Accelerate Your IMS Data!

Tim Willging, Rocket Software

Anthony Ciabattoni, Rocket Software

Session: 17750

Agenda

- DB2 Analytics Accelerator for z/OS overview
- Loading IMS data to the Accelerator
- Accelerator Loader – External ‘Dual’ Load
- General Accelerator Loader Details
- Summary

IBM zEnterprise and Analytics Accelerator

Transactional Processing

Analytics Workload

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

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

PureData System for Analytics

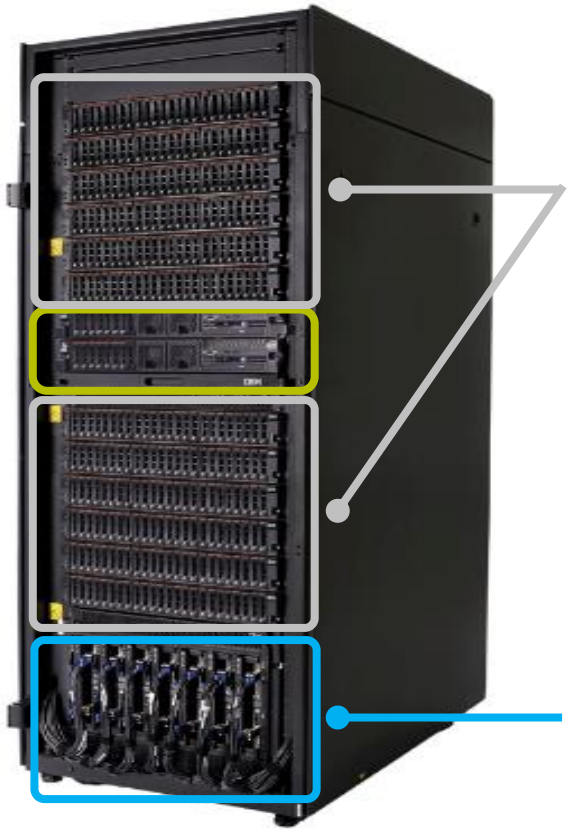

SMP Hosts
DB2 Analytics Accelerator Server
SQL Compiler, Query Plan, Optimize
Administration

Disk Enclosures
Slice of User Data
Swap and Mirror partitions
High speed data streaming
High compression rate

Snippet Blades™ (S-Blades, SPUs)
Processor &
streaming DB logic
High-performance database
engine streaming joins,
aggregations, sorts, etc.

DB2 for z/OS Approach: Hybrid Database Management System

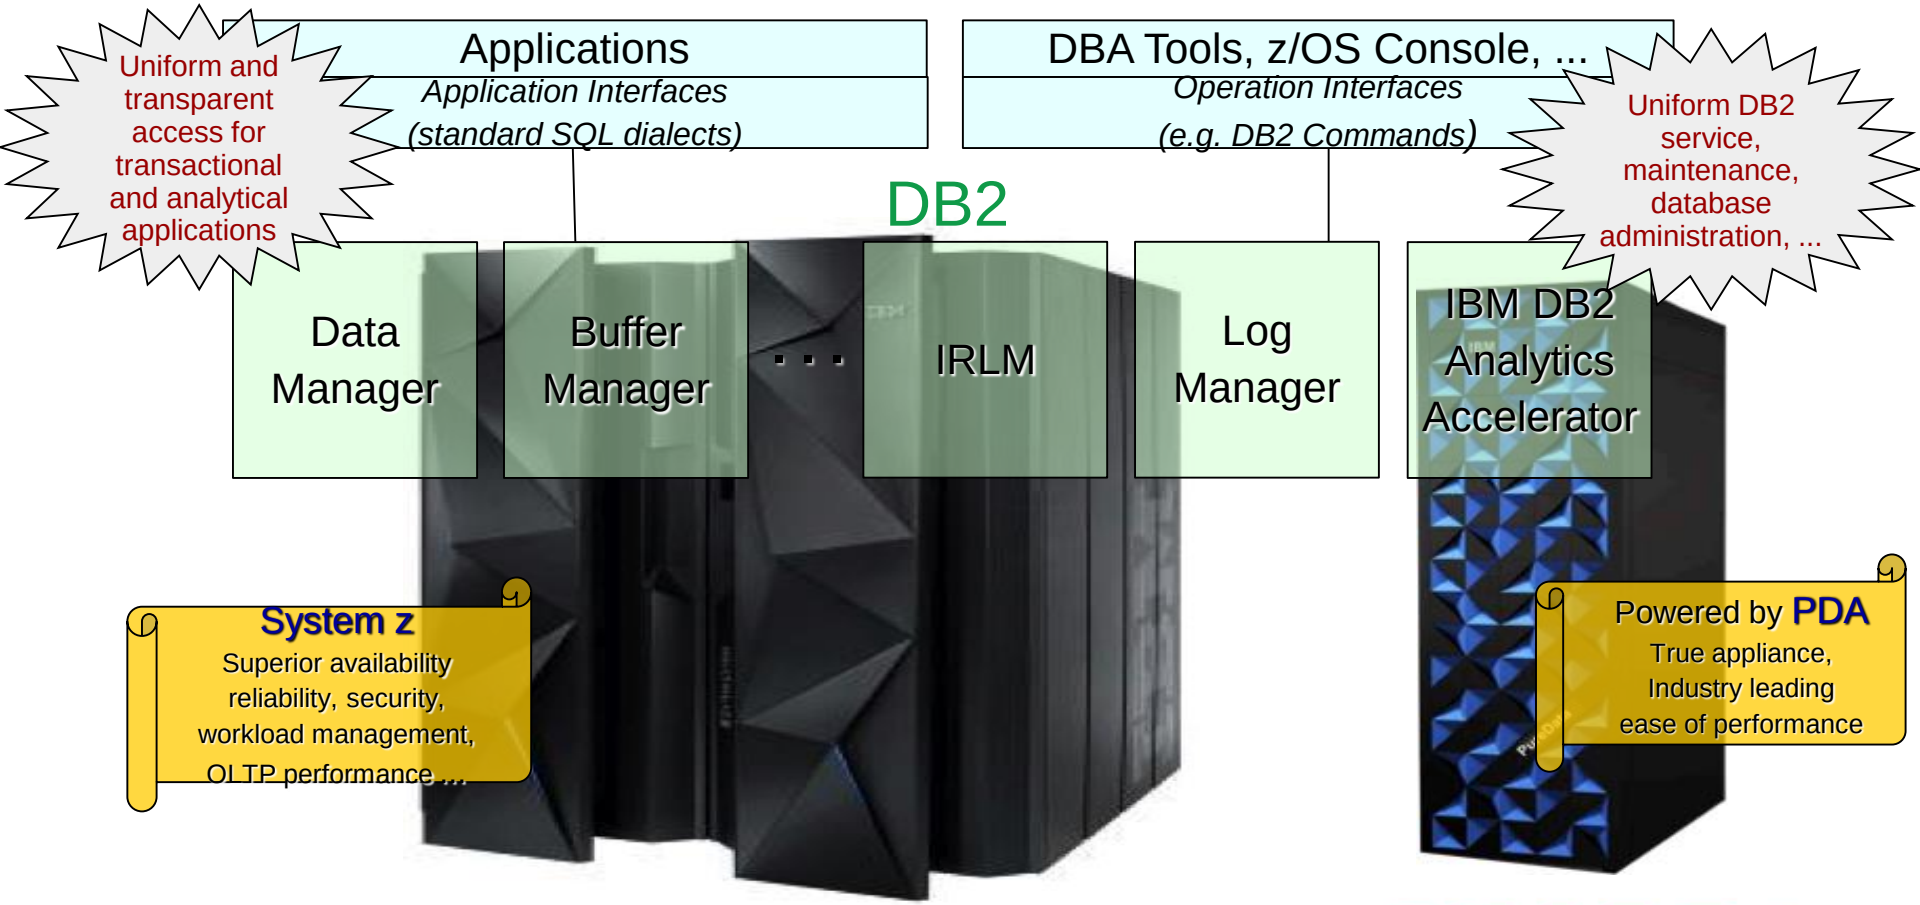

IBM DB2 Analytics Accelerator
Table Definition and Deployment

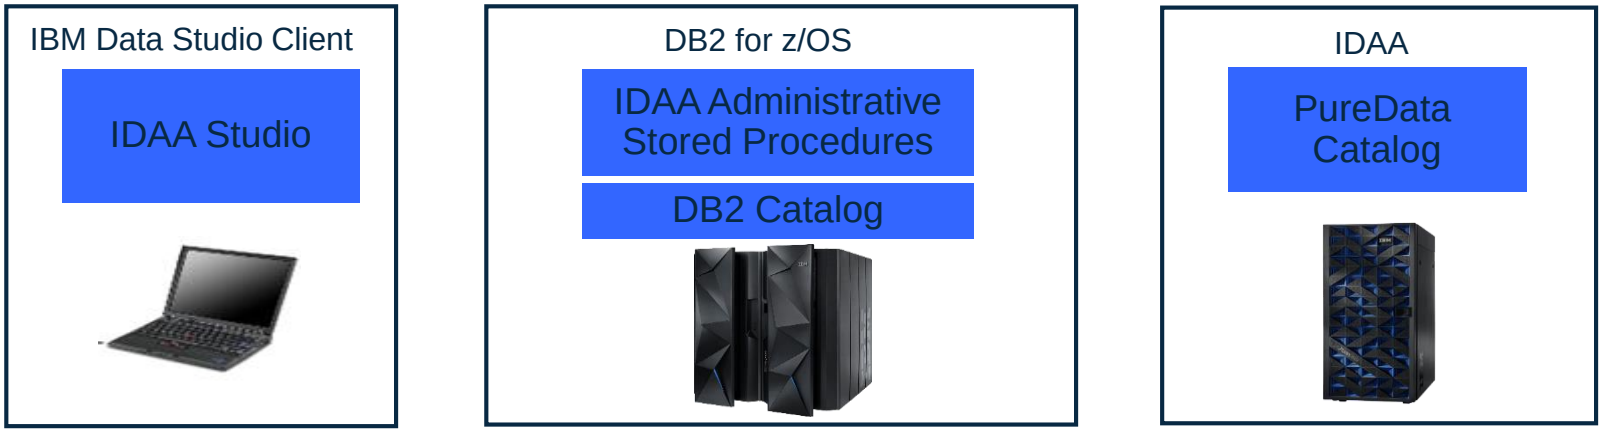

- The tables need to be defined and deployed to IDAA before data is loaded and queries sent to it for processing
  - Definition: identifying tables for which queries need to be accelerated
  - Deployment: making tables known to DB2, i.e. storing table metadata in the DB2 and PureData catalog
- IDAA Studio guides you through the process of defining and deploying tables, as wall 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

More users across the organization want access to business critical analytics

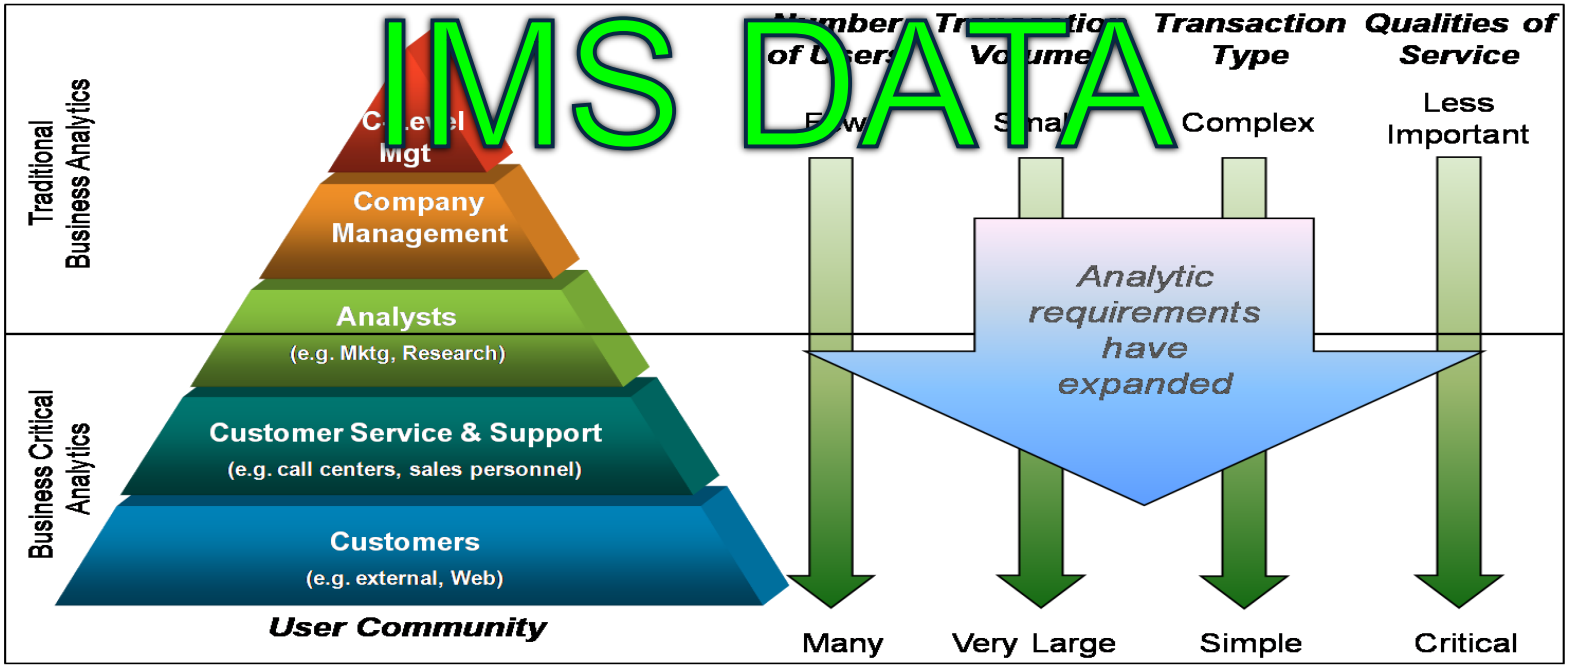

Query Execution Process Flow

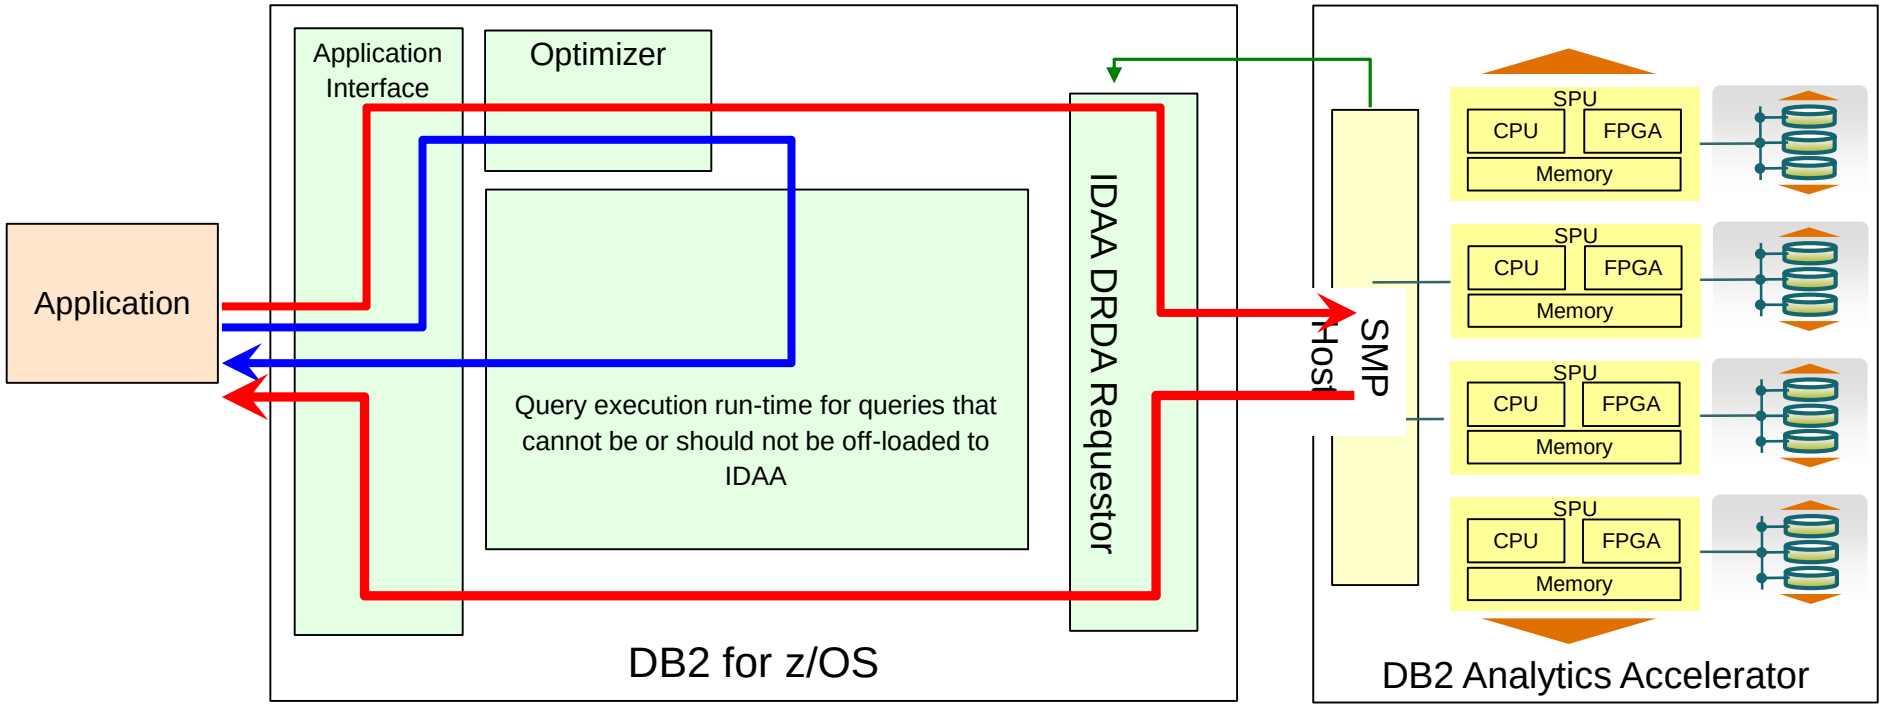

Heartbeat (DB2 Analytics Accelerator availability and performance indicators)

Queries executed without DB2 Analytics Accelerator

Queries executed with DB2 Analytics Accelerator

Queries executed with value of “ALL” may receive a SQL Error Code if the query cannot run on the accelerator

S-Blade Data Stream Processing

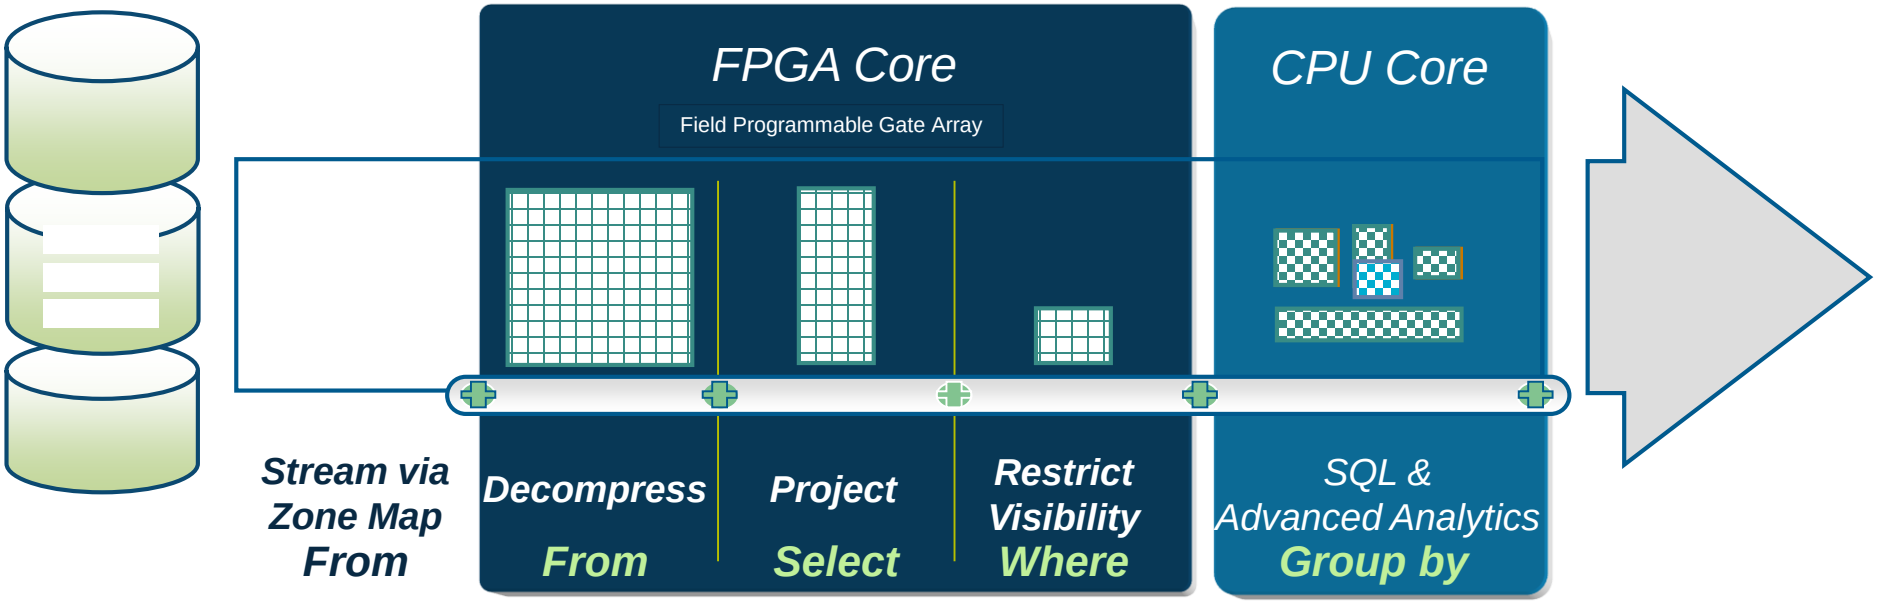


```
Select State, Age, Gender, count(*) From MultiBillionRowCustomerTable Where Date_Birth > ‘01/01/1960’
And State in (‘FL’, ‘GA’, ‘SC’, ‘NC’) Group by State, Age, Gender Order by State, Age, Gender
```

Query Routing Analysis

- Values for CURRENT QUERY ACCELERATION

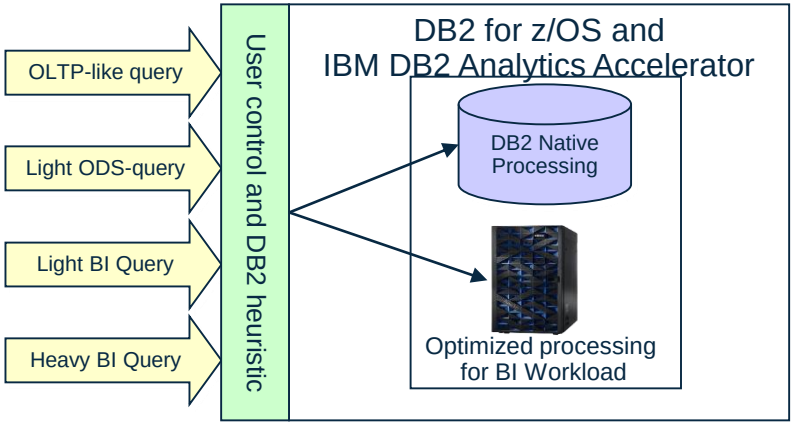

- Single and unique system for mixed query workloads
- Dynamic decision for most efficient execution platform
- New special register QUERY ACCELERATION
- New heuristic in DB2 optimizer
- Accelerator Only Tables available in IDAA V4.1 PTF5

Value	Description
NONE	No query is routed to the accelerator
ENABLE	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. 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
ENABLE WITH FAILBACK	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. 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
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

Accelerator Only Tables – New in 4.1 PTF5 !

- New table type in DB2
  - CREATE TABLE xx.xx … IN ACCELERATOR accel-name
- Data does not reside in DB2
- All queries routed to accelerator
  - Insert / Update / Delete / Select
  - No logging or point in time recovery
- Use Cases
  - In database transformations (IDT)
  - Calculated result set
  - Temporary table for statistic and analytic tools
  - Data must be inserted to load
  - INSERT INTO xx.AOT SELECT xx,xx,xx from xx.xx

History of IMS Analytics

- Desire to combine IMS data with other data
  - Social, DB2 z/OS data, VSMA, SAS data, other DBMS, etc.
  - Some DB2 on z/OS, but mostly off z/OS
  - Data being sent to potentially many sources
- Security can be compromised
- Each project typically ends up with a new copy of the data
- Performance historically not keeping up without $$$$$

Accelerate IMS Access - Proposed Solution

- Leverage Analytics Accelerator
  - Metadata resides in DB2
  - Copy IMS Data into Accelerator Only? (AOTs)
- DB2 manages queries and controls access

Advantages:

- Data never leaves z/OS
- IMS workload unaffected
- Single server for z Analytics
- Join of IMS/DB2 data
- Less reason to ETL DB2/IMS data off platform

Basic Process

- Decide IMS data needed
- Decide extraction and mapping tools and process
- Currency required (Refresh Frequency)
- Map IMS data to relational model
- Create DB2 table that matches extracted record format
- Add table to Accelerator
- Extract IMS data
- Load extracted data to DB2 table
- Load data from DB2 into the accelerator
  - Possibly insert to Accelerator Only Table

Extraction Considerations and Methods

- Considerations
  - Availability requirements
  - Frequency of refresh?
  - Impact to OLTP workload
  - What data is needed?
    - Entire database, certain segments, multiple DBs?
  - Consistency of data?

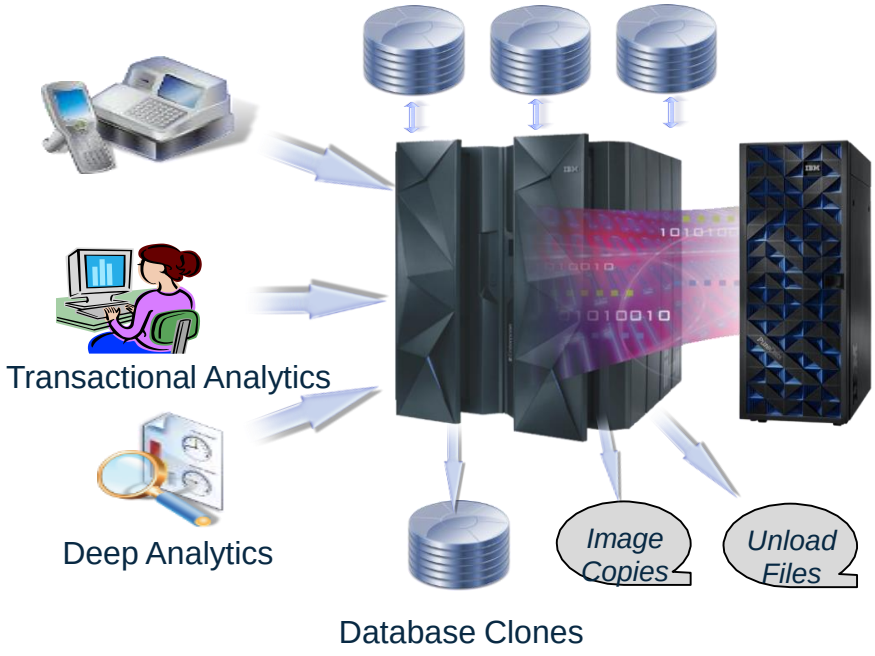

Extraction Considerations and Methods

- Extraction Tools and Methods
  - Custom IMS Application
    - Additional online workload
    - Data can still be changing
  - Database Clone (IMS Cloning Tool)
    - Group of databases at a point in time
  - Image Copies/Unload Files
    - Additional knowledge of structure needed
  - Mapping and ETL Tools
    - IMS Explorer
    - Data Stage, Informatica
    - Data Virtualization
    - IMS Catalog via JDBC
    - Other tools

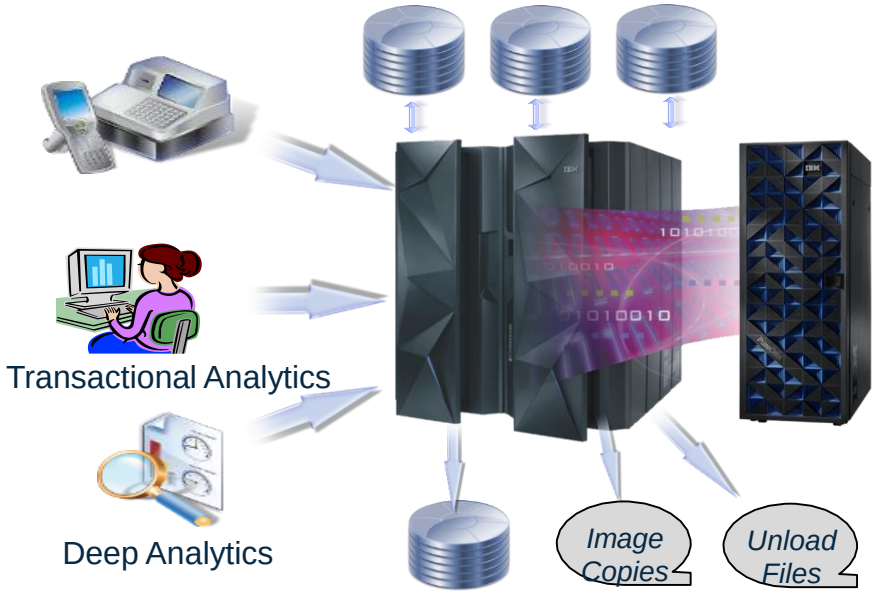

Mapping and Transforming Data

- Segment -> Table
  - Field -> Column
- Data type not required by IMS
  - Many times FIELD only defined for sequence fields
  - Data content not enforced by IMS
    - Data cleansing required?
- Where are field descriptions defined?
  - IMS Catalog
  - Copy books
  - JAVA Classes
- Non-unique or non-keyed segments

IMS Database

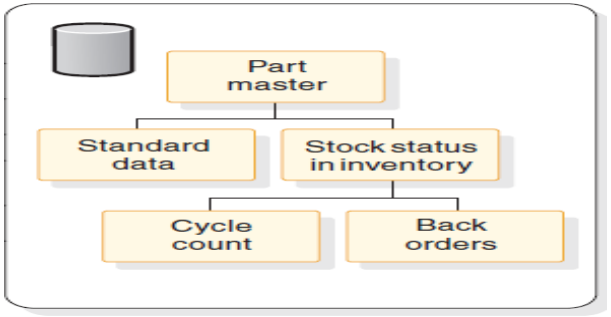

DB2 Tables

&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

Mapping IMS Data to Tables

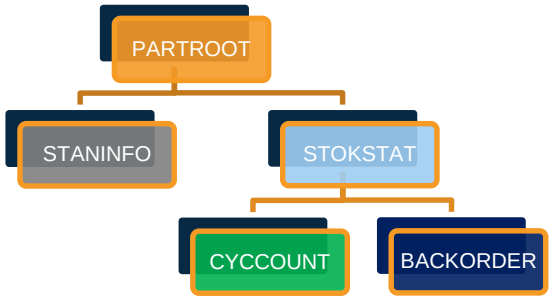

DBD:DI21PART

Table-name PARTROOT Column-names PARTKEY …	Table-name STANINFO Column-names STANKEY …	Table-name STOKSTAT Column-names STOCKKEY …	Table-name CYCCOUNT Column-names CYCLEKEY …	Table-name BAKCORDER Column-names BACKKEY …

IMS/DB definitions

DataBase Definition (DBD)

```
DBD NAME=DI21PART,ACCESS=(HISAM,VSAM)
DATASET DD1=DI21PART,DEVICE=3380,OVFLW=DI21PARO,
SIZE=(2048,2048),RECORD=(678,678)
SEGM NAME=PARTROOT, PARENT=0,BYTES=50, FREQ=250
FIELD NAME=(PARTKEY,SEQ),TYPE=C,BYTES=17,START=1
SEGM NAME=STANINFO,PARENT=PARTROOT,BYTES=85, FREQ=1
FIELD NAME=(STANKEY,SEQ),TYPE=C,BYTES=2,START=1
SEGM NAME=STOKSTAT, PARENT=PARTROOT, BYTES=160, FREQ=2
FIELD NAME=(STOCKEY,SEQ),TYPE=C,BYTES=16,START=1
SEGM NAME=CYCCOUNT, PARENT=STOKSTAT, BYTES=25, FREQ=1
FIELD NAME=(CYCLKEY,SEQ),TYPE=C,BYTES=2,START=1
SEGM NAME=BACKORDR, PARENT=STOKSTAT, BYTES=75, FREQ=0
FIELD NAME=(BACKKEY,SEQ),TYPE=C,BYTES=10,START=1
DBDGEN
FINISH
END
```

Program Specification Block (PSB)

```
DBPCB01  PCB    TYPE=DB,DBDNAME=DI21PART,PROCOPT=GOT,
KEYLEN=43
SENSEG NAME=PARTROOT
SENSEG NAME=STANINFO,PARENT=PARTROOT
SENSEG NAME=STOKSTAT,PARENT=PARTROOT
SENSEG NAME=CYCCOUNT,PARENT=STOKSTAT
SENSEG NAME=BACKORDR,PARENT=STOKSTAT
PSBGEN LANG=COBOL,PSBNAME=DFSSAM07
END
```

Flattening IMS Database Records

- Concatenated Keys
  - Concatenated key fields not stored with segment data
  - Key fields needed for each row to maintain referential integrity
- OCCURS clauses
  - Multiple instances of a field in a single instance of a segment
  - Multiple ‘rows’ should be generated

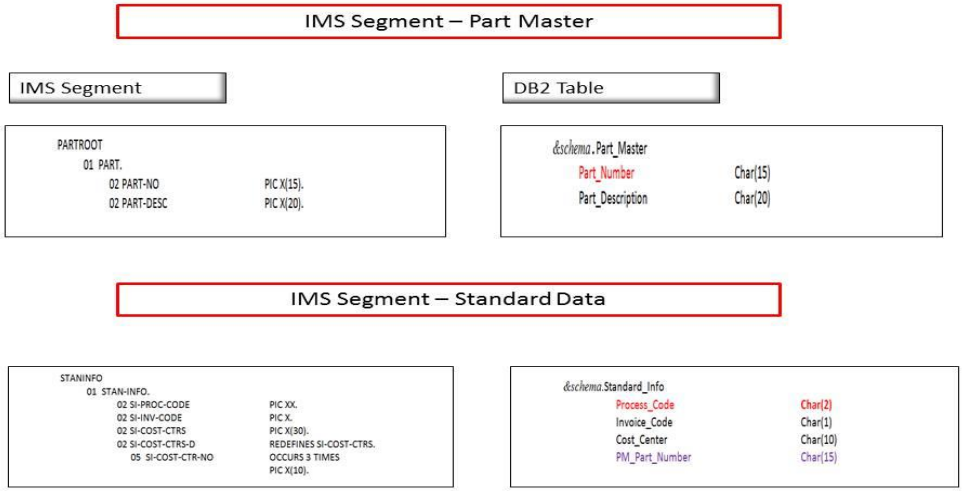

STANINFO (Standard data)
01 STAN-INFO.
02 SI-PROC-CODE PIC XX. <-Key
02 SI-INV-CODE PIC X.
02 SI-COST-CTRS PIC X(30).
02 SI-COST-CTRS-D REDEFINES SI-COST-CTRS.
05 SI-COST-CTR-NO OCCURS 3 TIMES
PIC X(10).

Loading Transformed Data

- End result of transformation: Data in DB2 Load file format
- DB2 Load Utility can perform more transformations
- Load syntax needed to describe IMS data in file

```
(DEPTNO      POSITION (1:3)     CHAR(3),
DEPTNAME     POSITION (4:39)    CHAR(36),
MGRNO        POSITION (40:45)   INTEGER EXTERNAL(6),
HDATE        POSITION (46:55)   DATE EXTERNAL(10),
Etc…
```

- Custom process to insert data to AOTs?
  - Slower performance

IMS Data in DB2 Analytics Accelerator

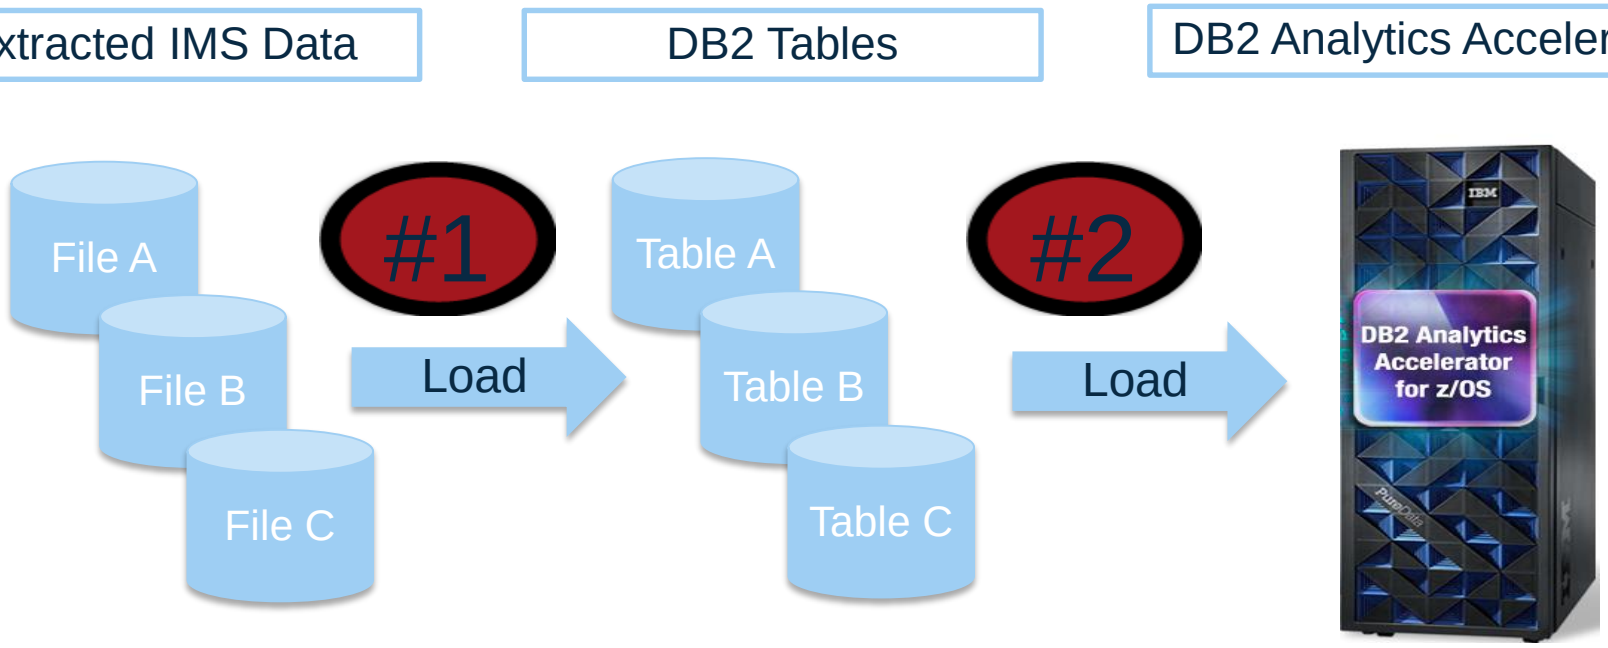

Two Step Load Process – Can be CPU Resource Intensive
Possibly insert directly or copy to Accelerator Only Table

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?

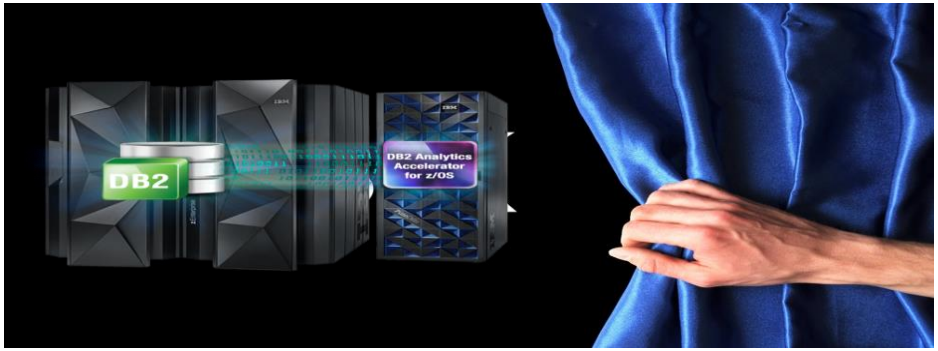

IBM Tools: Maximizing your Analytics Accelerator Investment

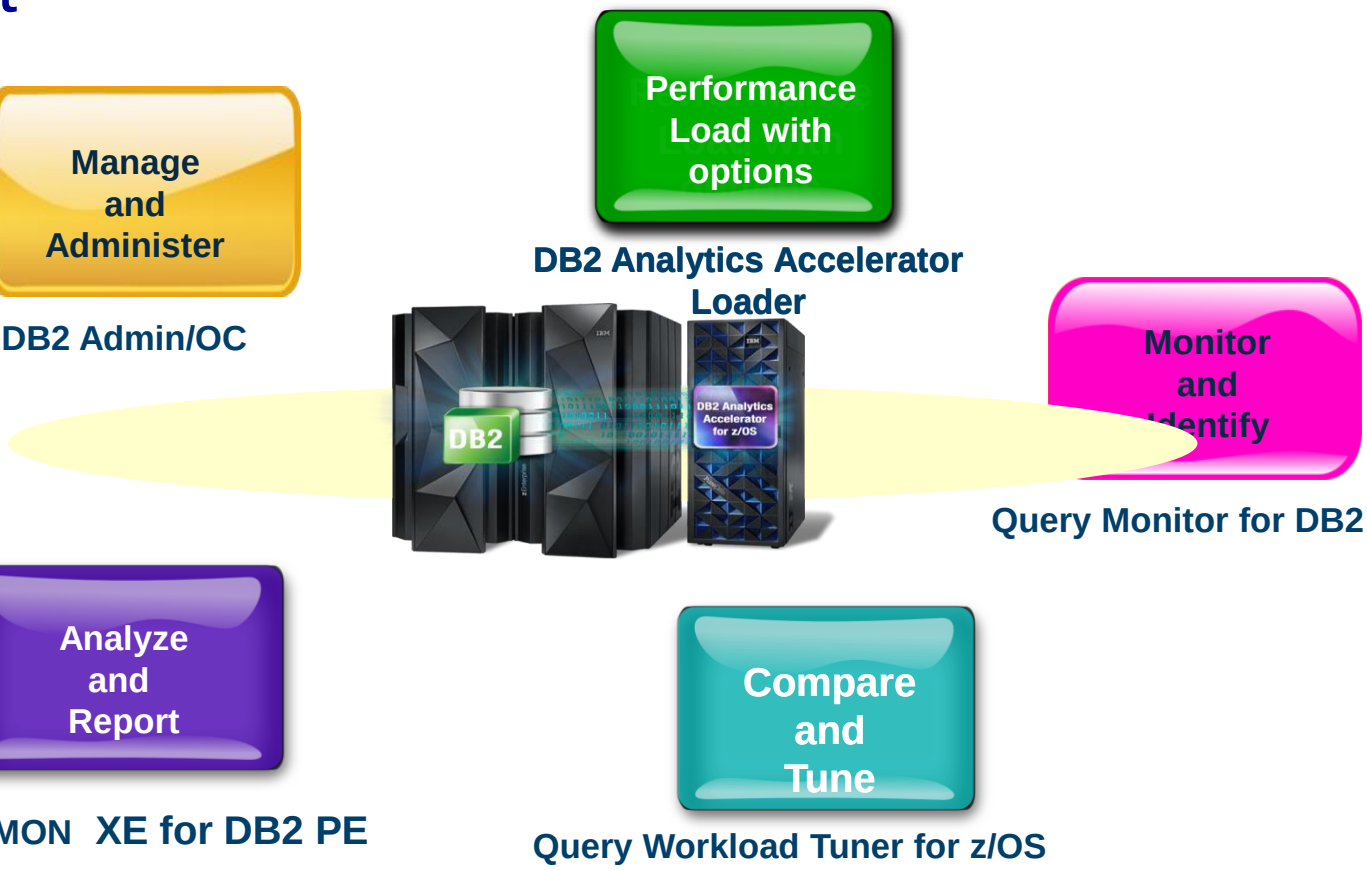

IBM DB2 Analytics Accelerator Loader:
What is External (Dual) Load

- 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 – Supports Accelerator Only Tables
     - 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.

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

DB2 Tables	Extracted IMS Data	DB2 Analytics Accelerator

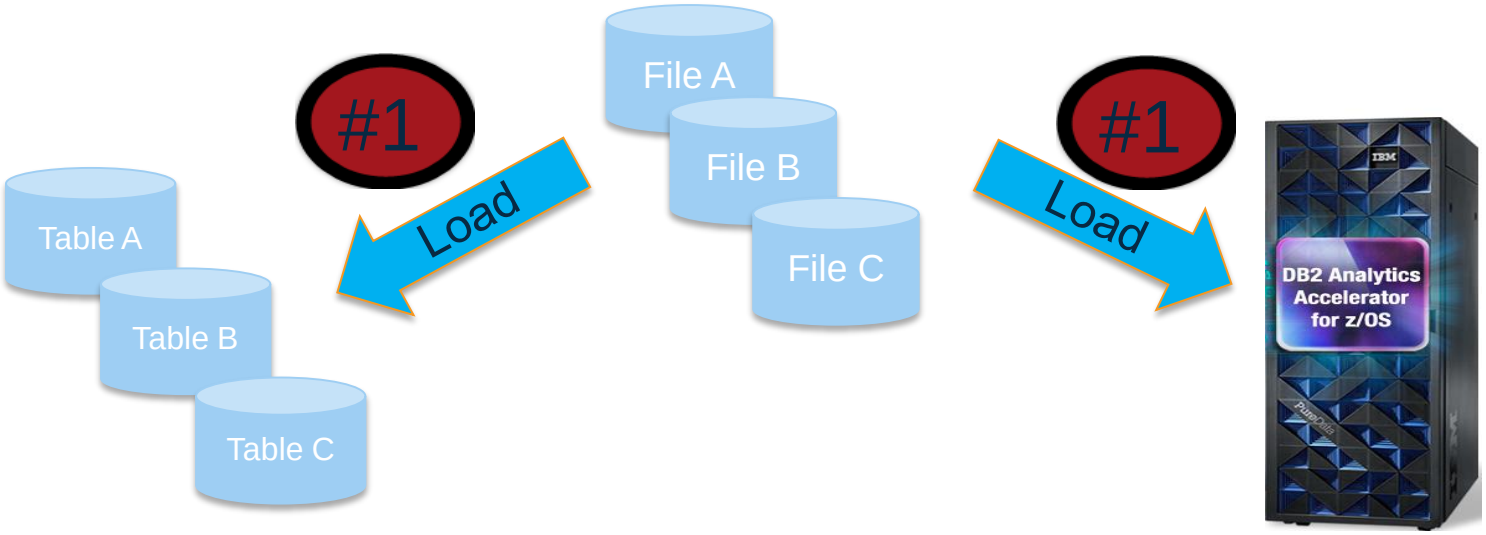

Parallel Load into DB2 and Accelerator – Faster Load Cycles – Reduce Costs
Use if data required in DB2

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

DB2 Tables	IMS Extracted External Data	DB2 Analytics Accelerator

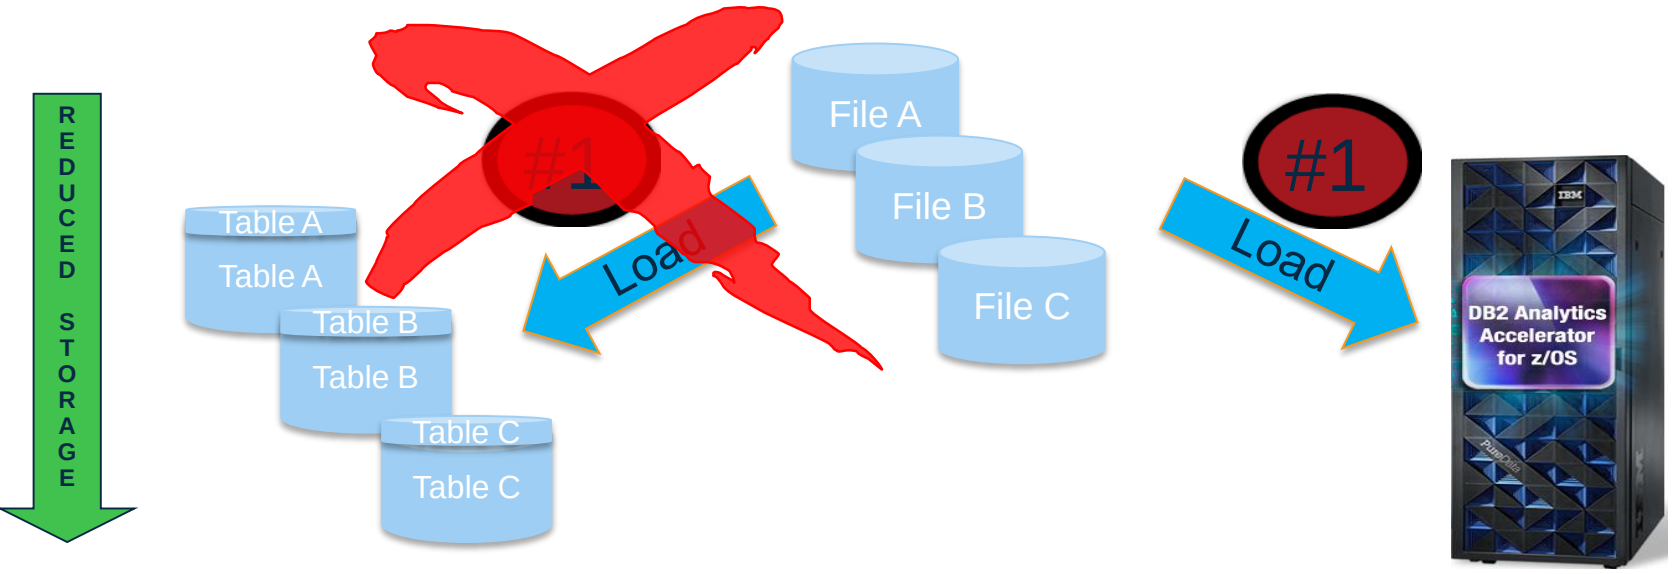

DB2 Analytics Accelerator Only Load – Reduced Elapsed Time – Reduced Cost – Reduced DASD
Load Directly to Accelerator Only Table!!!

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 or Accelerator Only Table
  - Up to 60% general purpose CPU reduction
  - DB2 Storage savings

Performance

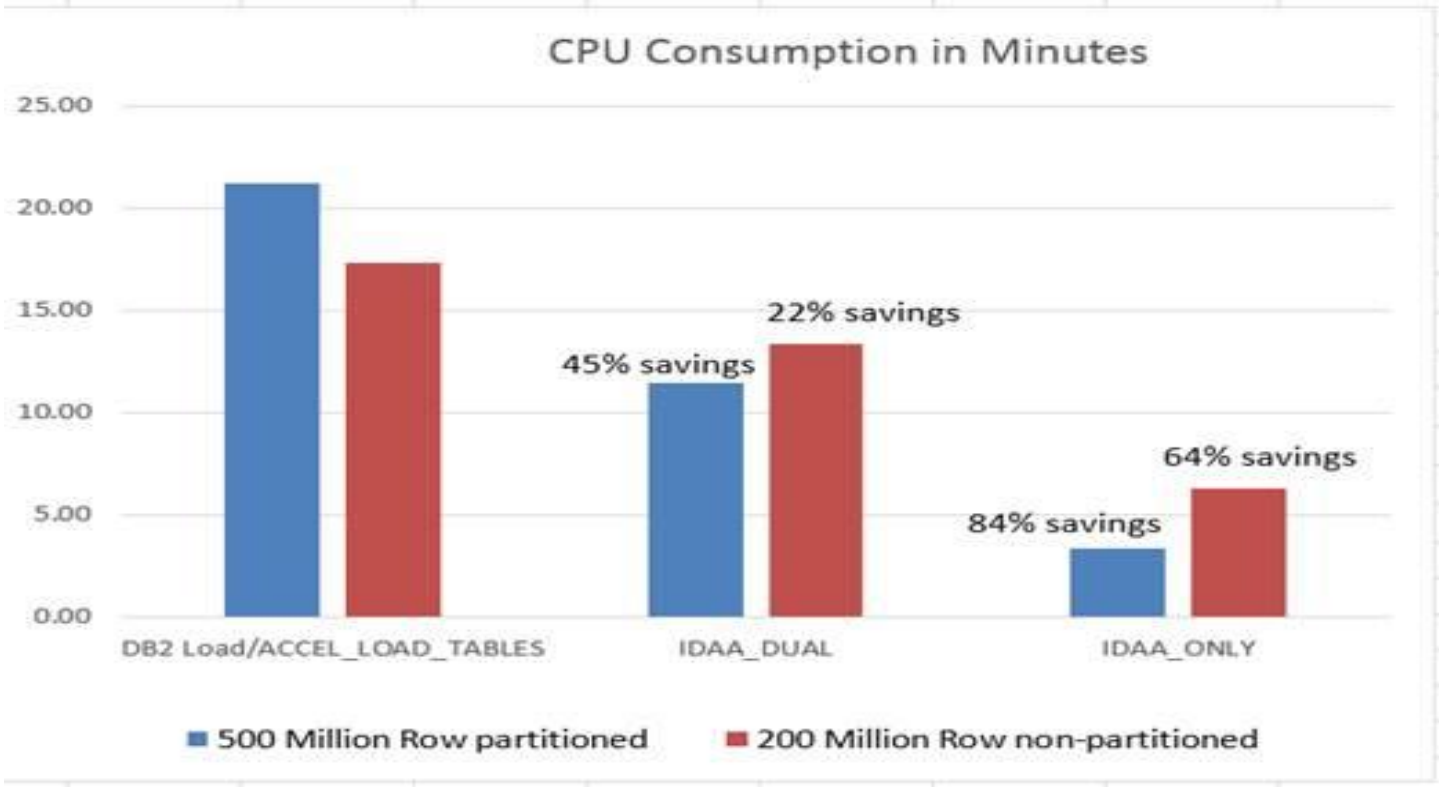

Performance

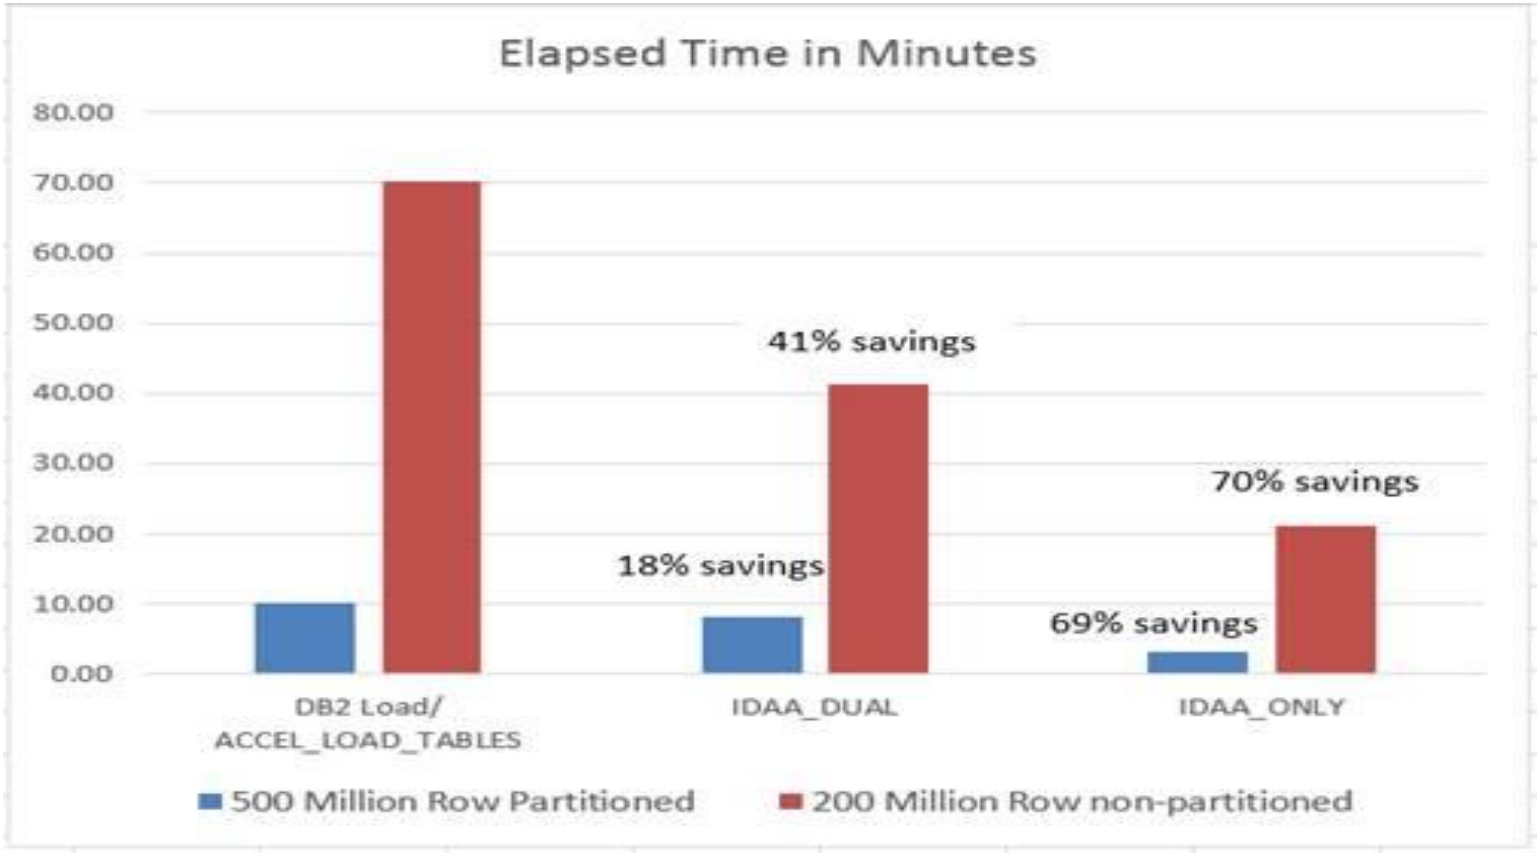

External Load ‘Accelerator-Only’ Considerations

- When should user consider loading accelerator only?
  - Data is maintained and updated elsewhere
  - DB2 is not required for data backup & recovery
  - All queries are qualified for accelerations
- Table must still exist in DB2 Catalog
  - If DB2/IDAA Table pair, DB2 table will be emptied upon load
    - Users should define small tablespaces (saves storage)
  - If AOT, DB2 table will never contain data
- Access to accelerated table remains via DB2
  - All DB2 security is honored

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

Current Process to Load Accelerator

Production DB2 Tables

DB2 Analytics Accelerator

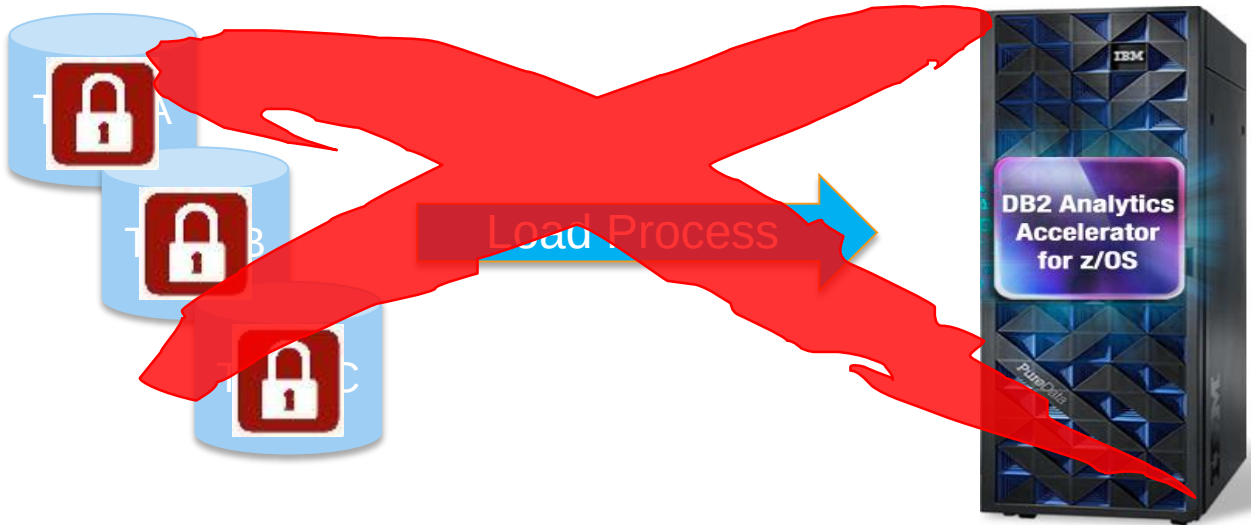

Loading Historical Data

May Be Locked During Load – Always Point in Time Loads

IBM DB2 Analytics Accelerator Loader
Group Consistent Load

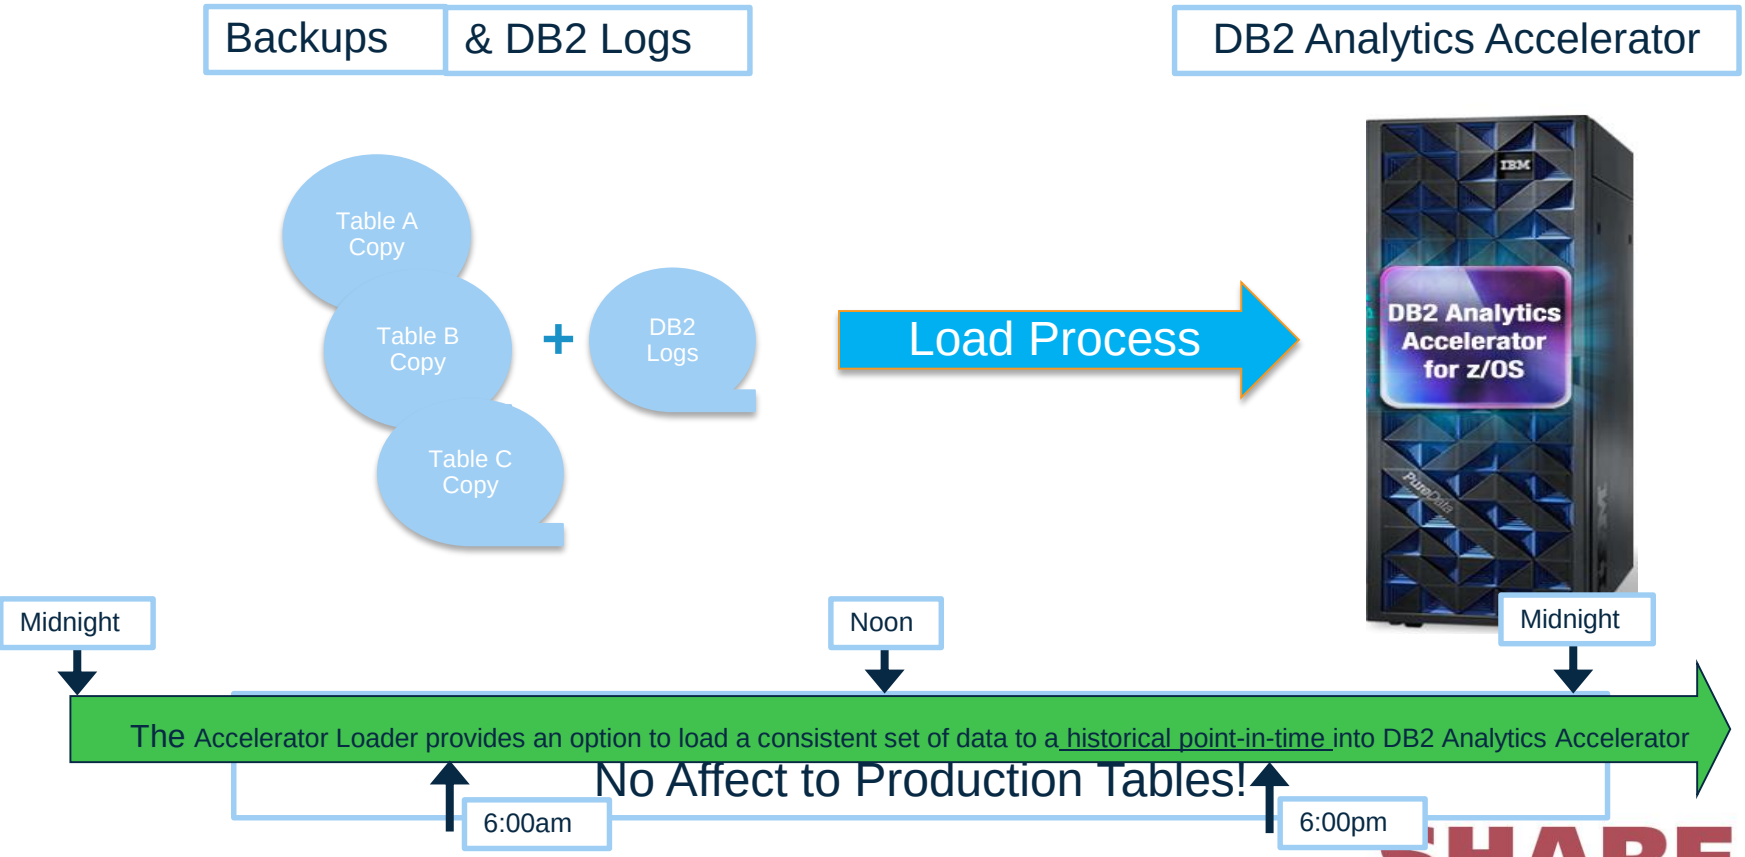

DB2 Analytics Accelerator Loader: Image Copy Load

- What is Image Copy Load?
  - Supplies the functionality to use an external or alternative image copy to load an Accelerator table
    - Source can be an OLTP image copy
    - Source image copy could be an alternative table on the same DB2 subsystem
    - Image copy data represents the point in time of the copy
- Image Copy Load use cases
  - Consolidate data from different DB2 systems into one accelerator
  - Historical load to alternate table in same system
    - Week-end / Month-end analysis
  - Use source ODS/Data Warehouse image copies to perform redirected loads/shadow table loads
    - Problem diagnosis
    - Perform analysis/data comparisons without impacting existing queries
  - Reuse of existing image copies

Current Process to Load Accelerator

Production OLTP Table

ODS/DW DB2 Table

DB2 Analytics Accelerator

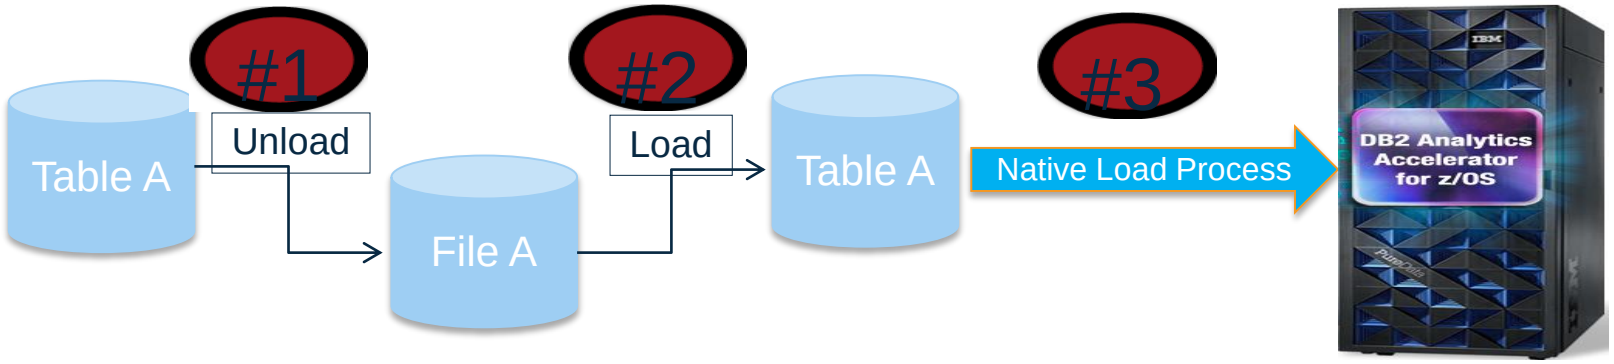

External DB2 Subsystem Table Load

DB2 Accelerator Loader – Image Copy Load

Production OLTP IC

ODS/DW DB2 Table

DB2 Analytics Accelerator

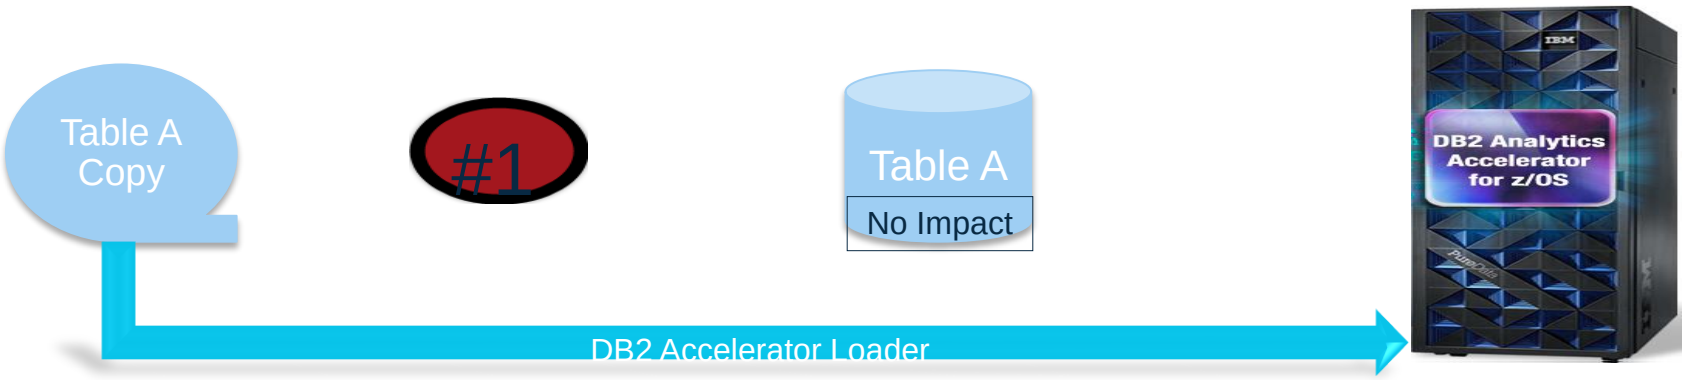

External DB2 Subsystem Image Copy Load
Load directly to Accelerator Only Table (Future APAR)

External Load vs Group Consistent Load vs IC 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 historical data needs to be loaded

- Image Copy Load
  - When an image copy is from an external DB2
  - A local image copy can be used to load an alternative DB2 table

Query Level Reporting of Accelerated Queries

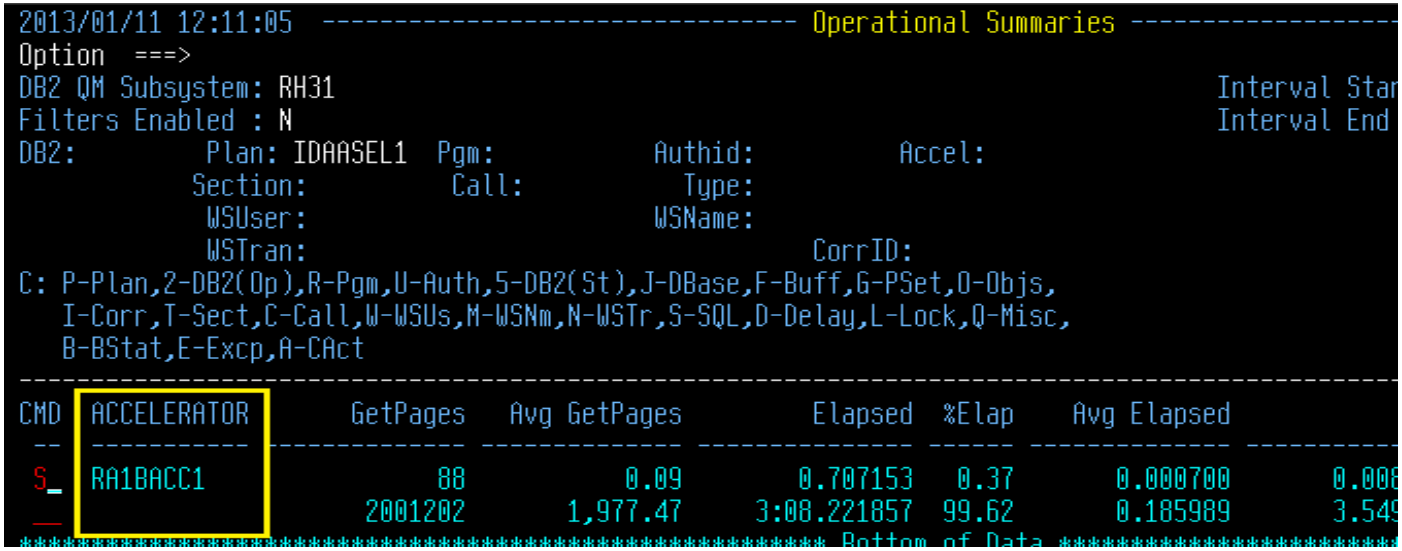

- Real time reporting of accelerated queries
- Summary of what ran in each accelerator (could be many)
- Summary of non-accelerated queries

Accelerator Modeling

- DB2 optimizer reports on query eligibility
  - Enabled at ZPARM level
- Query Monitor Accelerator Modeling allows customers to:
  - Report workload that ‘would-be’ accelerated
  - Report potential CPU/Elapsed savings if accelerator was attached
  - Report real DB2 CPU/Elapsed for queries that would be accelerated
  - Potential value in connecting accelerator to new DB2 system
- Easy to implement for existing Query Monitor customers
- Drill down to query level on what ‘would-be’ accelerated

IBM Query Monitor for DB2 on z/OS Accelerator Modeling

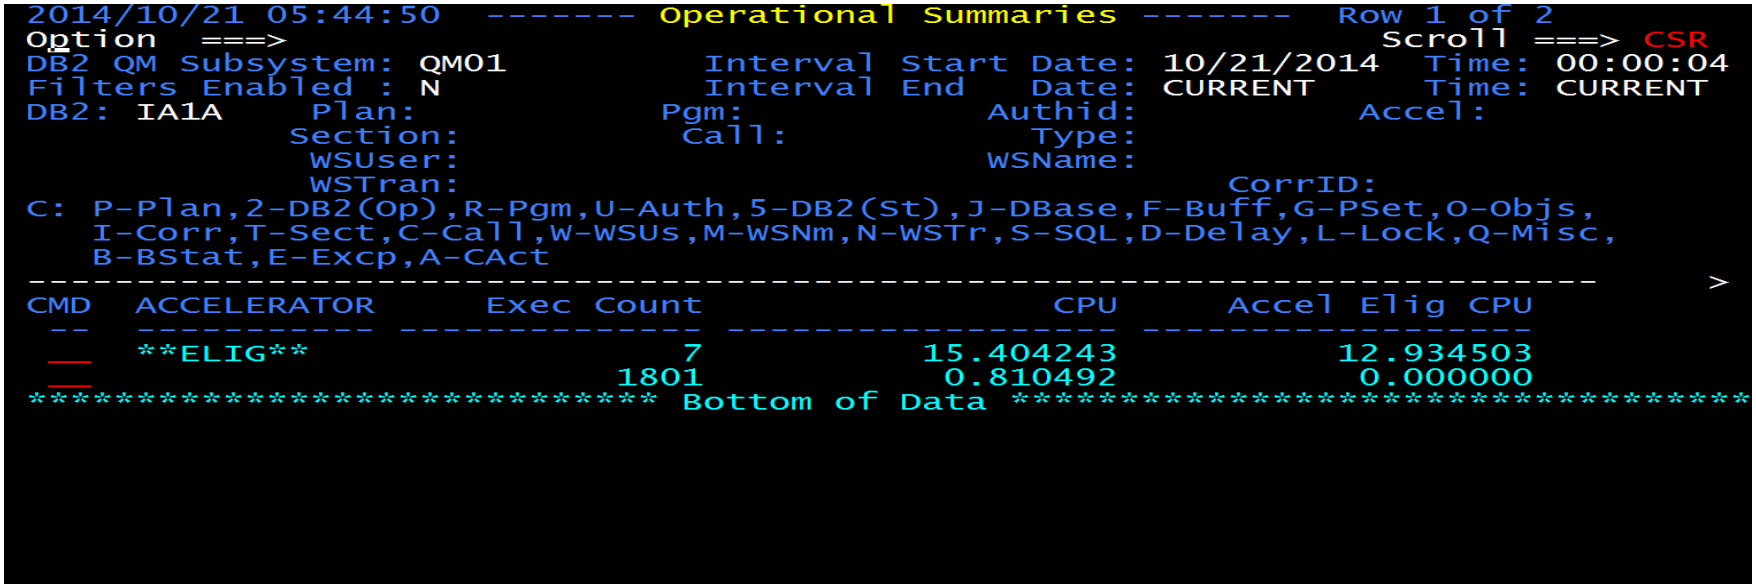

- Estimated CPU that would be saved if accelerator present

Basic Process Summary

- Decide IMS data needed
- Decide extraction and mapping tools and process
- Currency required (Refresh Frequency)
- Map IMS data to relational model
- Create DB2 table that matches extracted record format
- Add table to Accelerator
- Extract IMS data
- Load extracted data to DB2 table
- Load data from DB2 into the accelerator
- If using Accelerator Loader
  - Load directly to accelerator only table

Questions

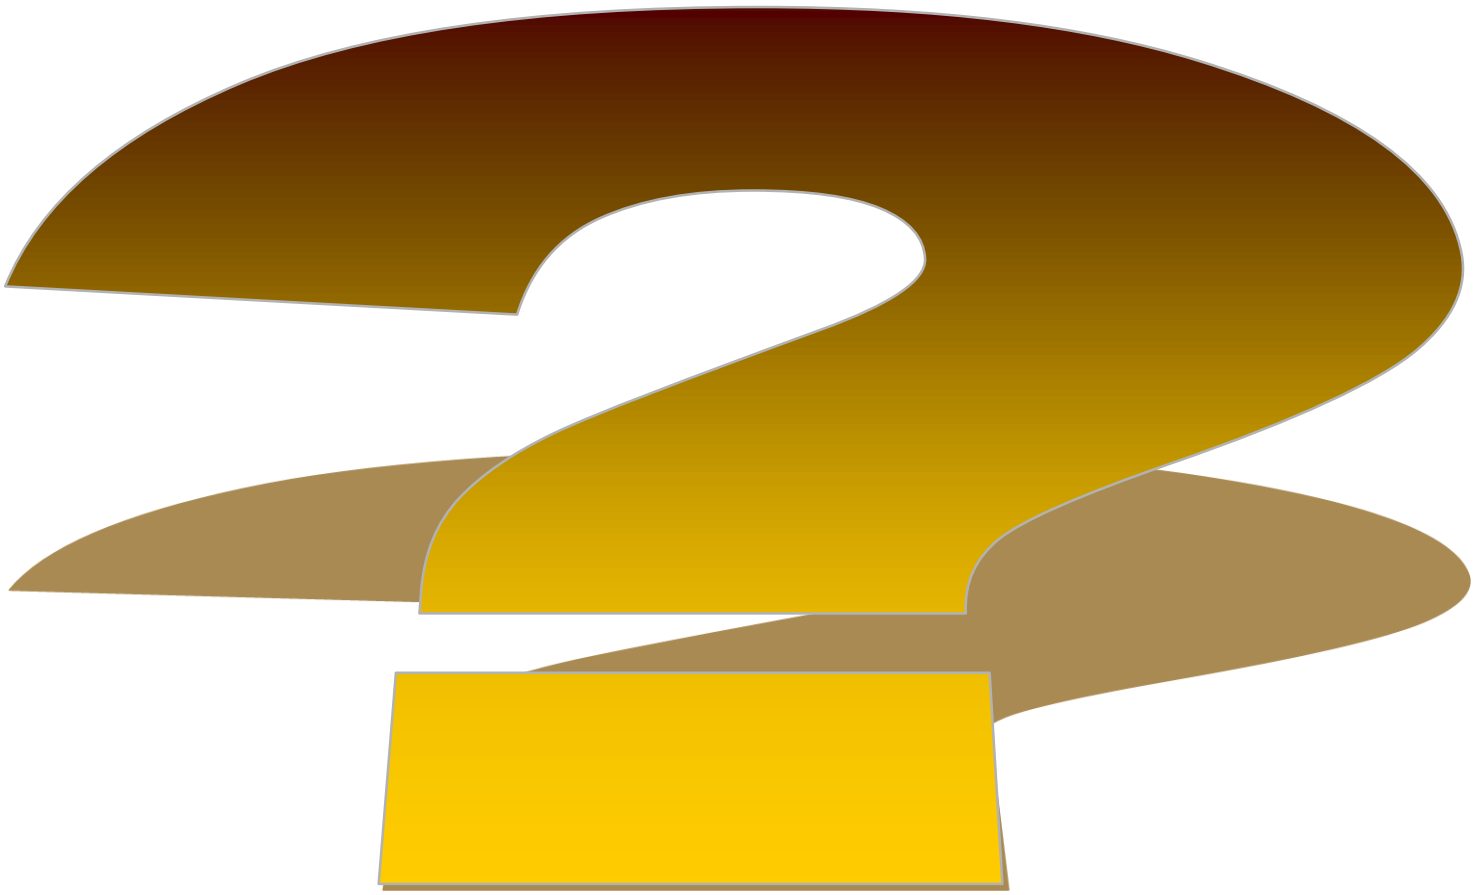

SHARE
Educate · Network · Influence

Complete your session evaluations online at www.SHARE.org/Orlando-Eval

SHARE in Orlando 2015