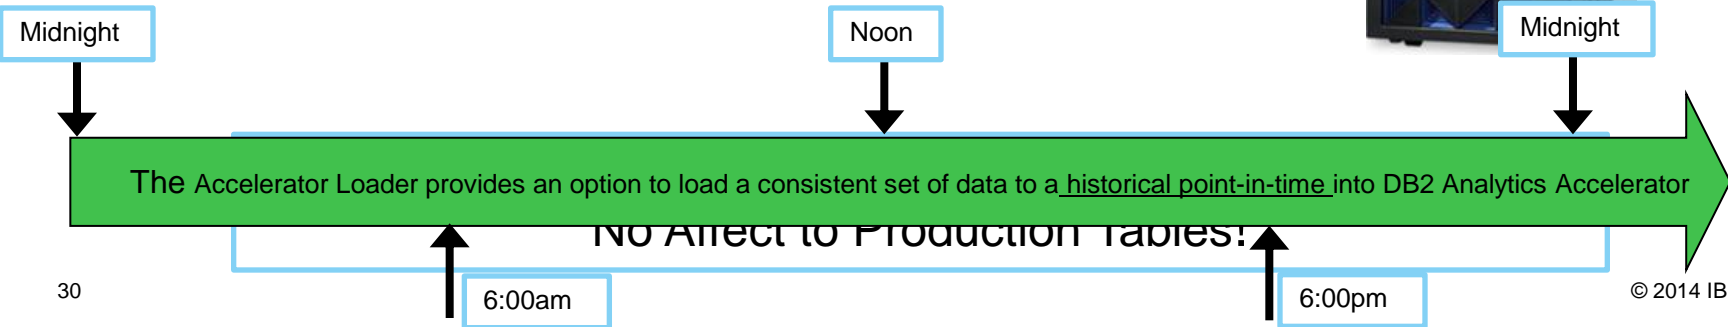
Group Consistent Load

Backups & DB2 Logs

DB2 Analytics Accelerator



Load Process



Group Consistent Load: Internal Details

1. Users specify point in time to load
2. Loader selects appropriate image copies before point in time
 - Supports full image copies
 - Incremental image copies
 - Inline image copies
 - FlashCopy image copies
3. Reads and merges image copies
4. Reads log records from image copy point to specified PIT
5. Sorts log records in Page/Time sequence
6. Applies sorted log records to image copy data page
7. Extracts table rows from page
 - Decompresses rows
8. Writes table rows to accelerator through USS pipe
 - Stored procedure ACCEL_LOAD_TABLES is used to open the pipe to the accelerator

Group Consistent Load: Considerations

- Direct load from image copies/logs from other DB2 systems not currently supported
 - OLTP → DB2 Analytics Accelerator
 - Targeting second half of 2014

- Currently does not integrate with CDC
 - Known requirement for future support
 - Targeting second half of 2014
 - Will allow loading replicated tables without locking

- Table alters currently require a post DB2 Reorg
 - Table definition is altered, a reorg is required before running consistent load

- No support today for multi-table tablespaces
 - Most tablespaces are single table
 - Can be added if market demands support

External 'Dual' Load or Group Consistent Load?

- External 'Dual' Load
 - When data is in a file
 - Users desire to load file into DB2, Accelerator, or both
 - Can be extracted from DB2 or other sources as input to the Loader

- Group Consistent Load
 - Run when data is already in DB2
 - When loading multiple related groups of tables
 - When customers require related data to be consistent

Profiles



Loader Profiles

```

LOADER                               Profile Display                               2014/04/14  14:12:44
Command ==> _____ scroll ==> PAGE

Commands:      CREATE
Line commands: B - Build  D - Delete  E - Edit  R - Rename  V - View  C - Copy

Profile like . . DEMO*
Creator like . . *
Profile type . . ALL (DUAL, CONSISTENT, or ALL)

Row 1 of 2  >

Cmd Name      Creator  Type      Share  Description
----->----->----->----->----->
DEMO GROUP CONSISTEN TSAXC    Consistent UPDATE  DEMO PROFILE GROUP CON
DEMO EXTERNAL LOAD  TSAXC    Dual     UPDATE  DEMO PROFILE EXTERNAL
***** Bottom of data *****

```

- Profiles are collections of Loader options and DB2 tables to load
- Can be used to build JCL to run Loader functions
- Profile/Load types:
 - Group Consistent Load
 - External (Dual) Load

Profiles – Group Consistent Load JCL Example

```
//SYSINHLO DD *  
  IDAA CONSISTENT LOAD  
  (  
    TEMPLATE  
    (  
      NAME TEMPSAMP  
      DSN '&US..&DB..&SN..&SS..T&TI.'  
    )  
  )  
  GROUP  
  (  
    SPACE  
    (  
      CREATOR 'USEIDAA'  
      NAME 'GLWTDPT'  
    )  
    SPACE  
    (  
      CREATOR 'USEIDAA'  
      NAME 'GLWTEMP'  
      PARTITION 1  
    )  
    SPACE  
    (  
      CREATOR 'USEIDAA'  
      NAME 'GLWTEMP'  
      PARTITION 2  
    )  
    TO_CURRENT  
    NEW_COPY  
    FCCOPYDDN TEMPSAMP  
  )  
  ACCELNAME QDS5ACC1  
  PARALLEL '4,4'  
  BUFFERS_IN_31_BIT  
  LOG_COPY_PREFERENCE R1R2A1A2  
  USER_INDICATOR HLO  
)
```

Profiles – Dual Load JCL Example

```

//SYSIN DD *
TEMPLATE ISYSREC
  DSN &US..IDSD.&DB..&TS..P&PA.
  DISP(SHR,KEEP,KEEP)
LOAD DATA INDDN
IDAA_DUAL ON QDS5ACC1
ACCEL_LOAD_TASKS 4
REPLACE
KEEPDICTIONARY
LOG NO
SORTDEVT SYSALLDA SORTNUM 10
ERRDDN ISYSERR
MAPDDN ISYSMAP
WORKDDN (ISYSUT1, ISORTOUT)
INTO TABLE
  "TSAXC"."RANDOM_TABLE"
  PART 1
  INDDN ISYSREC
  (
    PART_ID POSITION(1) INTEGER EXTERNAL(8),
    PART_KEY POSITION(9) INTEGER EXTERNAL(8),
  )
INTO TABLE
  "TSAXC"."RANDOM_TABLE"
  PART 2
  INDDN ISYSREC
  (
    PART_ID POSITION(1) INTEGER EXTERNAL(8),
    PART_KEY POSITION(9) INTEGER EXTERNAL(8),
  )
/*

```

- DB2 Analytics Accelerator_DUAL directive tells Loader to load both DB2 and Accelerator
- DB2 Analytics Accelerator_ONLY directive tells Loader to load only Accelerator
 - NOTE! DB2 table data will be deleted on a Accelerator-only run

Summary

- DB2 Analytics Accelerator Loader
 - Improves business availability
 - Helps facilitate greater Accelerator usage
 - Facilitates loading external source data
 - Facilitates the power of querying/joining data from different sources
 - Reduces cost of loading data into Accelerator
 - Reduces CPU by exploiting zIIP
 - Loads to the Accelerator and DB2 in parallel
 - Reduces DASD requirements
 - Simplifies the process of loading data to the Accelerator

Questions



Thank
YOU

The image features the words "Thank YOU" in a large, 3D, light blue font. Each letter of the word "Thank" contains a different person's face, and the word "YOU" also contains three different faces. The faces are diverse in age and ethnicity. The letters have a slight shadow, giving them a three-dimensional appearance.

Appendix



Profiles



Loader Profiles

```

LOADER                               Profile Display                               2014/04/14 14:12:44
Command ==> _____ scroll ==> PAGE

Commands:      CREATE
Line commands: B - Build  D - Delete  E - Edit  R - Rename  V - View  C - Copy

Profile like . . DEMO*
Creator like . . *
Profile type . . ALL (DUAL, CONSISTENT, or ALL)
                                                    Row 1 of 2  >

Cmd Name          Creator  Type          Share
----->         ----->  ----->     ----->
_____ DEMO GROUP CONSISTEN TSAXC    Consistent  UPDATE      DEMO PROFILE GROUP CON
_____ DEMO EXTERNAL LOAD   TSAXC      Dual        UPDATE      DEMO PROFILE EXTERNAL
***** Bottom of data *****

```

- Profiles are collections of Loader options and DB2 tables to load
- Can be used to build JCL to run Loader functions
- Profile/Load types:
 - Group Consistent Load
 - External (Dual) Load

Consistent Load Options

```

LOADER                               Consistent Load Options                2014/04/14  14:14:33
Command ==>> _____

Commands: TABLES - Edit tables list  ACCELERATOR - select accelerator

Creator      . . . : TSAXC              Name      . . . : DEMO GROUP CONSISTENT LOAD
Share option . . UPDATE                Description . . DEMO PROFILE GROUP CONSISTENT  >

Utility processing options:
Accelerator name . . . . . QDS5ACC1
Load time      . . . . . CURRENT      (CURRENT, SPECIFIED, or QUIESCE)
RBA or LRSN end point . . . . . _____ (hexadecimal value)
Timestamp end point . . . . . _____ (YYYY-MM-DD-hh.mm.ss.nnnnnn)
Time zone of timestamp . . . . . LOCAL (LOCAL or GMT)
Quiesce end point . . . . . _____ (1-999)
Continue on errors . . . . . NO      (Yes/No)

FlashCopy options:
Use FlashCopy . . . . . YES (Yes/No) Only valid for current
Use FlashCopy DSN template . . . . . YES (Yes/No) Update NO (Yes/No)

Log read and log apply options:
SYSCOPY scan operating mode . . . . . ZPARAM (LOCAL, RECOVER, ZPARAM, or USER)
SYSCOPY selection preference . . . . . LPLBRPRBFC (LPLBRPRBFC in any order)
Log reader copy preference . . . . . R1R2A1A2 (R1R2A1A2 in any order)
Buffers in 31 bit storage . . . . . YES (Yes/No)
Number of PARALLEL log read . . . . . 4 (0-16 tasks)
Number of PARALLEL log apply . . . . . 4 (1-10 tasks)
  
```

- Choose Load Time
- Optional FlashCopy Options
- Log Read / Apply Options

Profiles – Consistent Load Select Tables

```

LOADER                               DB2 Table List                               2014/04/14  14:19:45
Command ==>>                          scroll ==>> PAGE

Commands:      ADD - Add table(s)
Line commands: D - Delete table  RIS - select RI tables
               RIA - Add all RI tables

Creator . . . : TSAXC                Name . . . . : DEMO GROUP CONSISTENT LOAD
Share option . : UPDATE              Description . : DEMO PROFILE GROUP CONSISTENT

                                Row 1 of 8
Cmd Table Name                  Part  Creator                  Database Tablespace
----->----->----->----->----->----->----->----->----->----->
___  GLWTDPT                      N/A   USEIDAA                    USEIDAA  GLWSDPT
___  GLWTEPA                      N/A   USEIDAA                    USEIDAA  GLWSEPA
___  GLWTPJA                      N/A   USEIDAA                    USEIDAA  GLWSPJA
___  GLWTPRJ                      N/A   USEIDAA                    USEIDAA  GLWSPRJ
___  GLWTEMP                      1     USEIDAA                    USEIDAA  GLWSEMP
___  GLWTEMP                      2     USEIDAA                    USEIDAA  GLWSEMP
___  GLWTEMP                      3     USEIDAA                    USEIDAA  GLWSEMP
___  GLWTEMP                      4     USEIDAA                    USEIDAA  GLWSEMP
***** Bottom of data *****

```

- Specify all the tables to be loader to consistent point in time
- Masking is supported to select tables
- Automatically add RI-related tables

Profiles – Consistent Load Build JCL

```

LOADER          Build Accelerator Loader JCL          2014/04/14 14:16:07
Command ==>>> _____ Scroll ==>>> PAGE

Commands:      ADD - Add job card line
Line commands: D - Delete job card line

Press F3 to build JCL or F12 to cancel.

Generated JCL data set name:
  Data set name . . TSAXC.HLO.JCLLIB
  Member name . . . GCONSIST (if data set is partitioned)

Processing options
Enter "/" to select option
  _ Specify new data set allocation parameters
  / Review generated JCL
  _ Warn if JCL already exists
  _ Warn if JCL was edited after generation

                                           Row 1 of 5

Cmd Job Card Information
-----
//GCLDEMO  JOB  , 'TSAXC', CLASS=A, MSGCLASS=X, NOTIFY=TSAXC,
//                REGION=256M
//JOBLIB   DD  DISP=SHR, DSN=QA1A.SDSNEXIT
//                DD  DISP=SHR, DSN=HLO.QA0110.SHLOLOAD
//                DD  DISP=SHR, DSN=DSN.VA10.SDSNLOAD
***** Bottom of data *****

```

- Build JCL into a PDS or dataset name

Profiles – Group Consistent Load JCL Example

```
//SYSINHLO DD *  
  IDAA CONSISTENT LOAD  
  (  
    TEMPLATE  
    (  
      NAME TEMPSAMP  
      DSN '&US..&DB..&SN..&SS..T&TI.'  
    )  
    GROUP  
    (  
      SPACE  
      (  
        CREATOR 'USEIDAA'  
        NAME 'GLWTDPT'  
      )  
      SPACE  
      (  
        CREATOR 'USEIDAA'  
        NAME 'GLWTEMP'  
        PARTITION 1  
      )  
      SPACE  
      (  
        CREATOR 'USEIDAA'  
        NAME 'GLWTEMP'  
        PARTITION 2  
      )  
      TO_CURRENT  
      NEW_COPY  
      FCCOPYDDN TEMPSAMP  
    )  
    ACCELNAME QDS5ACC1  
    PARALLEL '4,4'  
    BUFFERS_IN_31_BIT  
    LOG_COPY_PREFERENCE R1R2A1A2  
    USER_INDICATOR HLO  
  )
```

Profiles – Dual Load Options

```

LOADER                               Load from External Options                               2014/04/14  14:21:29
Command ===> _____

Commands: TABLE - select DB2 table  ACCELERATOR - select accelerator
          COLINFO - Edit table column info

Creator   . . . : TSAXC                Name . . . . . : DEMO EXTERNAL LOAD
Share option . . UPDATE                Description . . DEMO PROFILE EXTERNAL LOAD  >

Schema . . . . . : USEIDAA
Table name . . . : GLWTEMP
Partition . . . . : ALL

Target options:
Load target . . . . . B (A - Accelerator, B - Both accelerator and DB2)
Accelerator name . . . . QDS5ACC1

Required load options:
Input data set name . . . : TSAXC.UNLOAD.USEIDAA.GLWTEMP.A
Input member . . . . . : (if data set is partitioned)
Input DSN template . . . : &US..IDSD.&DB..&TS..P&PA. Update NO (Yes/No)
Table column info DSN . . : TSAXC.PUNCH.USEIDAA.GLWTEMP.A
Table column info member . . (if data set is partitioned)

DB2 load options:
Parallel load . . . YES (Yes/No)
Load tasks . . . . 4 (1-20)
Utility ID . . . . DUAL
KEEPDICTIONARY . . YES (Yes/No)
ENFORCE . . . . . YES (Yes/No)
LOG . . . . . NO (YES, NO, or NOCOPYPEND)
NUMRECS . . . . . 200000 (Integer or blank)
SORTDEVT . . . . . SYSALLDA (Device type or blank)
SORTNUM . . . . . 10 (2-255 or blank)

ERRDDN template DD name . . ISYSERR Update NO (Yes/No)
MAPDDN template DD name . . ISYSMAP Update NO (Yes/No)
SYSUT1 template DD name . . ISYSUT1 Update NO (Yes/No)
SORTOUT template DD name . . ISORTOUT Update NO (Yes/No)

```

Profiles – Dual Load Table Selection

```

LOADER                               DB2 Table selection                               2014/04/14  14:27:43
Command ===> _____                Scroll ===> PAGE

Commands:      DEFAULT - Default sort
Line commands: S - Select a table

Table creator like . . . TSAXC* _____ > DB2 SSID: QAA5
Table name like . . . RANDOM* _____ >
                                           Row 1 of 21
Cmd Table Name                               Part  Creator                               Database Tablespace
-----
RANDOM_TABLE                                ALL   TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                1     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                2     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                3     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                4     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                5     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                6     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                7     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                8     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                9     TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                10    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                11    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                12    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                13    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                14    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                15    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                16    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                17    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                18    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                19    TSAXC                                DBRAND  TSRAND
RANDOM_TABLE                                20    TSAXC                                DBRAND  TSRAND
***** Bottom of data *****

```

- Only one table (many parts) can be selected for Dual Load function

Profiles – Template DSN Builder

```

LOADER                               Input DSN Template                2014/04/14  14:35:08
Command ==>>> _____

```

```

Commands: SHOW - Show DSN

```

```

Template name . . ISYSREC

```

```

Enter codes to create a data set name specification:

```

```

Qualifier code . . __ Free form literal . . _____

```

```

Current data set name qualifier string:

```

```

&US..IDSD.&DB..&TS..P&PA.

```

```

Valid data set name specification codes are

```

1. Database	10. Time (HHMMSS)	19. Unique
2. Space name	11. Hours (HH)	20. SSID
3. Partition	12. Minutes (MM)	21. User ID
4. DSNUM	13. Seconds (SS)	22. Job name
5. Date (YYYYDDD)	14. Local/Recovery (L/R)	23. Step name
6. Year (YYYY)	15. Primary/Backup (P/B)	24. Utility ID
7. Month (MM)	16. Copy type (Full/Incr)	25. Utility name
8. Day (DD)	17. Listdef	26. Use freeform literal
9. Julian day (DDD)	18. Sequence	27. Substring qualifier

- ISPF interface helps builds dataset names for templates
- For FlashCopy and DB2 Load templates

Profiles – Dual Load JCL Example

```
//SYSIN DD *
TEMPLATE ISYSREC
  DSN &US..IDSD.&DB..&TS..P&PA.
  DISP(SHR,KEEP,KEEP)
LOAD DATA INDDN
IDAA_DUAL ON QDS5ACC1
ACCEL_LOAD_TASKS 4
REPLACE
KEEPDICTIONARY
LOG NO
SORTDEVT SYSALLDA SORTNUM 10
ERRDDN ISYSERR
MAPDDN ISYSMAP
WORKDDN (ISYSUT1,ISORTOUT)
INTO TABLE
  "TSAXC"."RANDOM_TABLE"
  PART 1
  INDDN ISYSREC
  (
    PART_ID          POSITION(1)    INTEGER EXTERNAL(8),
    PART_KEY         POSITION(9)    INTEGER EXTERNAL(8),
  )
INTO TABLE
  "TSAXC"."RANDOM_TABLE"
  PART 2
  INDDN ISYSREC
  (
    PART_ID          POSITION(1)    INTEGER EXTERNAL(8),
    PART_KEY         POSITION(9)    INTEGER EXTERNAL(8),
  )
/*
```

- DB2 Analytics Accelerator_DUAL directive tells Loader to load both DB2 and Accelerator
- DB2 Analytics Accelerator_ONLY directive tells Loader to load only Accelerator
 - NOTE! DB2 table data will be deleted on a Accelerator-only run

Technical Limitations



Technical Limitations of Accelerator Loader V1.1

- External Data Load Considerations
- DB2 LOAD utility considerations
 - Data is converted to internal format, essentially running the LOAD into DB2 as FORMAT INTERNAL
- The Following syntax is currently restricted:
 - ASCII
 - CCSID
 - CONTINUEIF
 - DECFLOAT_ROUNDMODE
 - DEFAULTIF (Support coming soon)
 - EBCDIC (This is default)
 - FORMAT (Only External currently supported)
 - IDENTITYOVERRIDE
 - IGNOREFIELDS
 - INCURSOR
 - NOSUBS
 - NULLIF (Support coming soon)
 - PRESORTED
 - RESUME (Accelerator Restriction)
 - ROWFORMAT
 - SHRLEVEL CHANGE
 - STRIP
 - TRUNCATE
 - UNICODE

Technical Limitations of Accelerator Loader V1.1 (cont)

- External Data Load Considerations
- Accelerator considerations
 - Data types not supported by the accelerator are not supported by Loader
 - Table must be defined in accelerator
 - AQT_MAX_UNLOAD_IN_PARALLEL is honored
 - You cannot use Loader to load individual partitions of a table that has the Accelerator status of 'InitialLoadPending'. A full table load must be performed first

Technical Limitations of Accelerator Loader V1.1 (cont)

- External Data Load Considerations
- Loader considerations
 - User is responsible for creating load file
 - Extracted data can come from various sources
 - IMS, VSAM, Oracle...etc
 - Field specifications are required in the LOAD utility syntax
 - Constraints are not validated when loading into the accelerator
 - DB2 LOAD will validate check constraints when loading into DB2
 - It is recommended a CHECK DATA with SCOPE ALL be run on the accelerator table if violations are found by DB2
 - Internally-formatted (Format Internal) data is not yet supported
 - Coded character set identifier (CCSID) conversion is not supported.
 - Only EBCDIC code pages are supported.
 - Trailing spaces in object names are not supported.
 - Ensure that any table space that you attempt to load was created with DEFINE YES, or, if created with DEFINE NO, that the underlying VSAM linear data sets have been created by an INSERT or a LOAD.
 - Generated columns are not supported

Technical Limitations of Accelerator Loader V1.1 (cont)

■ External Data Load Considerations

– Loader-supported Data Types

- BIGINT
- BINARY*
- CHAR
- DATE
- DECIMAL
- DOUBLE
- FLOAT
- INTEGER
- REAL
- SMALLINT
- TIME
- TIMESTAMP
- VARBINARY*
- VARCHAR

**When loading into the accelerator, Loader skips BINARY and VARBINARY data. When loading data into DB2, Loader converts BINARY and VARBINARY data to DB2 internal format and passes it to the DB2 LOAD utility.*

***Graphic column support is in progress*