


Mike Hood Technical Advisor

April, 2016

# DB2 Analytics Accelerator V5.1

## Technical Overview



# Agenda

- Value Proposition
- Technical Overview
- New Version 5.1 Functionality
- Managing and Monitoring
- Use Cases and Customer Experiences
- Strategy and Evolution

# Analytics have become Business Critical

## Prevent Fraud



Today, analytics are **integrated with transaction systems** running on the mainframe and are critical to the business

## Reduce Customer Churn



Failure of these applications for any length of time can result in lost business, customer turnover, reputational risk, etc.

## Cross-sell/up-sell to customers



These applications need to deliver insight in **real-time or near real-time** and integrate with business processes

## Operational Reporting



These applications may **support a large concurrent user population** with a high volume of requests

Business Critical Analytic applications require superior qualities of service, including a high degree of reliability, continuous availability, scalability, security and low data latency

# More users across the organization depend upon business critical analytics

Traditional Business Analytics

Business Critical Analytics



<i>Number of Users</i>	<i>Transaction Volume</i>	<i>Transaction Type</i>	<i>Qualities of Service</i>
Few	Small	Complex	Less Important
<b>Analytic requirements have expanded</b>			
Many	Very Large	Simple	Critical

**User Community**

# Hybrid transaction/Analytical processing



## The hybrid computing platform on z Systems

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

Delivers industry leading performance for mixed workloads

The unique heterogeneous scale-out platform in the industry

Superior availability, reliability and security

## DB2 Analytics Accelerator and DB2 for z/OS

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



# IBM DB2 Analytics Accelerator for z/OS

A workload optimized, appliance add-on to DB2 for z/OS that enables the integration of analytic insights into operational processes to drive business critical analytics & exceptional business value.

## SPEED

- Dramatically improve query response – up to 2000X faster – to support time-sensitive decisions
- Right-time. Low latency. Trusted. Accurate.

## SAVINGS

- Minimize data proliferation
- Lower the cost of storing and managing historical data
- Free up compute resources

## SIMPLICITY

- Simplify infrastructure, reduce ETL and data movement off-platform
- Non-disruptive installation

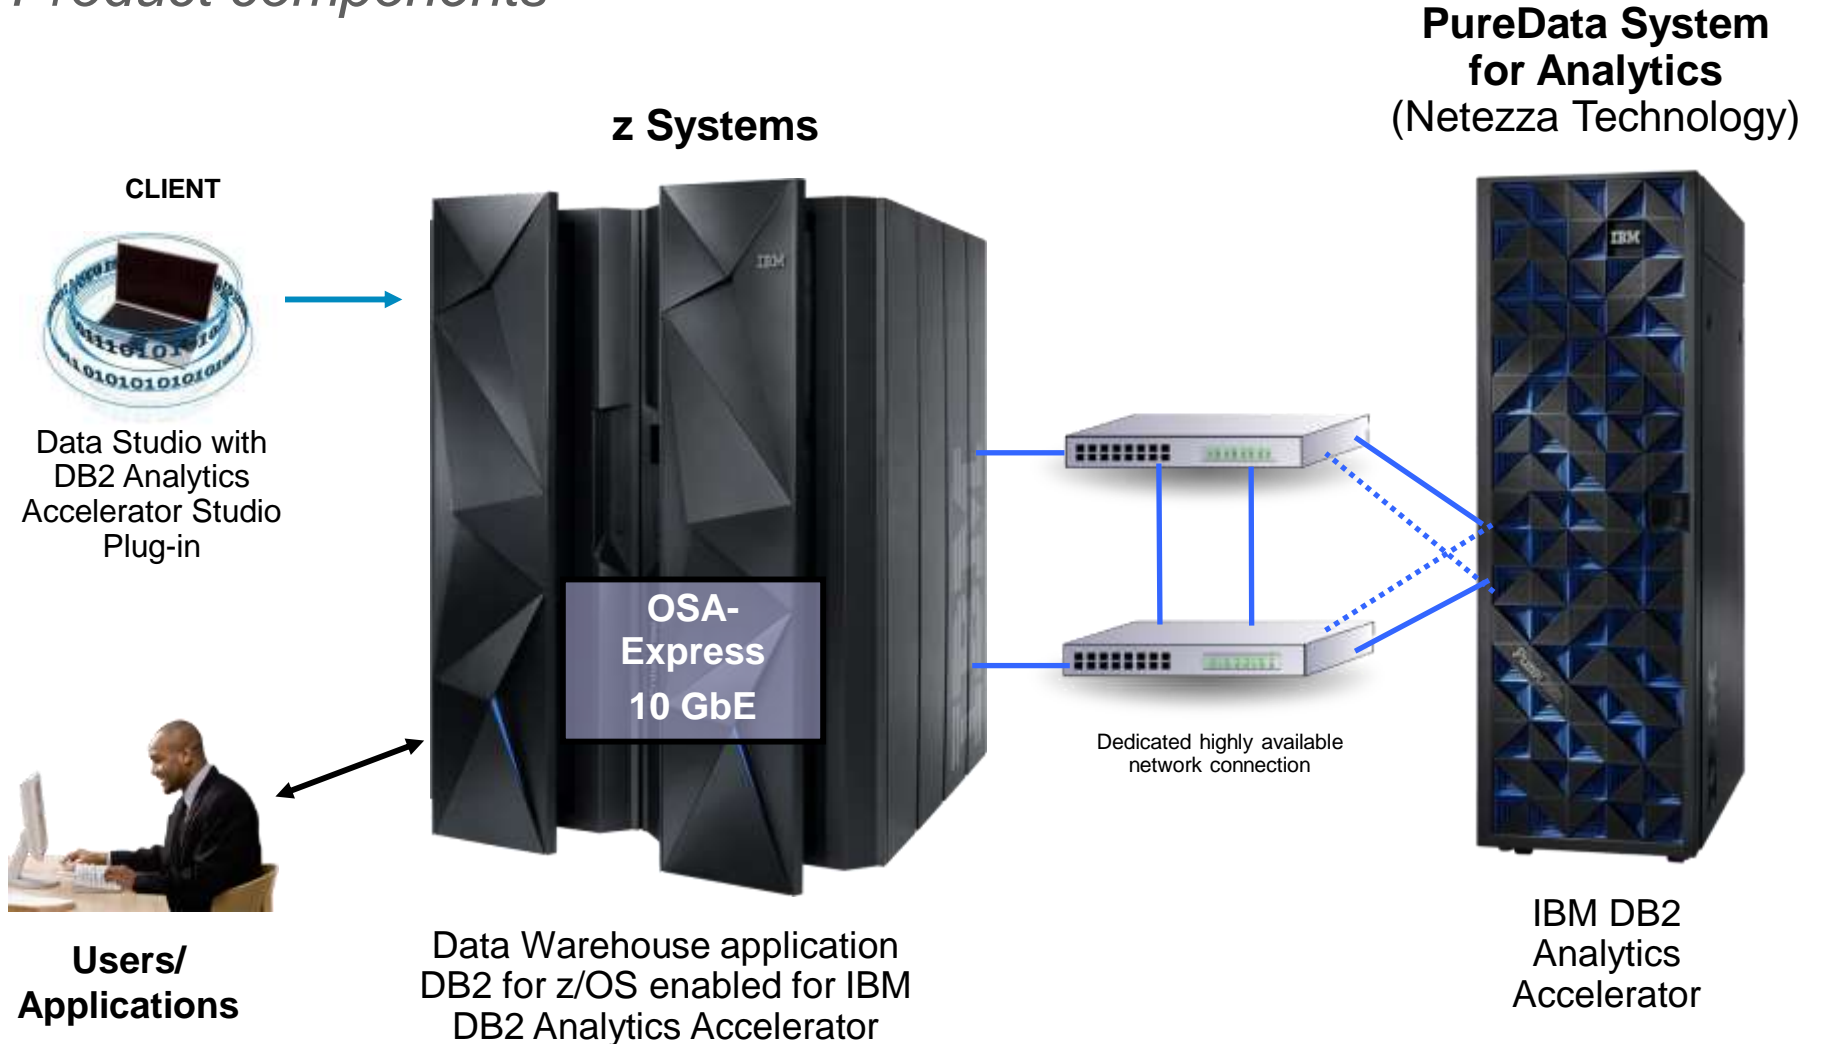
## SECURITY

- Safeguard valuable data under the control and security of DB2 for z/OS
- Protected. Secured. Governed.

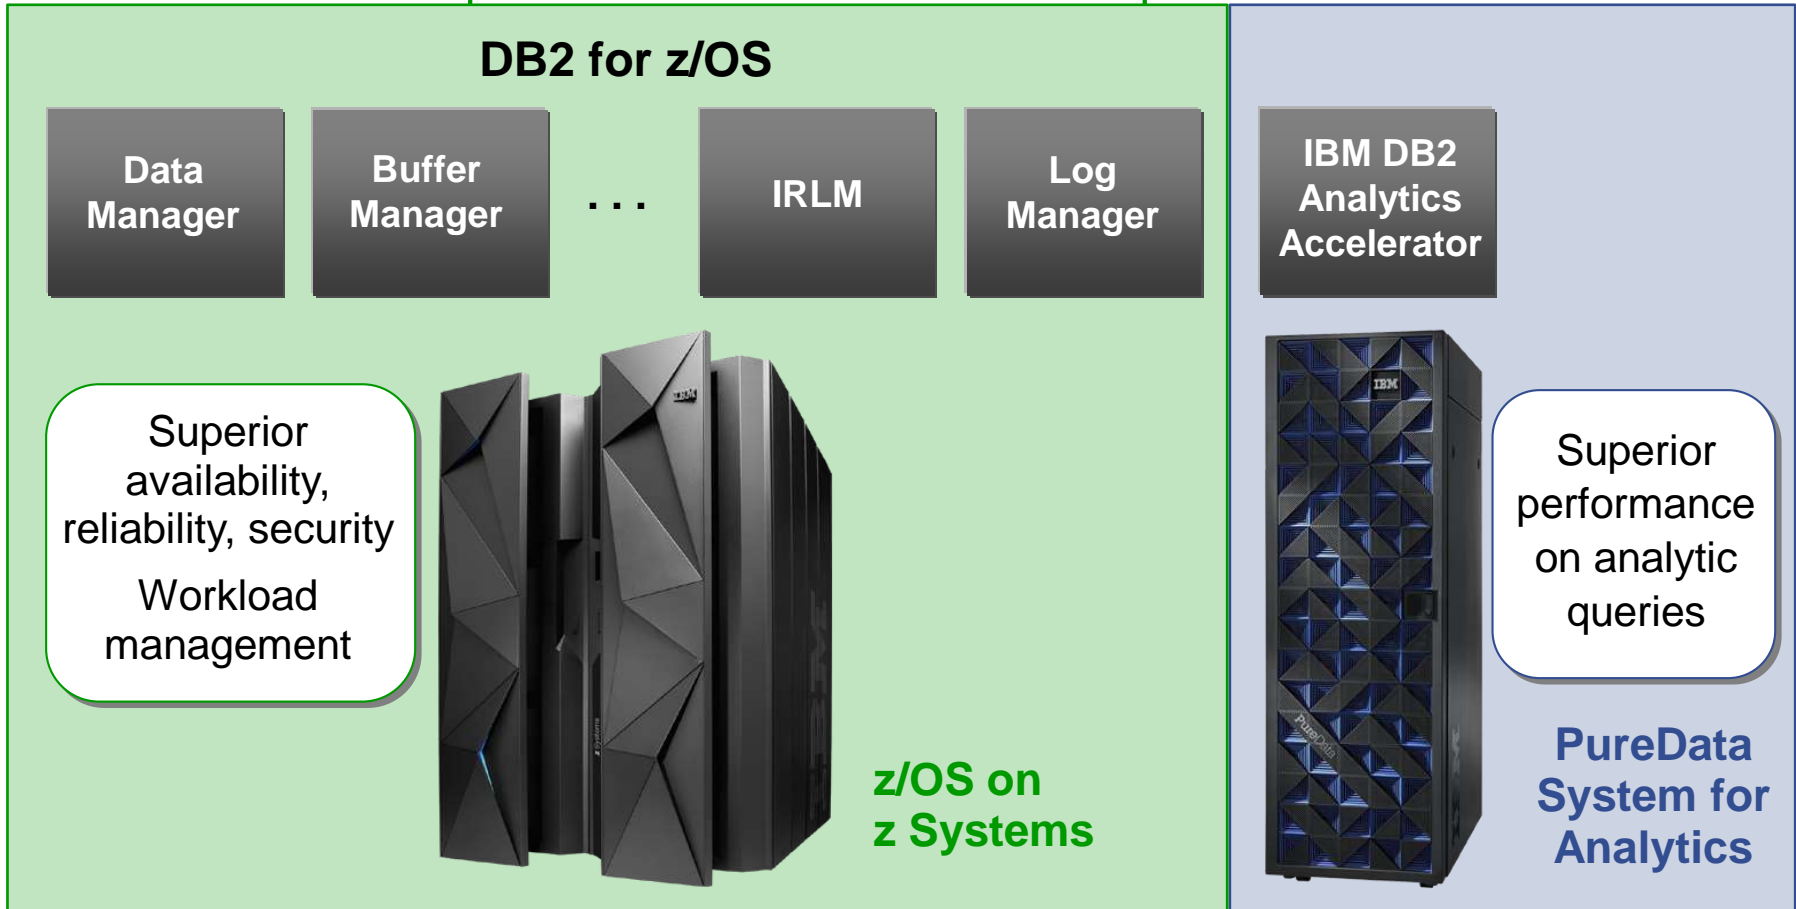



# IBM DB2 Analytics Accelerator

## Product components



# Deep DB2 integration within z Systems





# Connectivity options

Multiple DB2 systems can connect to a single Accelerator



A single DB2 system can connect to multiple Accelerators



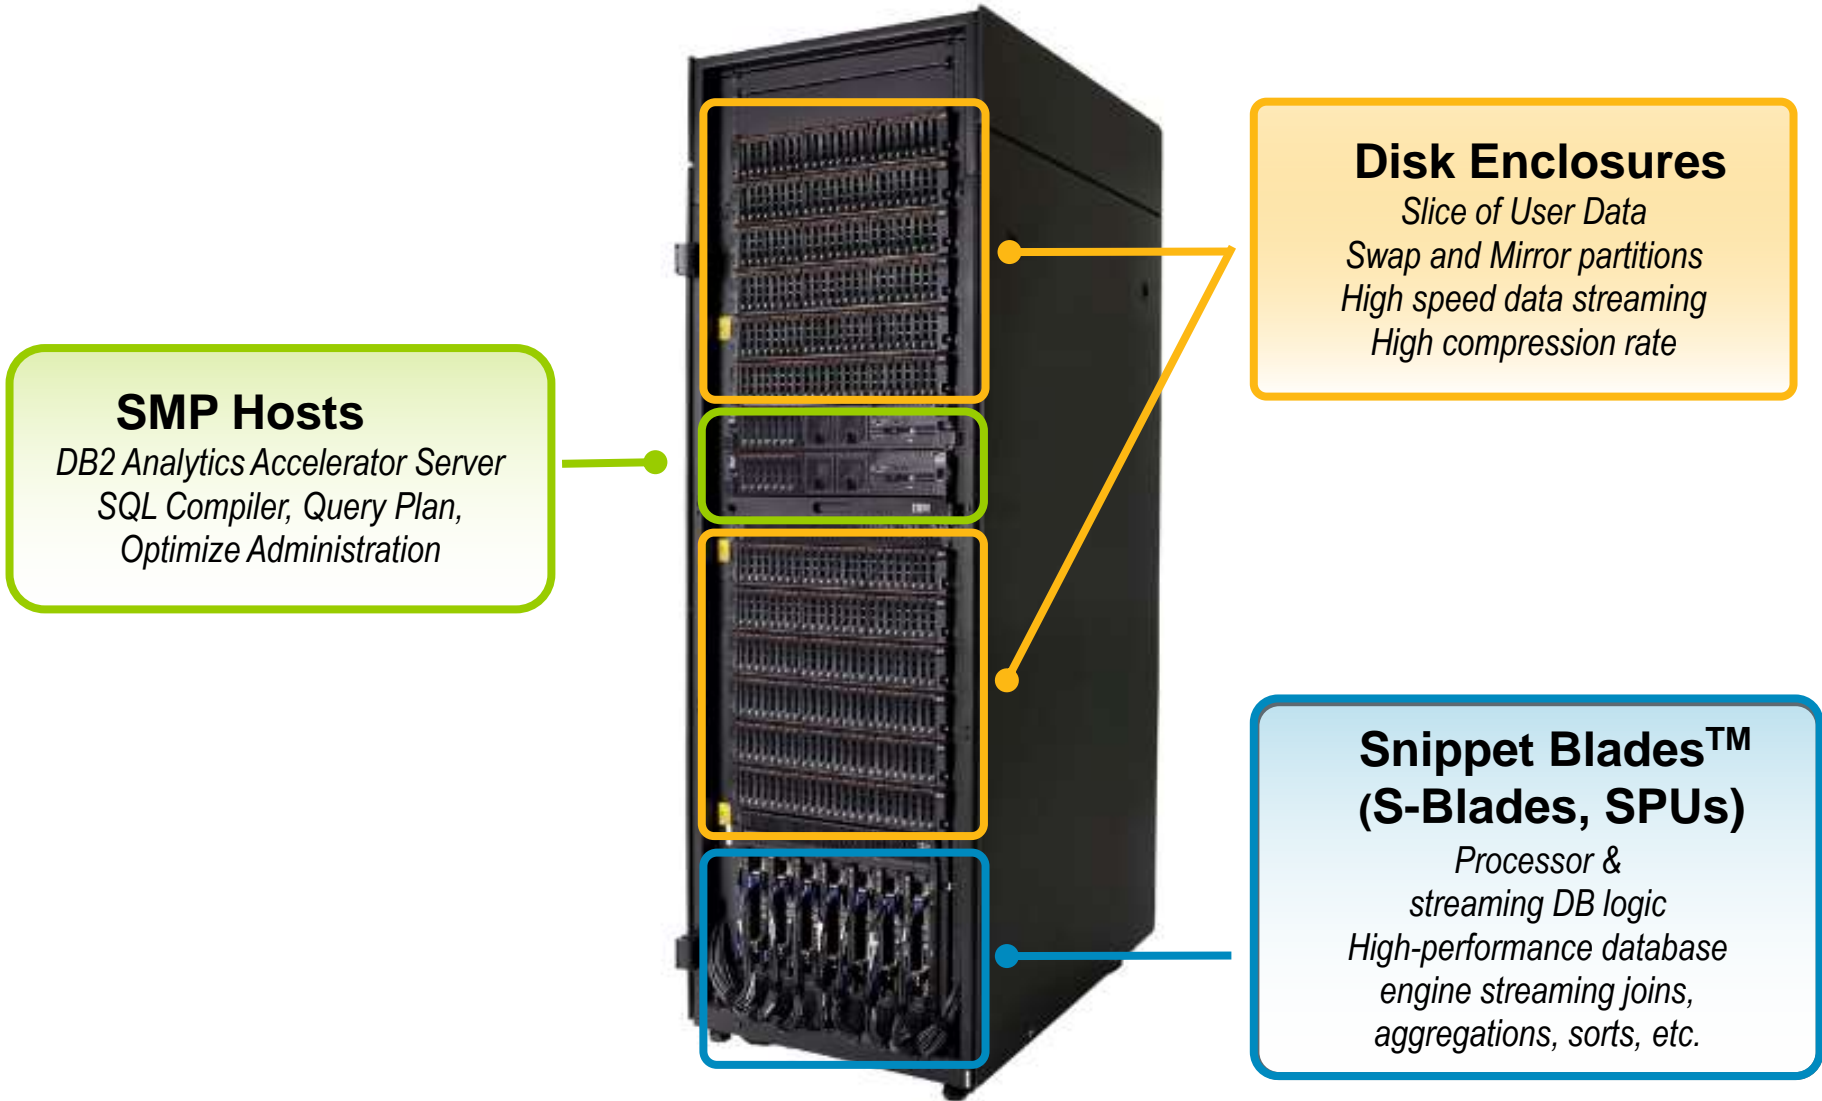
Multiple DB2 systems can connect to multiple Accelerators



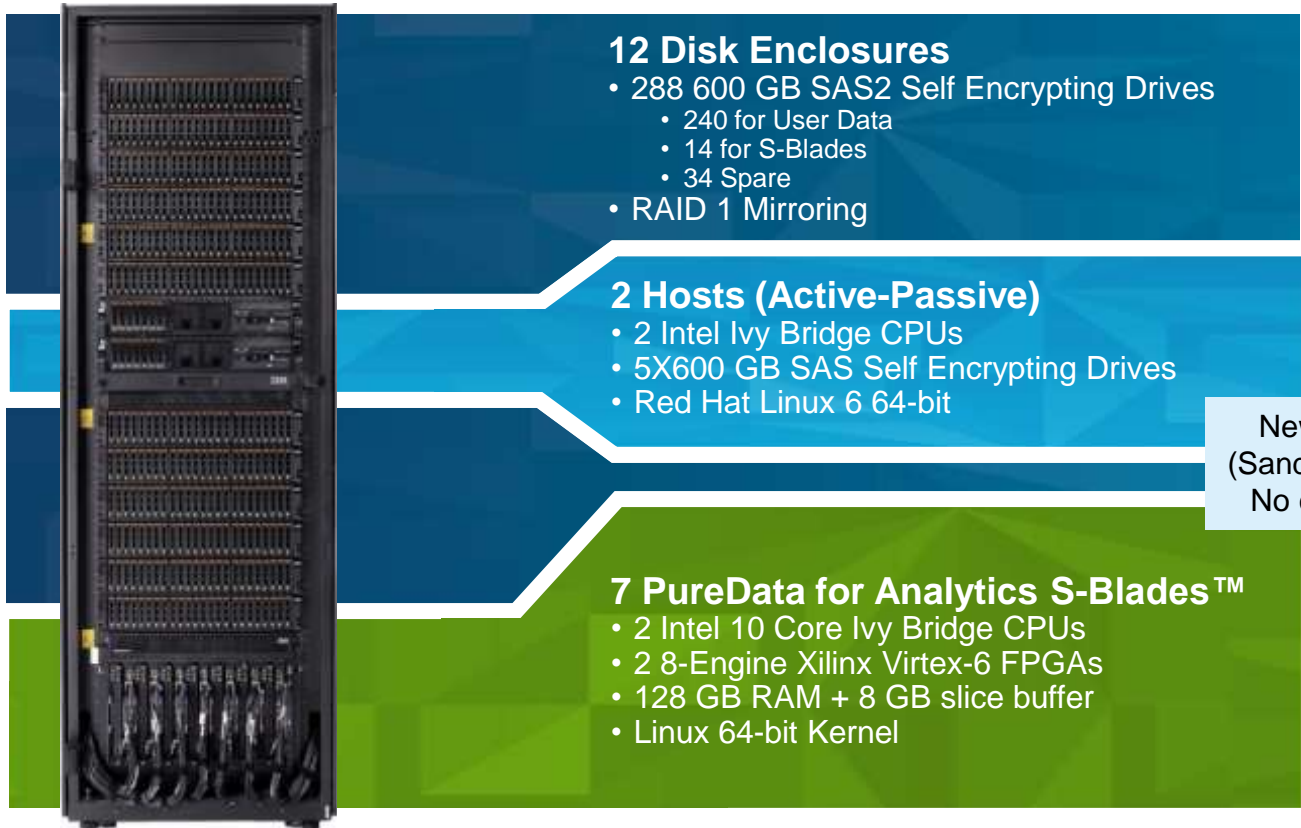
## Full flexibility for DB2 systems:

- residing in the same LPAR
- residing in different LPARs
- residing in different CECs
- being independent (non-data sharing)
- belonging to the same data sharing group
- belonging to different data sharing groups

# PureData System for analytics



# N3001 hardware overview



Scales up to 8 full Racks

Terabyte to Petabyte+ Capacity

New Intel Processor generation (Sandy -> Ivy) for hosts and Blades. No change for Disks and FPGAs

- User Data Capacity: 192 TB<sup>1</sup>
- Data Scan Speed: 478 TB/hr\*

- Power Requirements: 7.5 kW
- Cooling Requirements: 27,000 BTU/hr

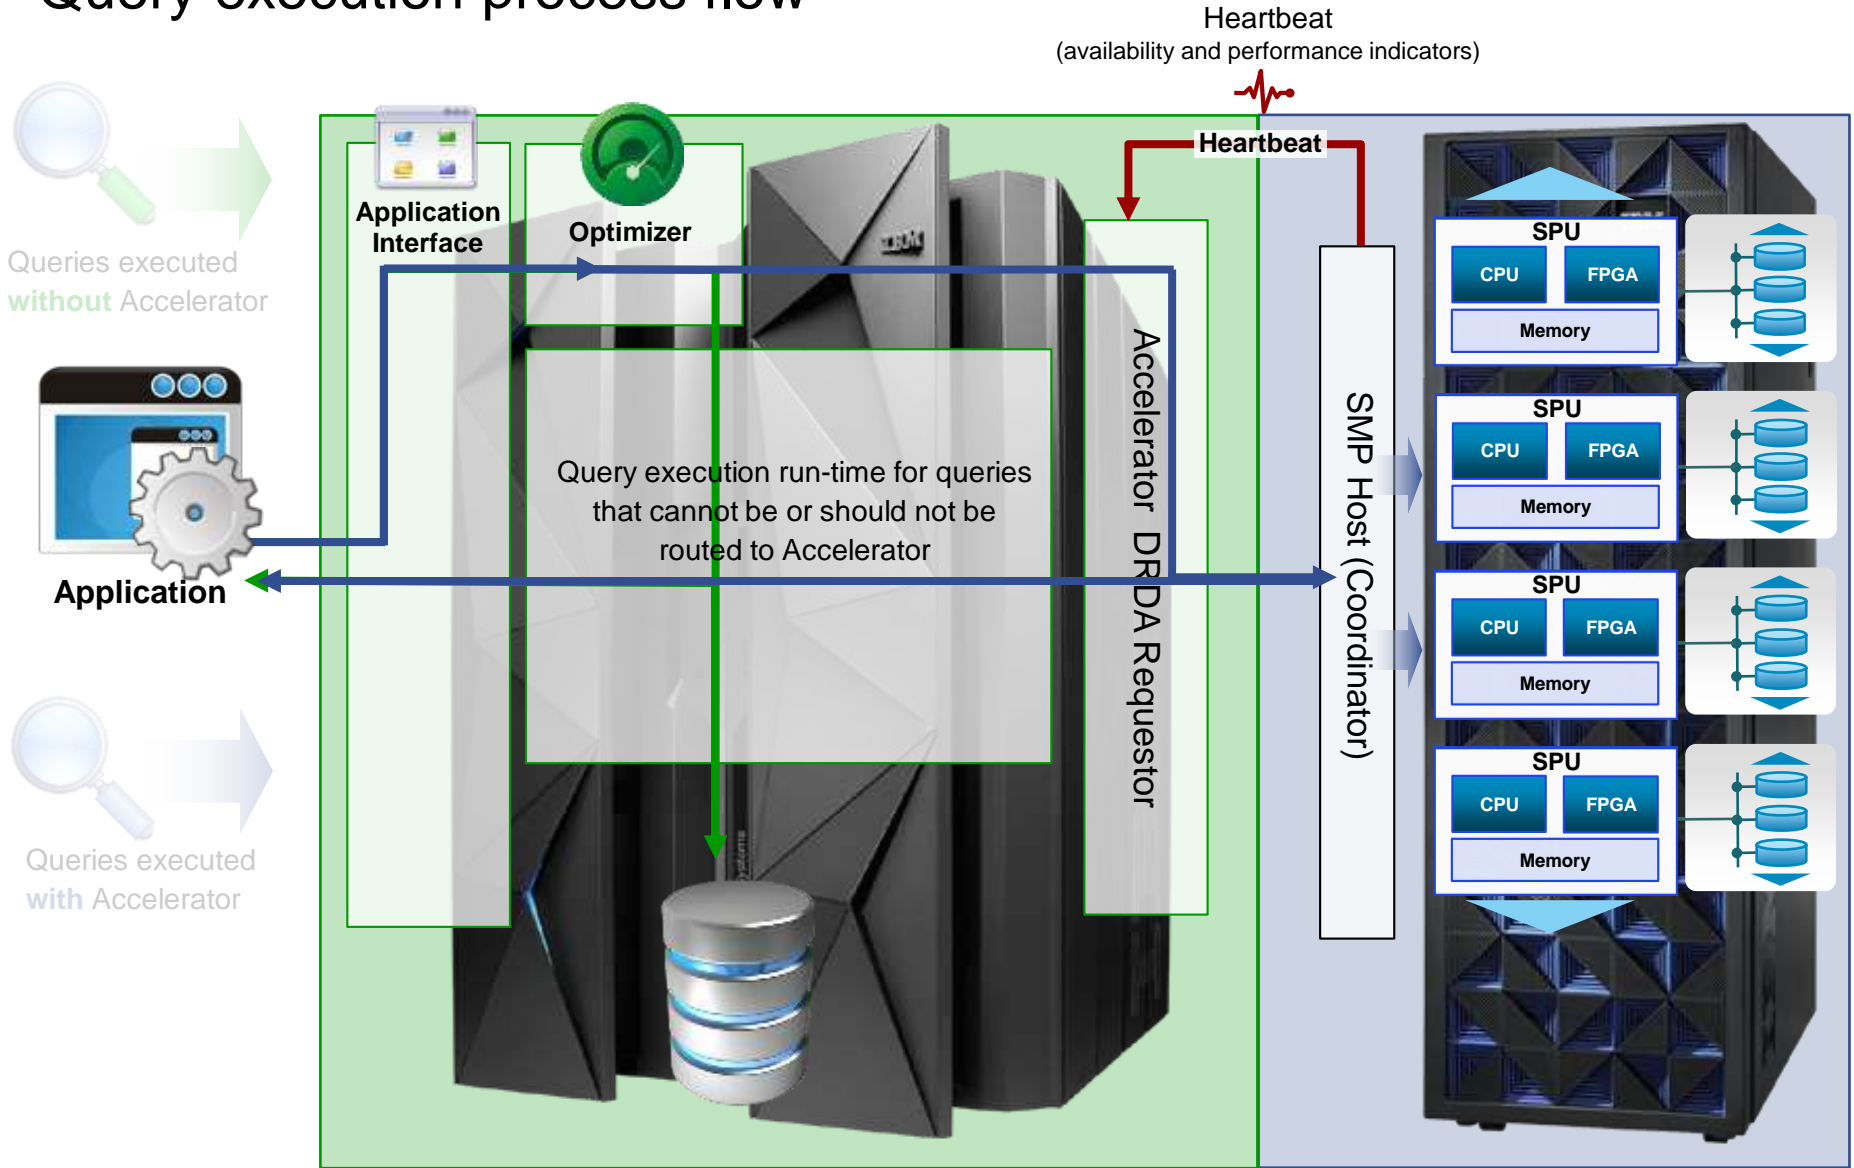
<sup>1</sup>Assuming 4X compression

## IBM DB2 Analytics Accelerator supported models

N3001 Model	001	002	005	010	020	040	080
<b>Cabinets</b>	N/A	¼	½	1	2	4	8
<b>S-Blades</b>	N/A	2	4	7	14	28	56
<b>SPU CPU Cores</b>	40	40	80	140	280	560	1120
<b>SPU FPGA Cores</b>	N/A	32	64	112	224	448	896
<b>Capacity (TB)</b>	4	8	24	48	96	192	384
<b>Effective Capacity (TB)*</b>	16	32	96	192	384	768	1536

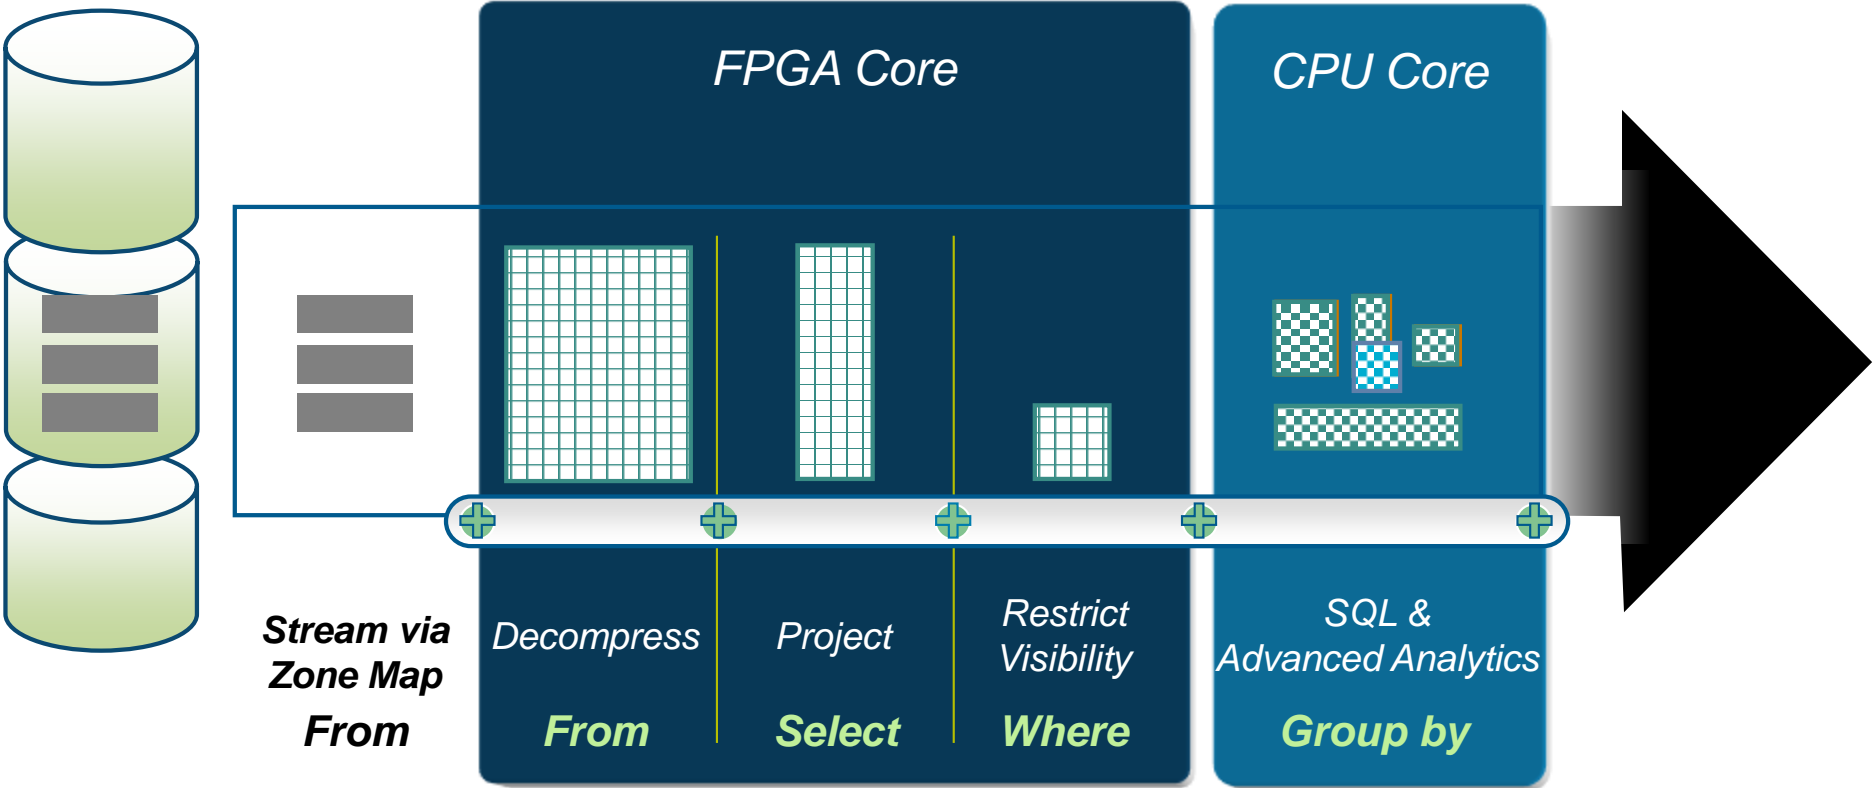
**Predictable, Linear Scalability throughout entire family**

# Query execution process flow





# S-Blade data stream processing



```

Select State, Age, Gender, count(*) From Multibuilding Row Row Count Table Birth Date Birth Date 1960'
And 01/1960' (AFB State) in ('SEL', 'Group by State, Age, Gender, State, State, Age, Gender
State, Age, Gender
    
```

## Routing criteria



- Dynamic and static queries can be accelerated
- DB2 Optimizer decides if query should be sent to Accelerator
  - Dynamic: At execution time
  - Static: At BIND time
- Whole query, not parts of query are accelerated
- Only read queries are considered for acceleration
- Queries within INSERT statements can be accelerated
- Prerequisites for query routing:
  - Accelerator is started
  - All used tables are available on Accelerator
  - Query routing option is specified
    - Via special register, BIND option or ZPARM

## Query routing options (dynamic)

- **Values for special register CURRENT QUERY ACCELERATION**
  - Needs to be set prior to query execution to enable or suppress routing of queries

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 SQLCODE will cause a failback to DB2 during PREPARE or first OPEN. No failback is possible after a successful OPEN of a query.</p>
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.
ALL	A query is routed to the accelerator. If it cannot be executed there, the query fails and a negative return code is passed back to the application.

## Query routing options (static)

- **Values for QUERYACCELERATION bind option**

- Needs to be set at bind time to either bind a query for acceleration or for execution in DB2

Value	Description
NONE	No query is bound for acceleration.
ENABLE	<p>A query is bound for acceleration if it satisfies the acceleration criteria including the cost and heuristics criteria. Otherwise, it is bound for execution 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 bound for acceleration if it satisfies the acceleration criteria including the cost and heuristics criteria. Otherwise, it is bound for execution in DB2.</p> <p>Under certain conditions the query will run on DB2 after it fails in the accelerator. In particular, any negative SQLCODE will cause a failback to DB2 during PREPARE or first OPEN. No failback is possible after a successful OPEN of a query. For a failback DB2 will do an <b>incremental rebind</b> before running the query.</p>
ELIGIBLE	A query is bound for acceleration if it satisfies the acceleration criteria irrespective of the cost and heuristic criteria. Otherwise, it is bound for execution in DB2.
ALL	<b>All</b> queries of the package are bound for acceleration if <b>all queries</b> satisfy the acceleration criteria irrespective of the cost and heuristic criteria. <b>Otherwise, the BIND will fail.</b>

# User interfaces for managing DB2 Analytics Accelerator

## ▪ DB2 stored procedures

- To add/remove accelerator, add/load/remove tables, get table/query/explain information, and more ..
- Execute from e.g. command line, via REXX or C-programs
  - Sample programs available

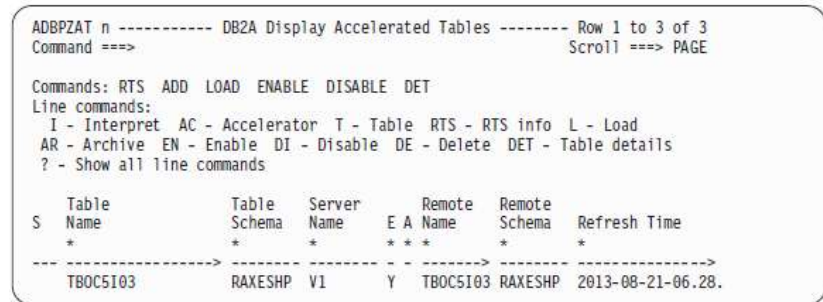
## ▪ DB2 Analytics Accelerator Studio

- Data Studio Plugin
- Externalizes all functions available through the stored procedures in a graphical user interface based on Eclipse
- Visual Explain



## ▪ DB2 11 Administration Tool\*

- Externalizes administration functions available through the stored procedures in ISPF panels
- Separately priced tool, not part of DB2 Analytics Accelerator





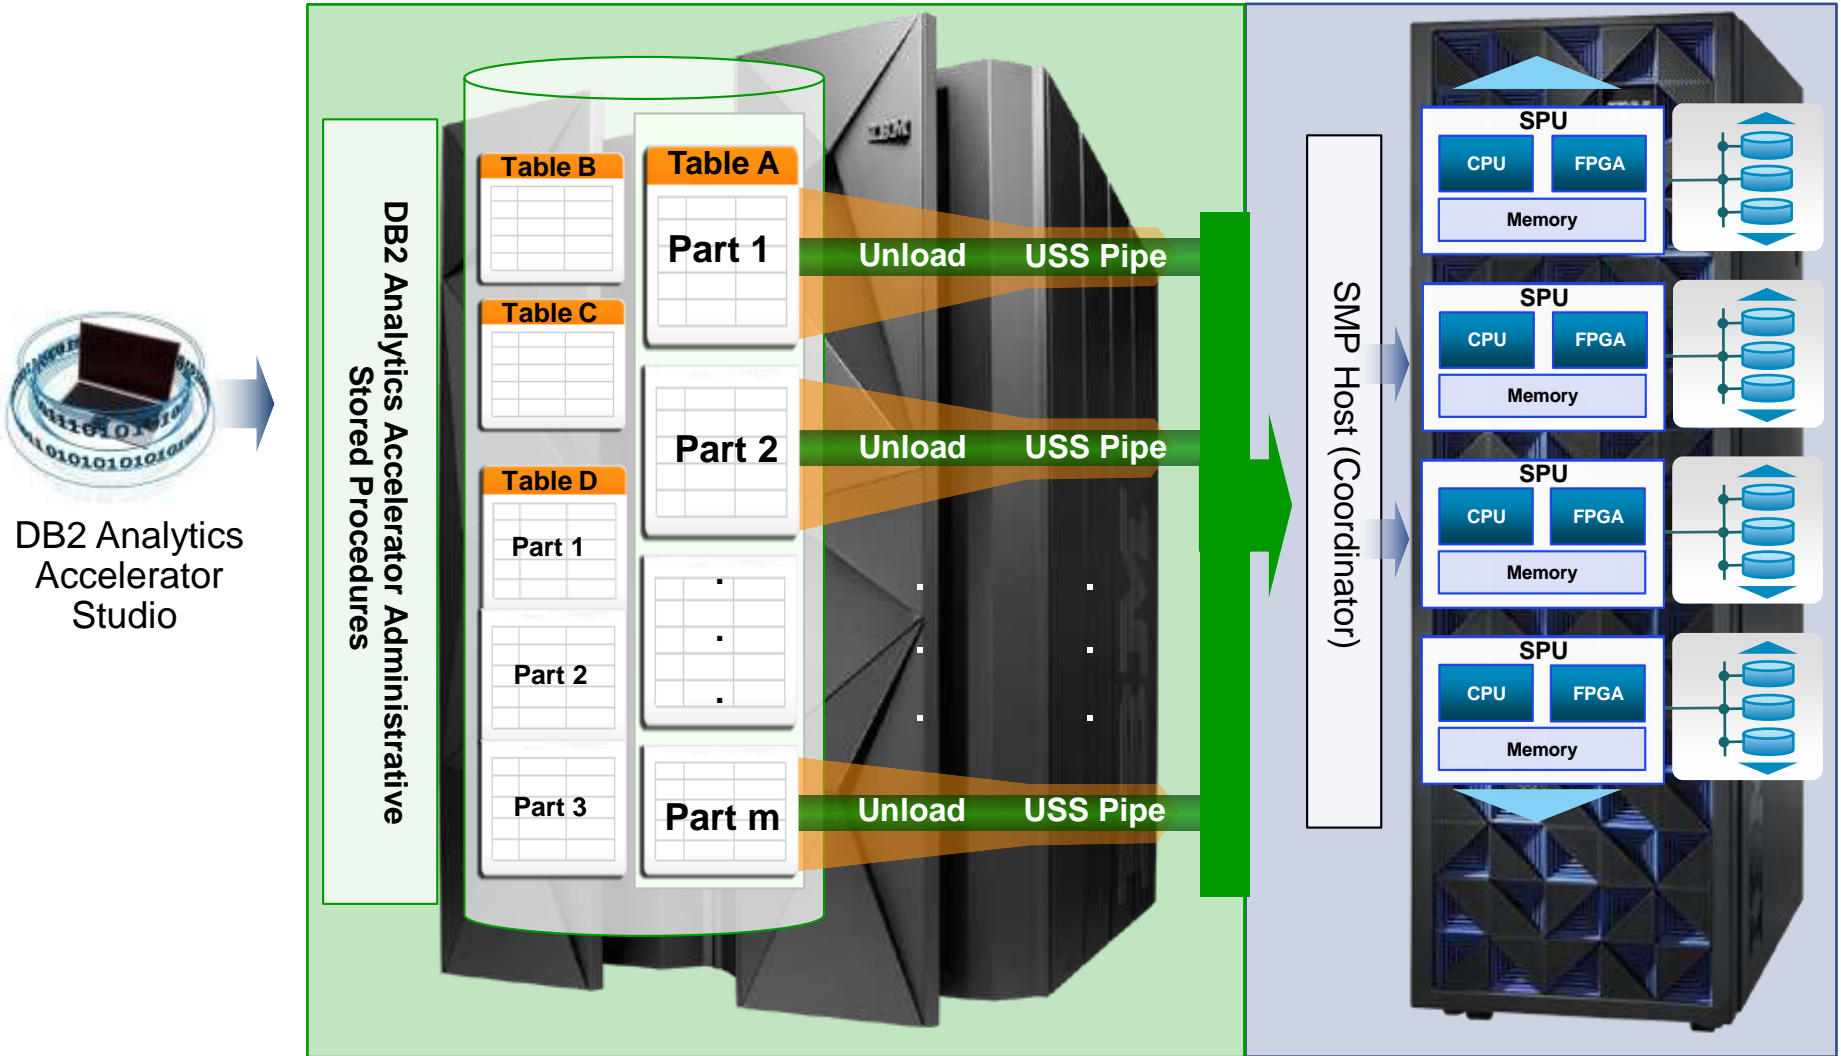
# Data load and update options with DB2 Analytics Accelerator

Synchronization options	Use cases, characteristics and requirements
<p><b>Full table load/refresh</b>  <i>The entire content of a database table is loaded/refreshed for accelerator processing</i></p>	<ul style="list-style-type: none"> <li>▪ Source table data is entirely replaced</li> <li>▪ Multiple sources or complex transformations</li> <li>▪ Smaller, un-partitioned tables</li> <li>▪ Reporting based on consistent snapshot</li> </ul>
<p><b>Table partition load/refresh</b>  <i>For a partitioned database table, selected partitions can be loaded/refreshed for accelerator processing</i></p>	<ul style="list-style-type: none"> <li>▪ Optimization for partitioned warehouse tables, typically appending changes “at the end”</li> <li>▪ More efficient than full table refresh for larger tables</li> <li>▪ Reporting based on consistent snapshot</li> <li>▪ Optionally: automatically load changed partitions only</li> </ul>
<p><b>Incremental Update</b>  <i>Log-based capturing of changes and propagation to Accelerator with low latency (typically few minutes)</i></p>	<ul style="list-style-type: none"> <li>▪ Scattered updates after “bulk” load</li> <li>▪ Reporting on continuously updated data (e.g., an ODS), considering most recent changes</li> <li>▪ More efficient for smaller updates than full table refresh</li> </ul>

## Accelerator data load

- Scope: Table or Partition
- Data is loaded or refreshed using ACCEL\_LOAD\_TABLES stored procedure
  - Tables are loaded sequentially
    - Tables can be loaded in parallel through separate invocation of stored procedure
  - Partitions are loaded in parallel
- Queries can be routed to accelerator when tables or partitions are being refreshed
- Changes to DB2 data can be prevented during load process using LOCKMODE option
  - Ensures consistent snapshot
- Up to ~ 3 TB/h load rate
  - Can vary depending on CPU resources, table partitioning, network bandwidth, concurrently running workload, etc.
- Supports change detection: DB2 automatically determines if table/partition was changed based on DB2 real time statistics
  - Otherwise skips the table/partition in the load request
- DB2 Analytics Accelerator Studio provides:
  - Load/refresh table or table partitions to accelerator
  - Indicates changed tables/partitions

# Accelerator data load



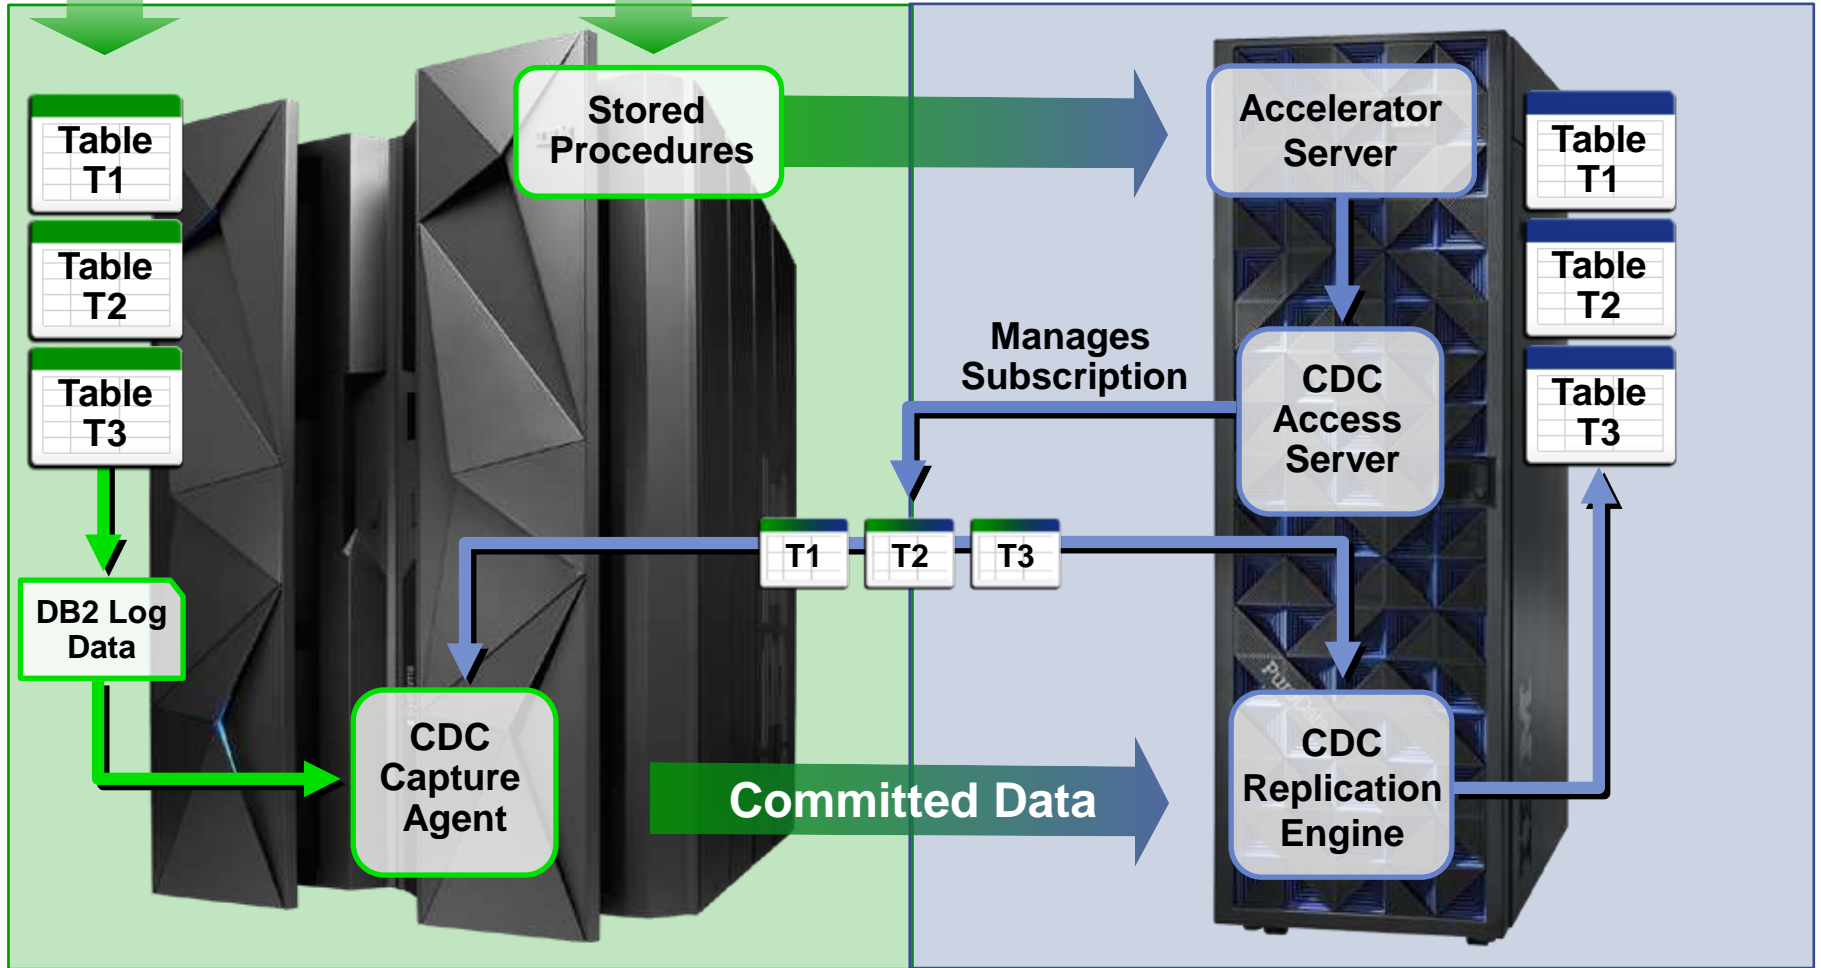
## Incremental update

- Keeps data in DB2 and the Accelerator in synch in near real-time
- Scope: Row
- Based on the Change Data Capture (CDC) component of IBM InfoSphere Data Replication
  - Part of the base DB2 Analytics Accelerator package
- INSERT/UPDATE/DELETE statements captured from DB2 log data and replicated to the Accelerator
- Management functionality integrated into Accelerator Stored Procedures and the DB2 Analytics Accelerator Studio
  - Enable/Disable tables for replication
  - Start/Stop replication
- Allows replication of up to 10 DB2 subsystems simultaneously

# Incremental update architecture

Applications executing I/U/D Statements on replicated tables

Accelerator Users enabling tables for replication





# Load/refresh versus incremental update

## Load/refresh



- Uses existing load mechanism based on UNLOAD utility
- Partition / table granularity
- High performance for larger data volumes
  - Throughput up to ~ 3TB/h, may vary
- High overhead for single-row changes
- Not suitable for frequent, arbitrary changes in the database
- Must be triggered explicitly, allows to preserve “snapshot semantics”
  - Change detection avoids unnecessary loads
- No dependency on DB2 log, can handle not logged changes

## Incremental update



- Uses separate log-reading/shipping technology
- Single row granularity
- High performance for shipping frequent, small and arbitrary changes
  - Throughput double digit GB/h, may vary
- Low latency (minutes)
- Allows to keep accelerator data continuously synchronized
- Not logged changes require explicit use of Accelerator load and explicit knowledge of affected tables or partitions


# Data load and update options with DB2 Analytics Accelerator Loader for z/OS\*

\* DB2 Analytics Accelerator Loader for z/OS, to be purchased separately


Synchronization options	Use cases, characteristics and requirements
<p><b>Group Consistent Load</b></p> <p><i>Loads groups (or sets) of operational DB2 tables to the accelerator at a user-specified time. One time specified for all tables.</i></p>	<ul style="list-style-type: none"> <li>▪ Uses DB2 image copies or logs as input</li> <li>▪ Transaction Consistent: Uncommitted transactions at the specified time are not loaded to accelerator, i.e. if Load is run after parent update but before child update, the update to parent table will not be loaded to accelerator</li> <li>▪ Maintain availability of tables during the load process, no tables locked</li> <li>▪ Loader can optionally create a new FlashCopy Image Copy</li> <li>▪ Load or refresh accelerator with zero impact to business critical data</li> <li>▪ Ability to load to accelerator with historical data</li> </ul>
<p><b>Dual Load</b></p> <p><i>Loads data from a file into both DB2 and the Accelerator in parallel.</i></p> <p><b>Accelerator only Load</b></p> <p><i>Loads data from a file directly into Accelerator (no load in DB2).</i></p>	<ul style="list-style-type: none"> <li>▪ Dual Load provides TCO savings with CPU and Elapsed Time improvements             <ul style="list-style-type: none"> <li>• Compared to loading to DB2 first and then loading to the Accelerator</li> </ul> </li> <li>▪ Accelerator only load provides significant CPU and DASD savings on the source (DB2)</li> <li>▪ Extracted data can come from other DB2 system, VSAM, IMS, Oracle, etc.</li> <li>▪ Exploits zIIP processor to reduce cost of loading to Accelerator</li> <li>▪ File must be compatible for input into the DB2 LOAD utility</li> <li>▪ Consider ramifications for backup/restore of DB2 data when using Accelerator only load – the data does not reside in DB2 anymore</li> </ul>

# DB2 Analytics Accelerator Loader for z/OS\*


Load data from image copy or other sources into DB2 and the Accelerator in parallel



Direct loading into Accelerator, point-in-time loading (i.e. roll-forward to point-in-time from log files)



Load data into the Accelerator only (meta data required in DB2)



# High Performance Storage Saver

Reducing disk storage cost by archiving data in the Accelerator and at the same time maintaining the excellent performance for analytical queries




- Supported for partitions of range partitioned tables
  - Either all or a subset of partitions can be archived
- Archive using ACCEL\_ARCHIVE\_TABLES stored procedure
  - Archiving to multiple Accelerators possible
- Archived partitions can be restored using ACCEL\_RESTORE\_ARCHIVE\_TABLES stored procedure
- Queries read archived partitions if GET\_ACCEL\_ARCHIVE special register or ZPARM is set in addition to QUERY\_ACCELERATION special register or ZPARM
- Queries return a combined result for recent partitions (loaded to the Accelerator) and archived partitions if both special registers or ZPARMs are set

# High Performance Storage Saver

*Storing historical data in Accelerator only*

**?**  
Query from Application to either Active only or Active and Historical



## Support for partitioned-by-range tables, for example:

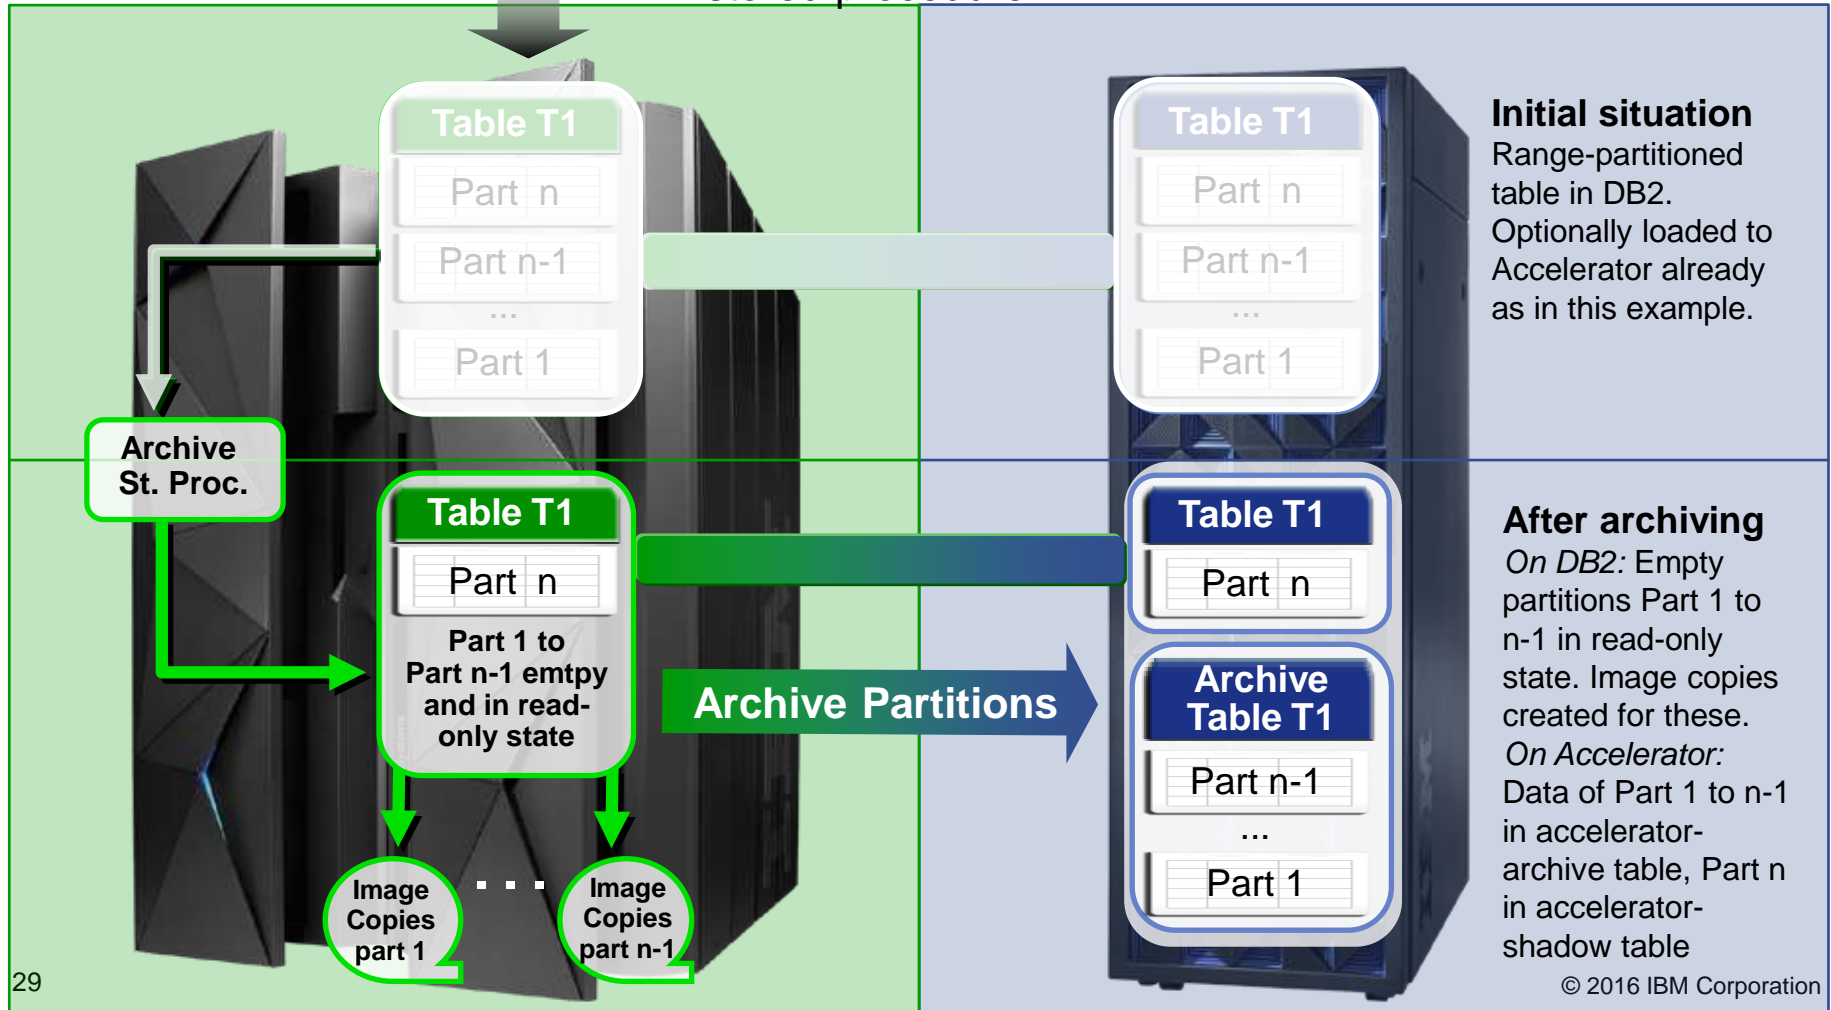
Time-based partitions, where only the recent partitions are used in a transactional context (frequent data changes, short running queries), but the entire table is used for analytics, regulatory/audit purposes, etc. (data intensive, complex queries)



# High Performance Storage Saver process



User requests to archive partitions 1 to n-1 using ACCEL\_ARCHIVE\_TABLES stored procedure



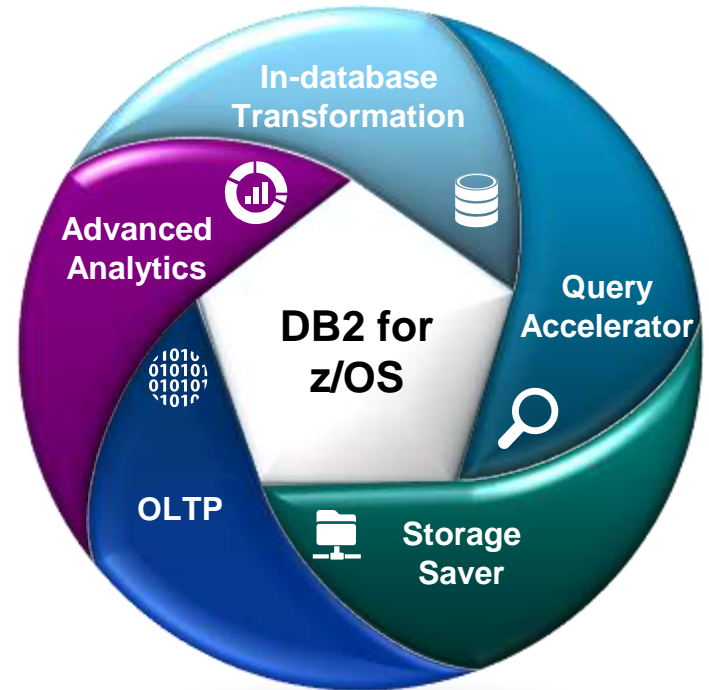
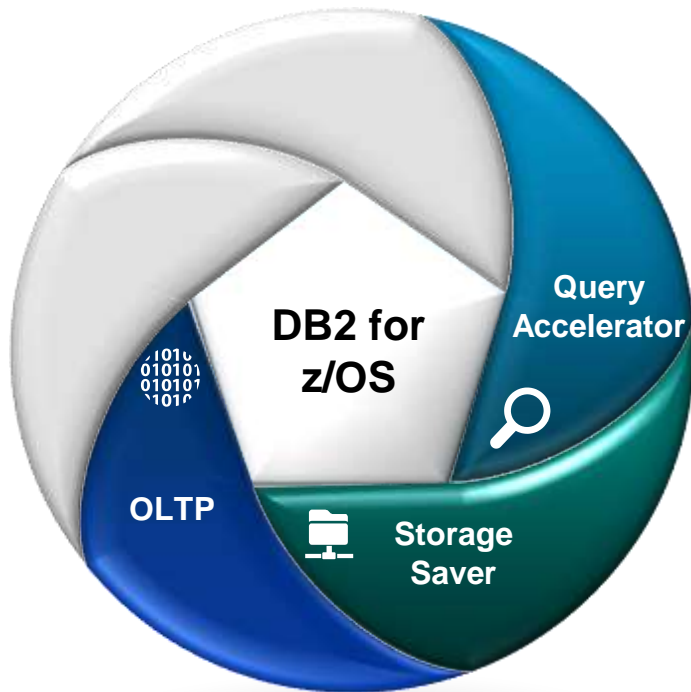



# Agenda

- Value Proposition
- Technical Overview
- **New Version 5.1 Functionality**
- Managing and Monitoring
- Use Cases and Customer Experiences
- Strategy and Evolution

# IBM DB2 Analytics Accelerator Version 5.1

Next iteration on the road to enable DB2 transition into a truly universal DBMS that provides best characteristics for both OLTP and analytical workloads.



## DB2 Analytics Accelerator Version 5.1

*Enabling real-time analytic solutions on a single, integrated system combining transactional data, historical data and predictive analytics*

- **Business agility** through **simplified architecture** with in-database transformation and multi-step processing
- **Real-time, actionable business processes** through in-database analytics
- **Insight into now to maximize business opportunities** through enterprise Incremental Update enhancements
- **Extended security** through encryption of data at rest and in motion while taking advantage of the **renowned built-in security of z Systems**
- **Enriched systems management capabilities** and improved serviceability through IBM Call Home



# Accelerator-only tables

## *Supporting in-database transformation and multi-step processing*

Introduction of **Accelerator-only tables (AOT)** to store intermediate or final results of data transformation or reporting processes



- Can be used to store a set of data in DB2 Analytics Accelerator only, not on DB2 for z/OS, without using the High Performance Storage Saver functionality
- Accelerate in-database data transformations and data movement processes
- Reduced need of data movement processes to other platforms for data transformation purposes
- Enables multi-step reporting on the Accelerator
- Saves disk space and CPU cost on z Systems currently used for transformations and reporting steps
- Allow data preparation steps for data mining and other advanced analytics to execute on the Accelerator

# 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



## Accelerator-only tables – Technical basics

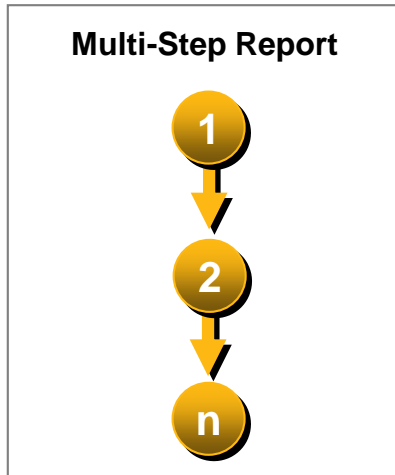
- AOTs are created and dropped using DB2 DDL statements (CREATE; DROP)
  - Accelerator must be started
  - QUERY ACCELERATION behavior may have any value during CREATE/DROP
  - Syntax:
    - *CREATE TABLE MYTABLE (...) IN ACCELERATOR <ACCEL1>;*
    - *DROP TABLE MYTABLE;*
- Recommended to create a database in DB2 to be used for the AOTs
  - *CREATE TABLE MYTABLE (...) IN ACCELERATOR <ACCEL1> IN DATABASE MYDB;*
  - Usual authorization necessary to create objects in database
- Queries using AOTs can only run on the Accelerator
  - QUERY ACCELERATION behavior must be set to ENABLE/ELIGIBLE/ALL
- AOTs can be subject to INSERT/UPDATE/DELETE operations on other accelerated tables archived tables or AOTs
- Dynamic and static SQL can be used with AOTs



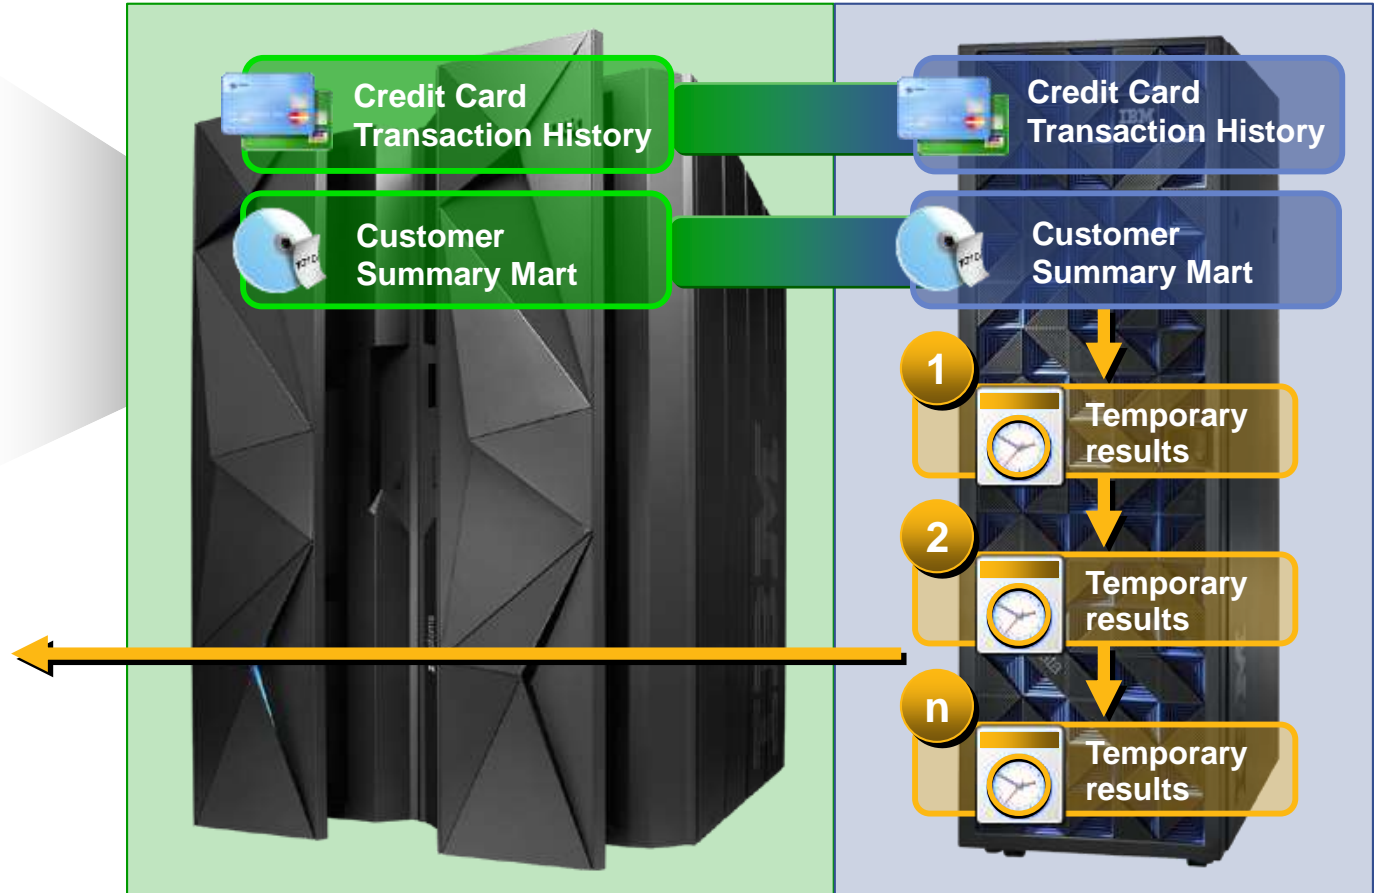
# Multi-step reporting applications with DB2 for z/OS

*With Accelerator-only tables: Temporary objects and processing on the Accelerator*

## Reporting Application

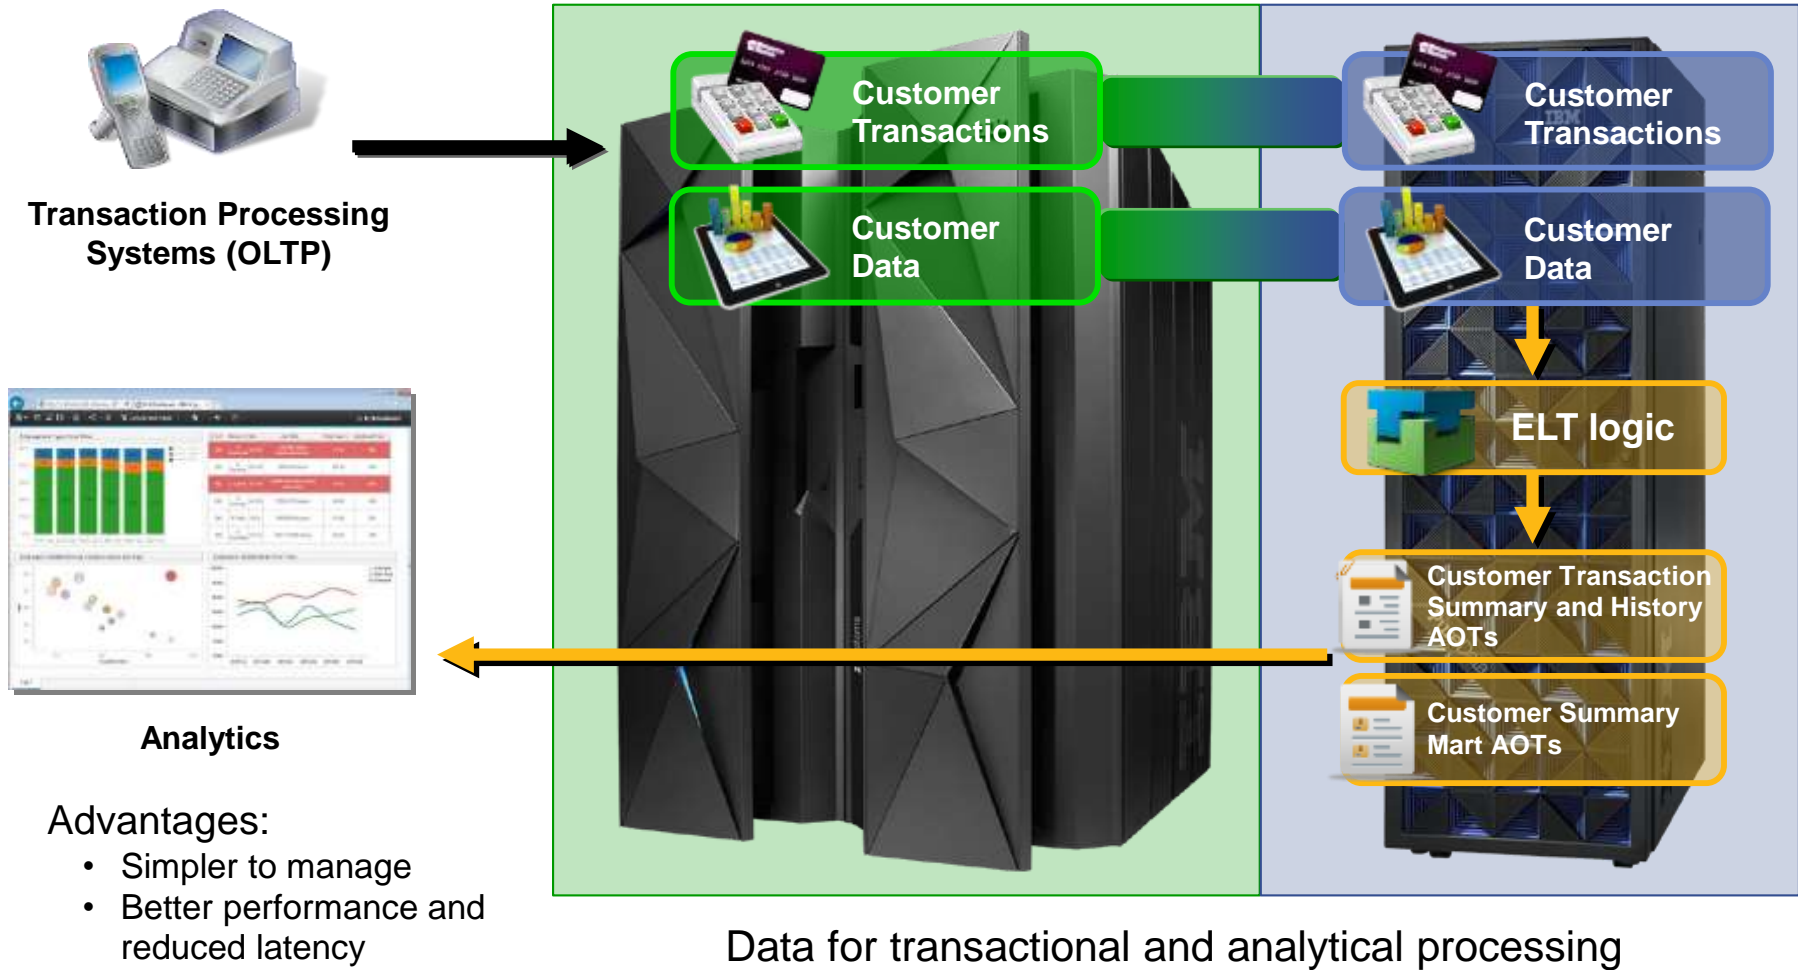


Reports and Dashboards



Data for transactional and analytical processing

# Using Accelerator-only tables and ELT logic in the Accelerator



## In-database analytics

*Enable acceleration of predictive analytics applications*

In-database analytics enables SPSS/Netezza Analytics (INZA) data mining and in-database modeling to be processed within the Accelerator



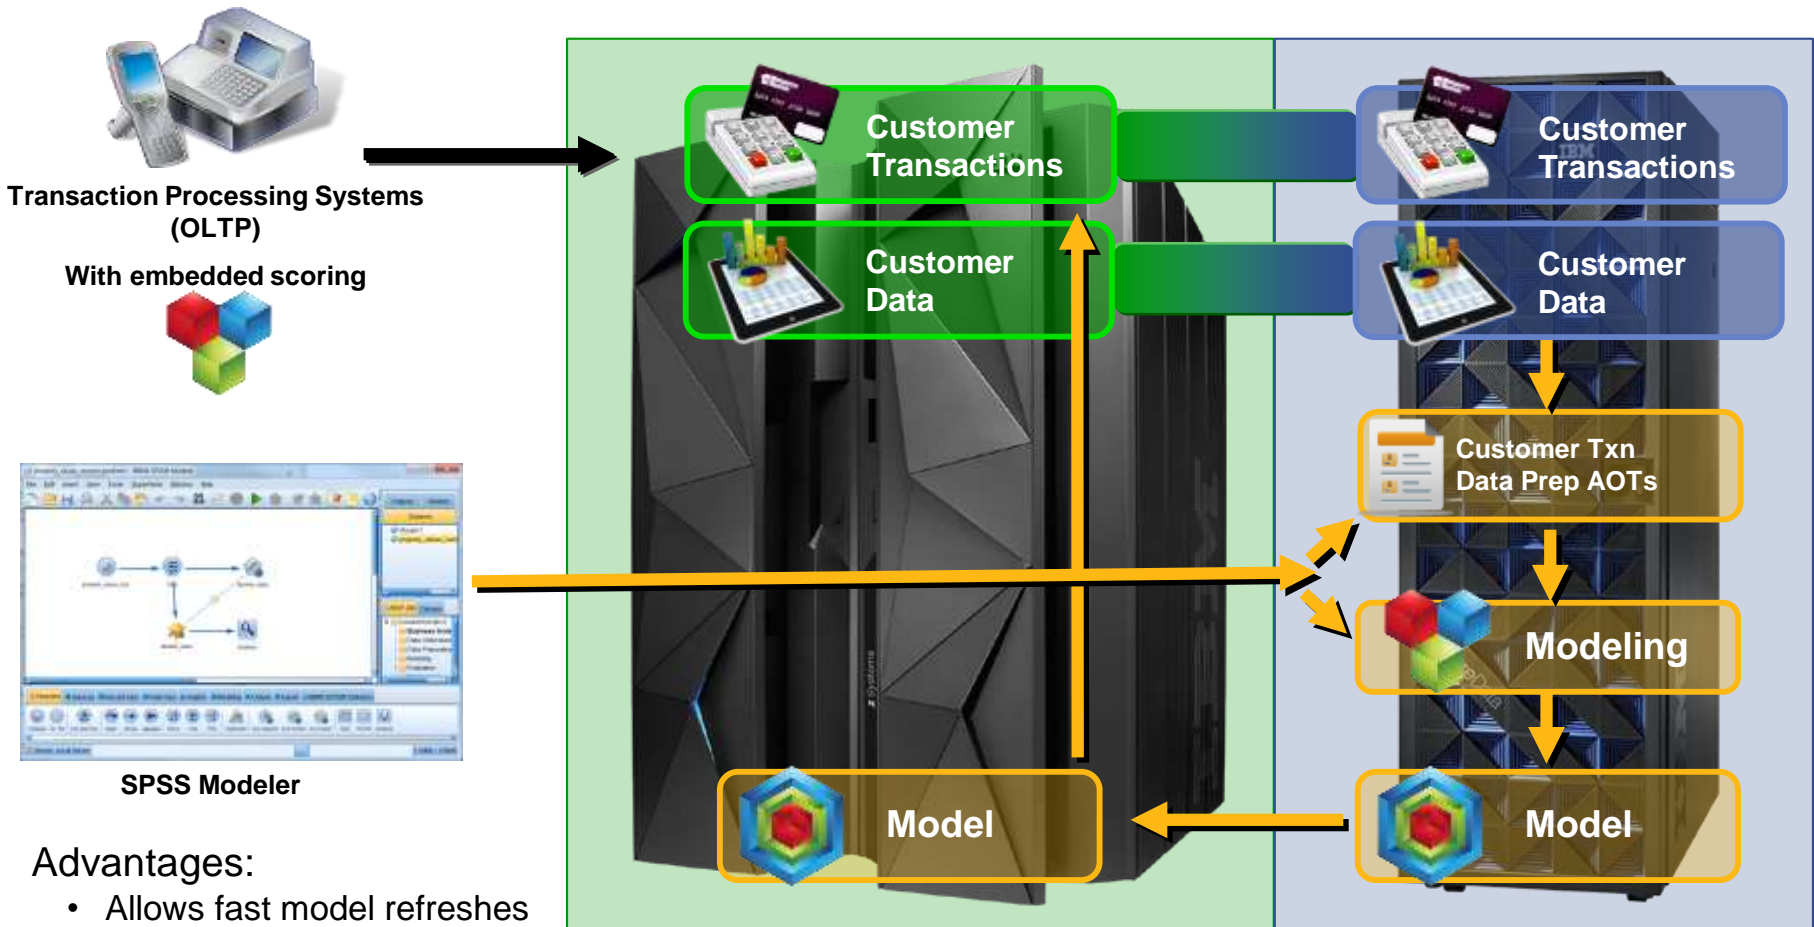
- Accelerates SPSS/Netezza Analytics (INZA) data mining and in-database modeling
- Allows frequent model refreshes to enable adequate scoring
- Reduced need of data movement processes to other platforms for predictive analytics purposes
- Supports the full lifecycle of a real-time analytics solution on a single, integrated system, combining transactional data, historical data, and predictive analytics

## In-database analytics – Technical basics

- Set of stored procedures of IBM Netezza In-Database Analytics Package (INZA) are available on the Accelerator for modeling purposes, including:
  - Decision Tree
  - Regression Tree
  - Naive Bayes
  - K-means Clustering and TwoStep Clustering
- Stored procedures use accelerator-shadow tables or accelerator-only tables as input and create accelerator-only tables as output
- DB2 for z/OS contains stored procedure wrappers to enable invocation of the stored procedures from SPSS Modeler 17
  - Actual stored procedure execution takes place on Accelerator

# In-database analytics

*Data preparation (using AOTs) and SPSS modeling in the Accelerator*



Data for transactional and analytical processing

**Advantages:**

- Allows fast model refreshes
- Ensures adequate scoring
- Better performance and reduced latency

# Agenda

- Value Proposition
- Technical Overview
- New Version 5.1 Functionality
- **Managing and Monitoring**
- Use Cases and Customer Experiences
- Strategy and Evolution

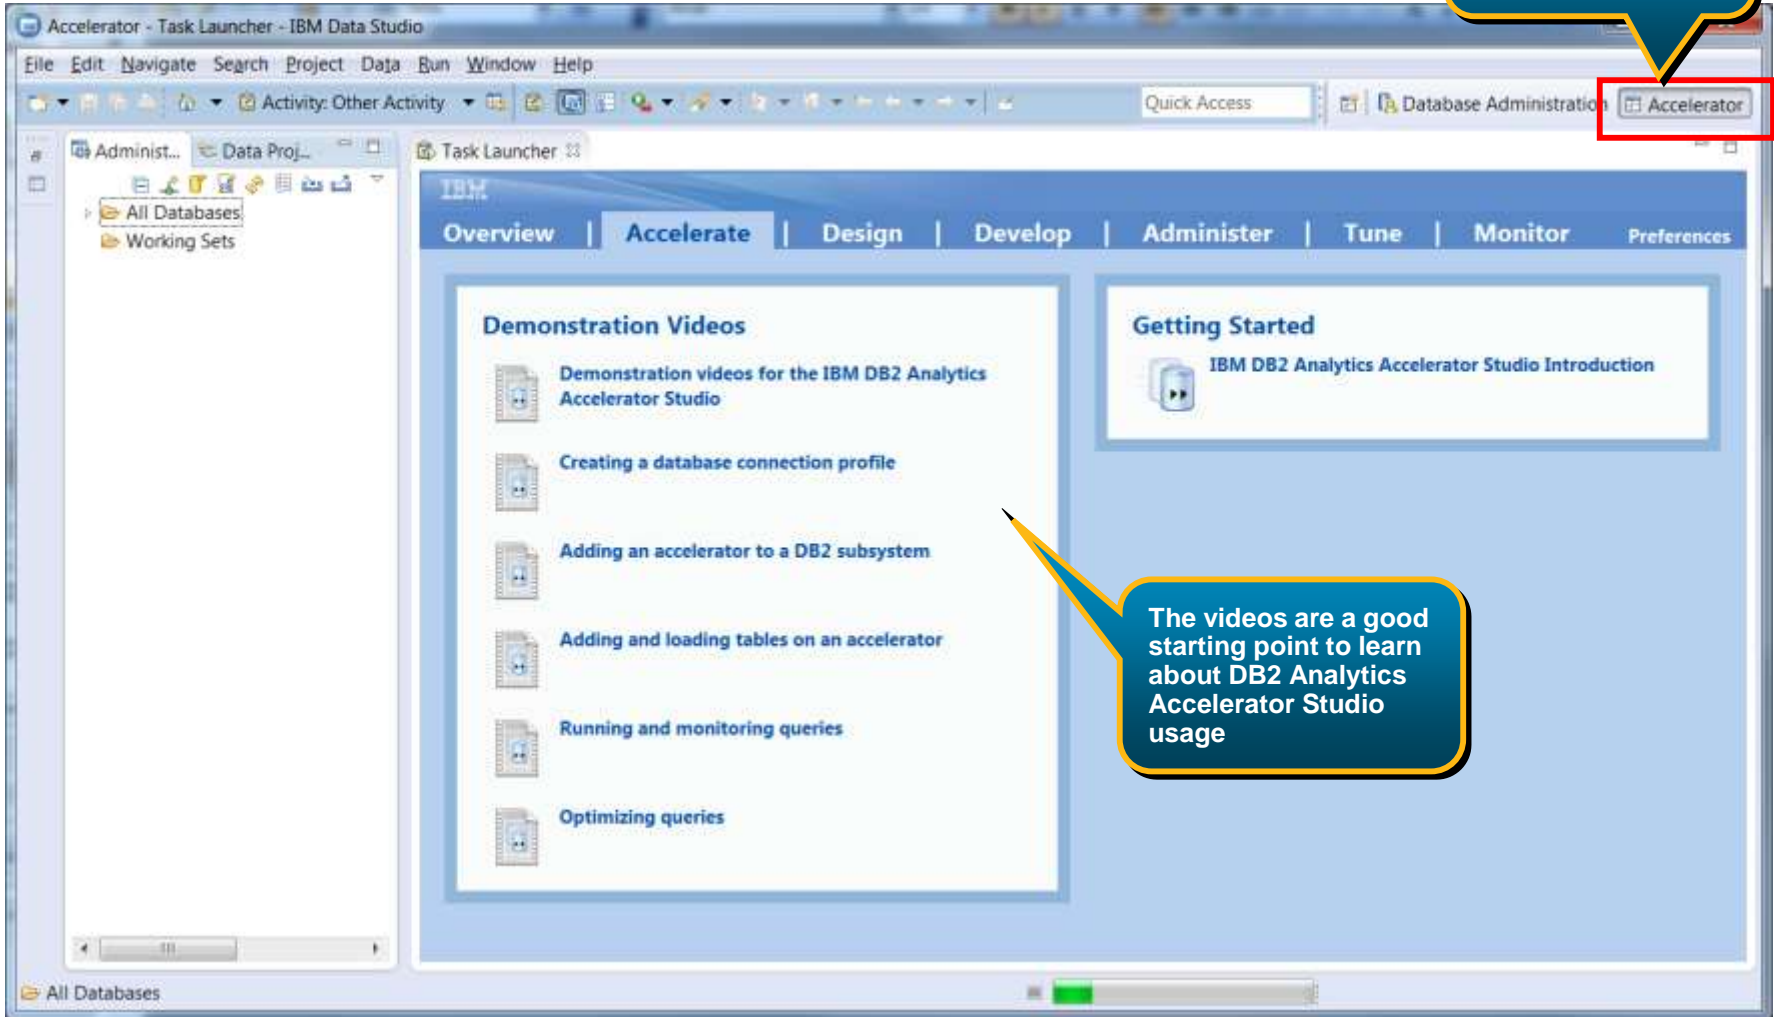


# DB2 Analytics Accelerator – Stored procedures

Name	Description
ACCEL_ADD_ACCELERATOR ACCEL_ADD_ACCELERATOR2	Pairing an accelerator to a DB2 subsystem
ACCEL_TEST_CONNECTION	Check connectivity from DB2 procedures to the accelerator
ACCEL_REMOVE_ACCELERATOR	Removing an accelerator from a DB2 subsystem and cleanup resources on accelerator
ACCEL_UPDATE_CREDENTIALS	Renewing the credentials (authentication token) in the accelerator
ACCEL_ADD_TABLES	Add a set of tables to the accelerator
ACCEL_ALTER_TABLES	Alter table definitions for a set of tables on the accelerator (distribution and organizing keys)
ACCEL_REMOVE_TABLES	Remove a set of tables from the accelerator
ACCEL_ARCHIVE_TABLES	Moves table partitions from DB2 to a storage save on the accelerator
ACCEL_RESTORE_ARCHIVE_TABLES	Restore the data of moved partitions to their original locations
ACCEL_GET_TABLES_INFO	List set of tables on the accelerator together with detail information
ACCEL_GET_TABLES_DETAILS	Collects information about a set of tables with regard to data changes or move operations with High Performance Storage Saver
ACCEL_LOAD_TABLES	Load/Reload/Update data from DB2 into a set of tables on the accelerator
ACCEL_SET_TABLES_ACCELERATION	Enable or disable a set of tables for query off-loading
ACCEL_SET_TABLES_REPLICATION	Enable or disable incremental updates for one of more tables on the accelerator
ACCEL_CONTROL_ACCELERATOR	Controlling the accelerator tracing, collecting trace and detail of the accelerator (software level etc.)
ACCEL_UPDATE_SOFTWARE	Update software on the accelerator (transfer versioned software packages or apply an already transferred package, also list software both on z/OS and accelerator side)
ACCEL_GET_QUERY_DETAILS	Retrieve statement text and query plan for a running or completed Netezza query
ACCEL_GET_QUERY_EXPLAIN	Generate and retrieve Netezza explain output for a query explained by DB2
42 ACCEL_GET_QUERIES	Retrieve active and/or history query information from accelerator

# Data studio – Accelerator perspective

Open Accelerator Perspective using menu Window->Open Perspective



The videos are a good starting point to learn about DB2 Analytics Accelerator Studio usage

# View Accelerator details and manage it

**Accelerator: STRIPER @ CC825**

Acceleration: Started [Stop](#) Credentials valid since: 1/16/14 2:03 PM [Update](#)  
 Status: Online Trace: DEFAULT / OFF [Configure](#) [Save](#) [Clear](#)  
 Used space: 2.73 GB of 22.9 TB Active queries: 0 (0 queued)  
 Replication: Started [Stop](#) Replication latency: High [Show events](#)

**Monitoring**

**Queries:** Successful: 192 Failed: 0  
**Queue:** Max. length: 0 Avg. wait time: 4 ms Max. wait time: 28.6 s  
**CPU cost:** Query execution: 1.44 s Data maintenance: 71.3 s Replication: 55.5 s

**About**

**Tables (93 of 155 loaded / 117 of 155 enabled for acceleration)**

[Add...](#) [Alter Keys...](#) [Remove](#) [Load...](#) Acceleration Storage Saver Replication Cancel Tasks

Name	Size	Acceleration	Last Load	Storage Saver Parti...	Replication Since	Distribution I
<b>BMB</b>	2 MB	7 of 8	7 of 8 tables	0 of 8 tables	0 of 8	-
BGDTMAE	-	Enabled	3/25/15 2:11 PM	-	Disabled	Random
BGDPEN	-	Enabled	3/25/15 2:11 PM	-	Disabled	Random

**Query Monitoring**

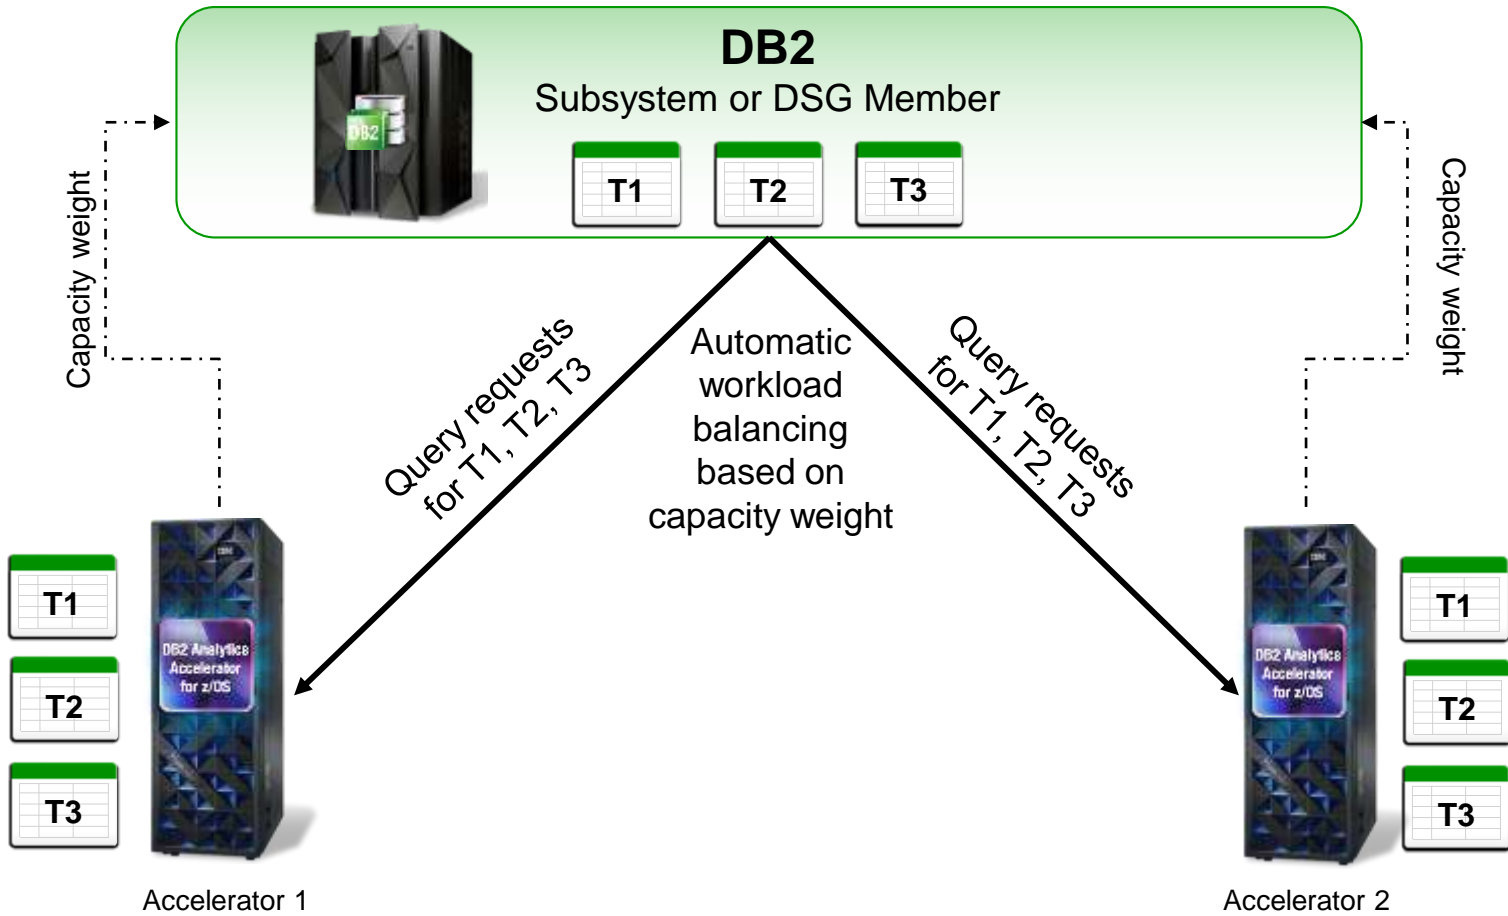
[Show SQL...](#) [Show Plan...](#) [Re-Run](#) [Cancel](#) View: All Queries Show: All

SQL Text	User ID	Start Time	State	Wait Time	Execution Ti...
SELECT SUM(SS_EXT_DISCOUNT_AMT) FROM TPCH6.SAL...	TPCH6	6/19/15 6:33:37 ...	Successful	0 seconds	0 seconds
SELECT COUNT(DISTINCT(SS_TICKET_NUMBER)) FROM T...	TPCH6	6/19/15 6:27:34 ...	Successful	0 seconds	0 seconds
SELECT COUNT(*) FROM TPCH6.SALES	TPCH6	6/19/15 6:22:47 ...	Successful	0 seconds	0 seconds

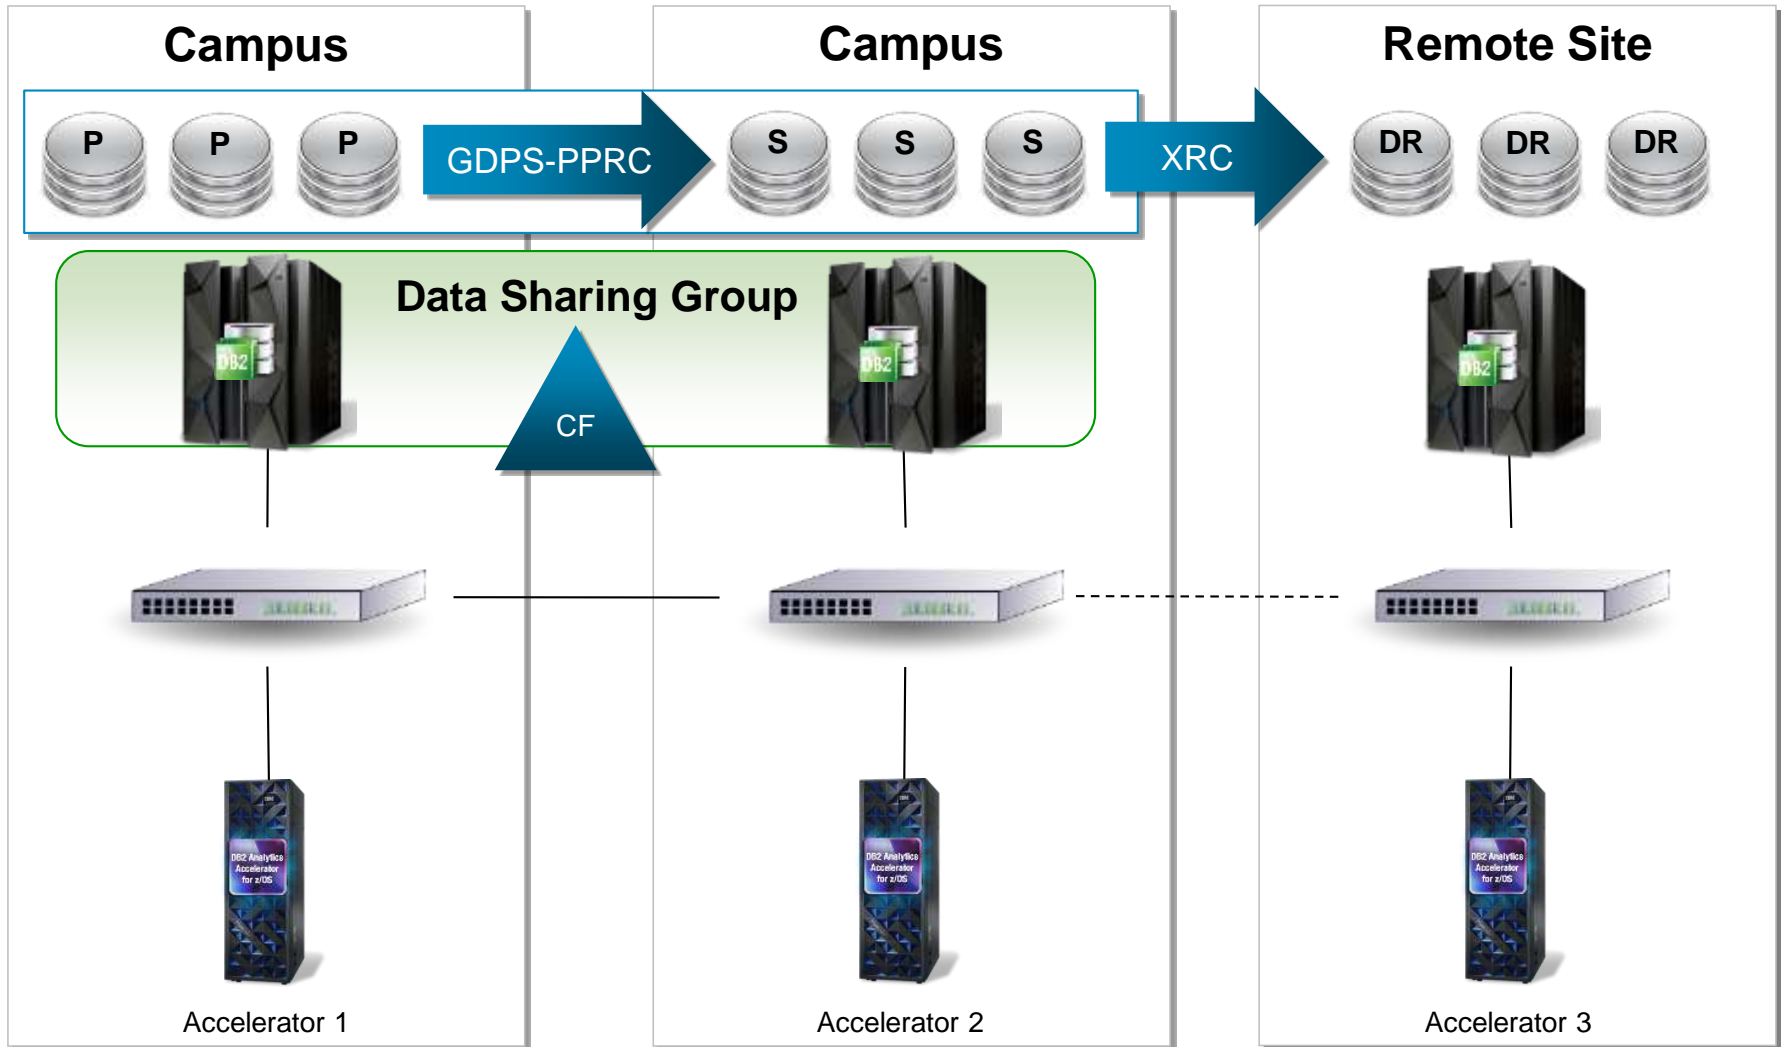
## Workload management concepts

Concept	What	Why	When
<b>Query Prioritization</b>	DB2 passes WLM importance level per query to the accelerator where it is mapped to a priority	Execute more important queries before less important queries that are sent from the same DB2 subsystem	An accelerator connected to a DB2 subsystem
<b>Workload Isolation</b>	User can configure resource allocation percentage on accelerator per connected DB2 subsystem	Ensure that specific DB2 subsystems have a guaranteed amount of available resources on a shared accelerator.	An accelerator connected to multiple DB2 subsystems
<b>Workload Balancing</b>	DB2 checks the utilization of each eligible accelerator before routing the query	Balance routing of queries equally across accelerators to ensure best utilization and performance;  High availability	Multiple accelerators connected to a DB2 subsystem

# High availability concepts – Workload balancing



# HA and DR with PPRC and XRC





# System monitoring



## ▪ Metrics available to monitor the following:

- Accelerator operational status
- Query execution (per query or aggregated)
  - Execution/CPU/wait times, queuing info,
  - # failed/success, # concurrently running, # rows/bytes returned
- Statistics
  - CPU utilization, disk space usage, replication progress and latency
- Accounting
  - Query elapsed/CPU/wait times per authid or per plan

## ▪ Monitoring metrics provided in:

- Statistics records from IFCID 2
- Accounting records from IFCID 3
- Query execution records from IFCIDs 316, 401

## System monitoring (contd.)

- Monitoring metrics can be viewed using:
  - DB2 Analytics Accelerator Studio

▼ Monitoring					
<b>Queries:</b>	Successful:	29,653	Failed:	6	
<b>Queue:</b>	Max. length:	0	Avg. wait time:	0 ms	Max. wait time: 28.6 s
<b>CPU cost:</b>	Query execution:	0:02:09	Data maintenance:	98.2 s	Replication: 0 ms

- DISPLAY ACCEL command
  - See example on a following slide
- DSN\_STATEMENT\_CACHE\_TABLE
  - Accumulated execution metrics per statement
- Tivoli OMEGAMON Performance Expert for z/OS
  - Statistics Report (see example on a following slide)
  - Accounting Report
  - Utility Trace Report
  - Real-Time Monitoring
- Other tools displaying metrics from specified IFCIDs

# System monitoring – DIS ACCEL DETAIL

DSNX810I @ DSNX8CMD DISPLAY ACCEL FOLLOWS -

DSNX830I @ DSNX8CDA 549

ACCELERATOR	MEMB	STATUS	REQUESTS	ACTV	QUED	MAXQ
-----	-----	-----	-----	-----	-----	-----
ZV4IDAA1	VA1A	STARTED	1662	0	0	0

LOCATION=ZV4IDAA2 HEALTHY

DETAIL STATISTICS

LEVEL = AQT04010

STATUS = ONLINE

FAILED REQUESTS = 0

AVERAGE QUEUE WAIT = 0 MS

MAXIMUM QUEUE WAIT = 0 MS

TOTAL NUMBER OF ACTIVE PROCESSORS = 4

AVERAGE CPU UTILIZATION ON COORDINATOR NODES = .01%

AVERAGE CPU UTILIZATION ON WORKER NODES = .00%

AVERAGE DISK IO UTILIZATION = .00%

NUMBER OF ACTIVE WORKER NODES = 1

TOTAL DISK STORAGE = 97713 MB

DISK STORAGE IN USE FOR THIS DB2 SYSTEM = 0 MB

DISK STORAGE IN USE FOR ALL DB2 SYSTEMS = 0 MB

TOTAL CPU FOR REQUESTS FOR THIS DB2 SYSTEM = 50560 MS

TOTAL CPU FOR DATA MAINTENANCE FOR THIS DB2 SYSTEM = 0 MS

TOTAL CPU FOR REPLICATION FOR THIS DB2 SYSTEM = 0 MS

DISPLAY ACCEL REPORT COMPLETE

DSN9022I @ DSNX8CMD '-DISPLAY ACCEL' NORMAL COMPLETION

# OMPE statistics report for IBM DB2 Analytics Accelerator V4

Distinguish accelerator data between subsystem related data and aggregated cross all subsystems

Q100	FOR SUBSYSTEM ONLY	QUANTITY	Q100	TOTAL ACCELERATOR	QUANTITY
QUERIES SUCCESSFULLY EXECUTED		1.00	QUERIES SUCCESSFULLY EXECUTED		1.00
QUERIES FAILED TO EXECUTE		1.00	QUERIES FAILED TO EXECUTE		1.00
CURRENTLY EXECUTING QUERIES		0.00	CURRENTLY EXECUTING QUERIES		0.23
MAXIMUM EXECUTING QUERIES		1.00	MAXIMUM EXECUTING QUERIES		1.00
CPU TIME EXECUTING QUERIES		1.290000	CPU TIME EXECUTING QUERIES		1.290000
CPU TIME LOAD/ARCHIVE/RESTORE		15:42.600000	CPU TIME LOAD/ARCHIVE/RESTORE		15:42.600000
CONNECTS TO ACCELERATOR		4.00	ACCELERATOR SERVER START	09/05/13 13:36:48.19	
REQUESTS SENT TO ACCELERATOR		6.00	ACCELERATOR STATUS CHANGE	09/09/13 11:47:05.22	
TIMED OUT		0.00			
FAILED		0.00	DISK STORAGE AVAILABLE (MB)		48000959.97
BYTES SENT TO ACCELERATOR		7618.00	IN USE FOR ACCEL DB - ALL DB2 (MB)		1932487.60
BYTES RECEIVED FROM ACCELERATOR		2707.00	IN USE FOR ACCEL DB - THIS DB2 (MB)		64322.60
MESSAGES SENT TO ACCELERATOR		33.00			
MESSAGES RECEIVED FROM ACCEL		33.00	MAXIMUM QUEUE LENGTH		0.00
BLOCKS SENT TO ACCELERATOR		0.00	CURRENT QUEUE LENGTH		0.00
BLOCKS RECEIVED FROM ACCELERATOR		2.00	AVG QUEUE WAIT ELAPSED TIME		0.021328
ROWS SENT TO ACCELERATOR		0.00	MAX QUEUE WAIT ELAPSED TIME		0.945941
ROWS RECEIVED FROM ACCELERATOR		53.00			
TCP/IP SERVICES ELAPSED TIME		28:18.061328	WORKER NODES		7.00
ELAPSED TIME IN ACCELERATOR		7.791182	WORKER NODES DISK UTILIZATION (%)		2.40
WAIT TIME IN ACCELERATOR		0.099476	WORKER NODES AVG CPU UTILIZATION (%)		23.14
			COORDINATOR CPU UTILIZATION (%)		8.71
			PROCESSORS		224.00
			DATA SLICES		240.00
CPU TIME FOR REPLICATION		N/P	CPU TIME FOR REPLICATION		N/P
LOG RECORDS READ		N/P	LOG RECORDS READ		N/P
LOG RECORDS FOR ACCEL TABLES		N/P	LOG RECORDS FOR ACCEL TABLES		N/P
LOG RECORD BYTES PROCESSED		N/P	LOG RECORD BYTES PROCESSED		N/P
INSERT ROWS FOR ACCEL TABLES		N/P	INSERT ROWS FOR ACCEL TABLES		N/P
UPDATE ROWS FOR ACCEL TABLES		N/P	UPDATE ROWS FOR ACCEL TABLES		N/P
DELETE ROWS FOR ACCEL TABLES		N/P	DELETE ROWS FOR ACCEL TABLES		N/P
REPLICATION LATENCY IN SECONDS		N/P			
REPLICATION STATUS CHANGE		N/P			

Capacity planning

Replication progress  
(ChangeDataCapture)

# Accelerator error messages

- Accelerator Error Messages are reported as DSNX881I messages in the z Systems SYSLOG
- Call Home functionality available to react on them

## DSNX881I

SYSLOG

Status Change (ID 1, Info)  
 Hardware Service Requested (ID 2, Error)  
 Hardware Restarted (ID 3, Info)  
 Disk Space Threshold (ID 4, Warning)  
 Runaway Query (ID 5, Info)  
 System Stuck in State (ID 6, Error)  
 Predictive Failure (ID 7, Warning)  
 ECC Errors (ID 8, Warning)  
 Regeneration Error (ID 9, Error)  
 Disk Error (ID 10, Error)  
 Hardware Temperature (ID 11, Warning)  
 System Temperature (ID 12, Error)  
 SPU Coredumps (ID 13, Warning)  
 Voltage Fault (ID 14, Error)  
 Network Interface State Change (ID 15, Information)  
 S-Blade CPU Core Change (ID 16, Information)  
 IMM Event (ID 17, Warning/Error )  
 Hardware Path Down (ID 18, Error)  
 Hardware needs attention (ID 19, Warning)

### Example:

```
0010 DSNX881I -D916 19 W 314751 (10-Dec-12, 18:24:17 EST) 596
0010 Hardware needs attention. Host = dwatf0312ma.boeblingen.de.ibm.com
0010 Hardware type = SPA Hardware ID = 1003 Location = 1st Rack, 2nd SPA
0010 -msg
```

Call Home

Accelerator

Error  
Event

z Systems

E-Mail

IBM  
PMR Gateway

Messages and meaning: <http://www-01.ibm.com/support/docview.wss?uid=swg27037905>)

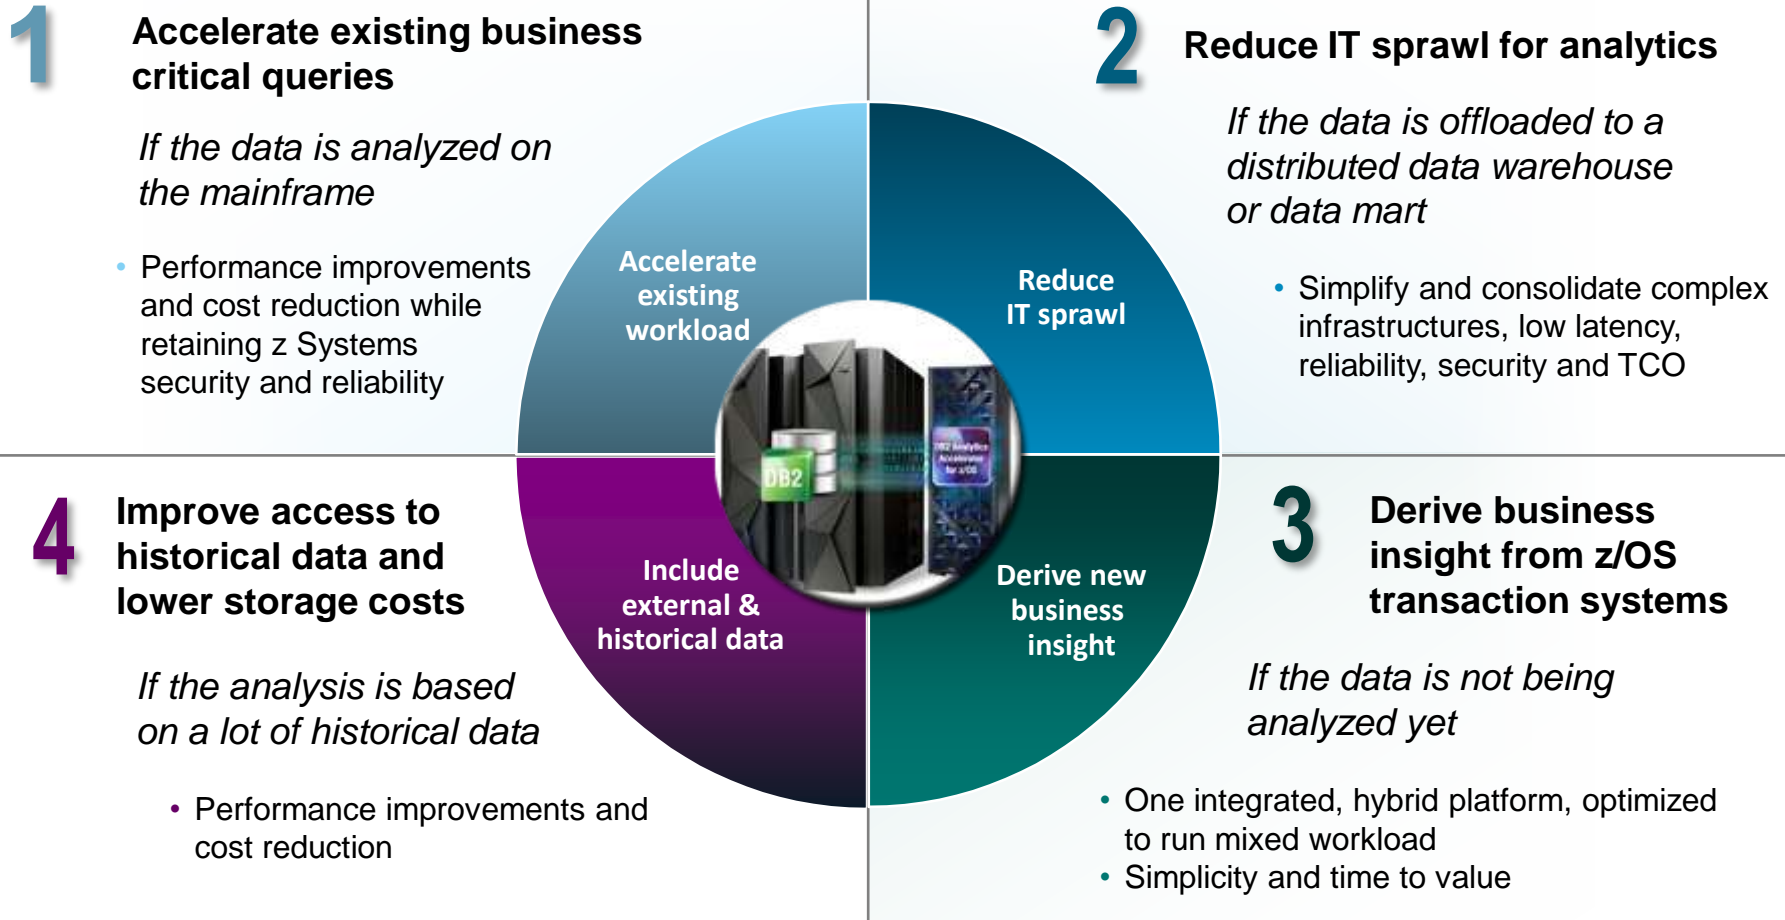
# Agenda

- Value Proposition
- Technical Overview
- New Version 5.1 Functionality
- Managing and Monitoring
- **Use Cases and Customer Experiences**
- Strategy and Evolution



# DB2 Analytics Accelerator – Four usage scenarios

*How organizations are leveraging the Accelerator today*



# New use cases with Accelerator-only tables

## 1 Accelerate existing business critical queries

*Adding application support for temporary objects*

- Multi-step Reporting
- QMF
- Campaign

## 2 Reduce IT sprawl for analytics

*In-database transformation*

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

## 4 Improve access to historical data and lower storage costs

*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

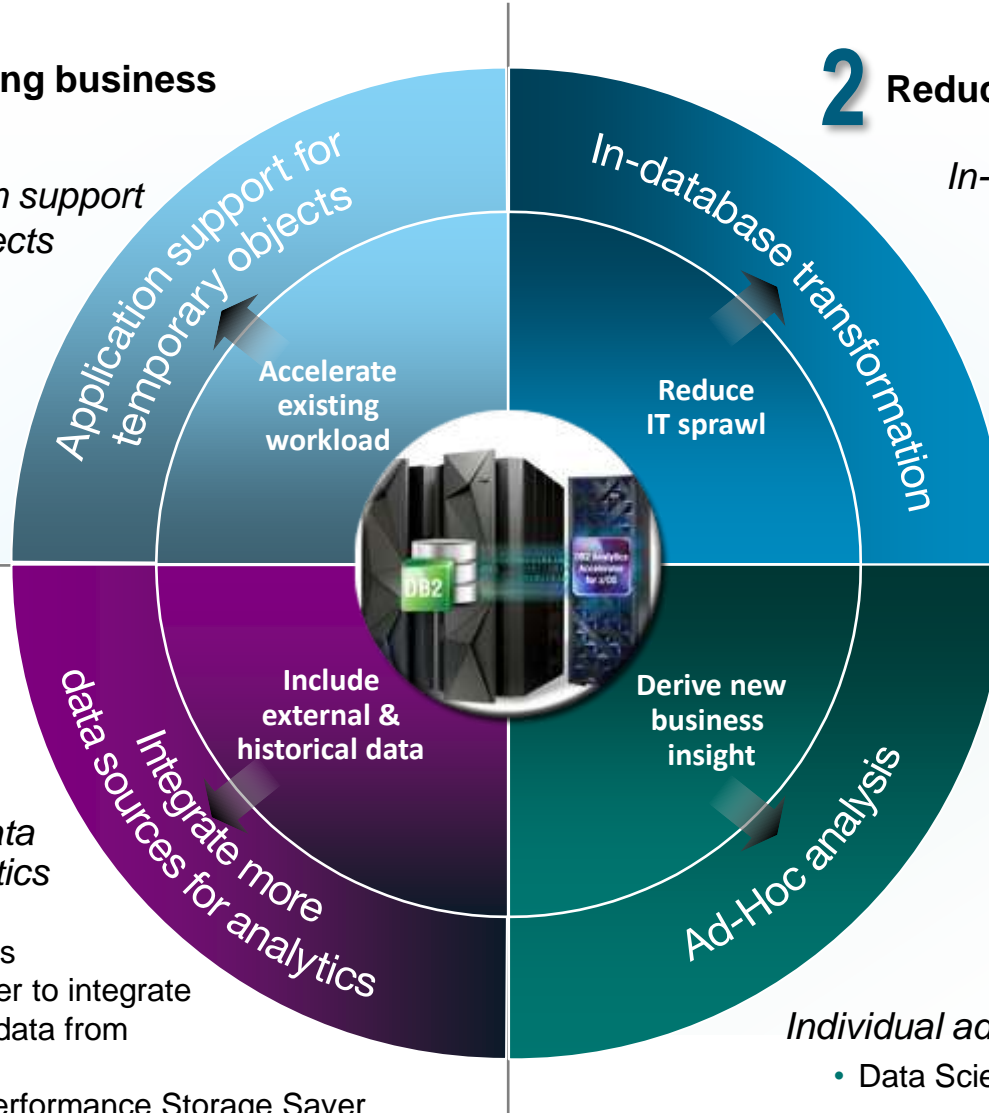
## 3 Derive business insight from z/OS transaction systems

*In-database analytics*

- Accelerate predictive analytics applications
- SPSS/INZA data mining and in-database modeling

*Individual ad-hoc analysis*

- Data Scientist Work Area



## Halkbank Case Study

Halkbank extends banking services on social and mobile channels securely and cost-effectively



<http://www-03.ibm.com/software/businesscasestudies/us/en/corp?synkey=E944346V10306Q58>

“ We wanted to improve the performance for complex analytic queries – both in order to support rapid and accurate decision making, and in order to meet deadlines for regulatory reporting, where we face penalties of tens of thousands of dollars for late reporting. We expect to be able to move 70 percent of our queries to DB2 Analytics Accelerator, accelerating queries, reducing our DB2 licensing costs on the mainframe, and freeing up processing resources for other workloads.

-- Cenk Niksarlı, Head of IT Infrastructure at Halkbank

## Swiss Mobiliar Case Study and Video

Swiss Mobiliar  
accelerates 50%  
of queries by a  
factor of 100



Case Study: <http://www-03.ibm.com/software/businesscasestudies/us/en/swzseries?synkey=T281533L28114L28>

Video: <http://www-03.ibm.com/software/businesscasestudies/en/us/swzseries?docid=ASHF-9GFDH9>

“

*DB2 Analytics Accelerator makes it possible for us to execute 90% of our queries 25x faster, and 50% of them 100x faster ...*

*Queries that used to take five hours to complete are now processed in just 20 seconds in the optimized mainframe environment—and we can run them any time, day or night, with no interruption to our production systems on the mainframe.*

-- Thomas Baumann, IT Performance Architect at Swiss Mobiliar

## Banca Carige Case Study and Video

Providing superior customer service on a secure platform



Case Study: [http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=AB&infotype=PM&appname=SWGE\\_ZS\\_ZS\\_USEN&htmlfid=ZSC03256USEN&attachment=ZSC03256USEN.PDF#loaded](http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=AB&infotype=PM&appname=SWGE_ZS_ZS_USEN&htmlfid=ZSC03256USEN&attachment=ZSC03256USEN.PDF#loaded)

Video: <https://www.youtube.com/watch?v=N7arHCI-CKk&list=PLIbvgyHNmNr1imFVslsklpUuhBHkmNHR2&index=1>

“ *DB2 Analytics Accelerator is key in our vision to provide the right level of data analysis to all users in the bank and possibly outside the bank.*

*By using DB2 Analytics Accelerator, we expect to reach ... all criteria we had in mind and create greater benefits for our users because our data are not moved...*

-- Daniel Cericola, IT & Governance at Banca Carige



## Swiss Re Case Study and Video

Speeding high-level decision making that boosts the bottom line



Case Study: [http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=AB&infotype=PM&appname=SWGE\\_ZS\\_ZS\\_USEN&htmlfid=ZSC03128USEN&attachment=ZSC03128USEN.PDF](http://www-01.ibm.com/common/ssi/cgi-bin/ssialias?subtype=AB&infotype=PM&appname=SWGE_ZS_ZS_USEN&htmlfid=ZSC03128USEN&attachment=ZSC03128USEN.PDF)

Video: [https://www.youtube.com/watch?v=xkcp\\_pJxT5E&feature=youtu.be](https://www.youtube.com/watch?v=xkcp_pJxT5E&feature=youtu.be)

“Our users are getting their reports as much as 70 percent faster—reports that took 10 hours are now available the same day, so user satisfaction has increased dramatically,” says Estermann. “Because those reports contain key analytics that guide pricing and decision making across our lines of business, the solution has the potential to sharpen the company’s competitive edge moving.”

-- Reto Estermann, Director, Information Technology at Swiss Re



# Agenda

- Value Proposition
- Technical Overview
- New Version 5.1 Functionality
- Managing and Monitoring
- Use Cases and Customer Experiences
- **Strategy and Evolution**

# Evolution and trends

Hardware evolution



More query acceleration



Capabilities for more use cases



Improved transparency and management



## Examples

- Faster hardware
- Self-encrypting disks

- Static SQL support
- Closing gaps in unsupported SQL

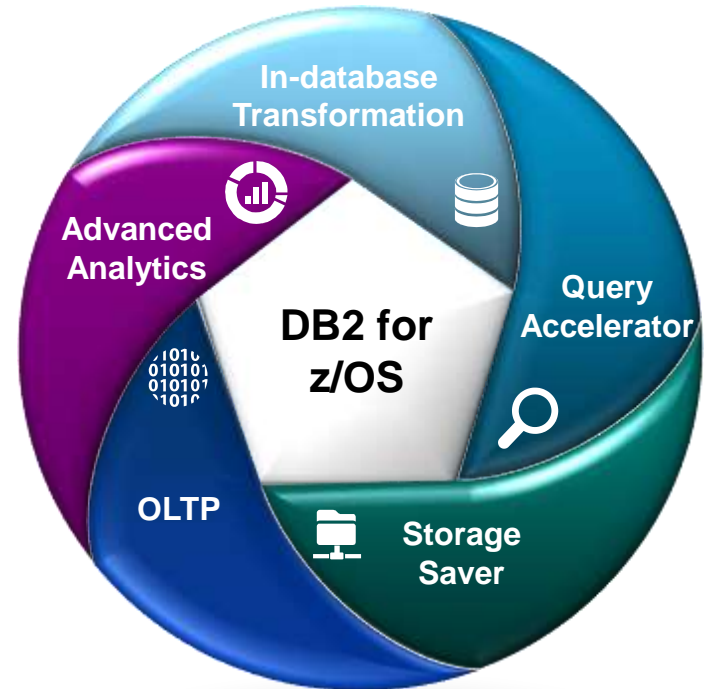
- Accelerator-only tables
- Improved data maintenance

- Call home
- Management of failures

# Strategy

Enable DB2 transition into a truly universal DBMS that provides best characteristics for both OLTP and analytical workloads.

- Complement DB2's industry leading transactional processing capabilities
- Provide specialized access path for data intensive queries
- Enable real and near-real time analytics processing
- Execute transparency to the applications
- Operate as an integral part of DB2 and z Systems
- Reusing industry leading PDA's query and analytics capabilities and take advantage of future enhancements
- Extend query acceleration to new, innovative usage cases, such as:
  - in-database transformations
  - advanced analytical capabilities
  - multi-temperature and storage saving solutions



Ultimately allow consolidation and unification of transactional and analytical data stores

# New developerWorks Community for IBM DB2 Analytics Accelerator

Join  
[ibm.biz/db2analyticsaccelerator](http://ibm.biz/db2analyticsaccelerator)



**IBM Bluemix** Develop in the cloud at the click of a button! Start building for free

IBM Sign in | Register

**IBM developerWorks** Technical topics Evaluation software Community Events


My home Forums Blogs **Communities** Profiles Podcasts Wikis Activities

IBM Champion program

**Communities** This Community Search

**IBM DB2 Analytics Accelerator Community**

▼ IBM DB2 Analytics Accelerator Community



Overview

▼ Community Description

**Welcome to the IBM DB2 Analytics Accelerator Community!**

IBM DB2 Analytics Accelerator is a workload-optimized appliance that integrates zEnterprise and PureData technologies to accelerate relevant data-intensive and complex DB2 for z/OS queries. Together, the DB2 Analytics Accelerator, DB2 for z/OS and zEnterprise form an integrated hybrid environment that can run transaction processing, complex analytical and reporting workloads concurrently and efficiently.

**Important Bookmarks**

- DB2 Analytics Accelerator Page
- DB2 Analytics Accelerator Page
- IBM Knowledge Center - IBM DB2 Analytics Accelerator for z/OS 4.1.0

**Members**

**Blog**

“ **New technote "How to assess queries for eligibility" available** 0

UteBaumbach | May 7 | Tags: queryacceleration queryassessment | 96 Visits

“ **Incremental Update - ONUTILITYACTION parameter** 0

Patric Becker | May 5 | Tags: cdc incremental update | 1 Comment | 94 Visits

[View All](#)

**Forums**

“ **Friendly arithmetic in DB2 for z/OS - SQLCODE +802 and the accelerator**

Latest post by UteBaumbach | Monday 5:57 PM | No replies

“ **How does the new 10 bytes RBA/LRSN affect tables archived via HPSS ?**

Latest post by UteBaumbach | Monday 5:55 PM | No replies

“ **CP improvements for loading data into IBM DB2 Analytics Accelerator**

Latest post by UteBaumbach | Monday 5:54 PM | No replies

**IBM**®

## Legal Disclaimer

- © IBM Corporation 2015. All Rights Reserved.
- The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.
- References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.