

A decorative graphic in the top left corner consists of several overlapping circles of various sizes. Each circle is divided into multiple segments of different colors, including shades of orange, yellow, purple, blue, and brown, creating a vibrant, abstract pattern.

# Getting the most from the DB2 Analytics Accelerator


Presenter

Title



# Agenda

- Trends in Analytics
- Business Critical Analytics
- DB2 Analytics Accelerator Overview
- Customer Experiences
- Best Practices for Assessing Workload and Requirements
- Delivering Analytic Applications on zEnterprise



# Trends in Analytics




# More data, more sources, more types . . .


Traditional structured data



Semi-structured data



Streaming



***ALL Data***

Information infrastructure



***ALL Platforms***



*Delivering insights to business users across the enterprise*

***Bring the Analytics to the Data***

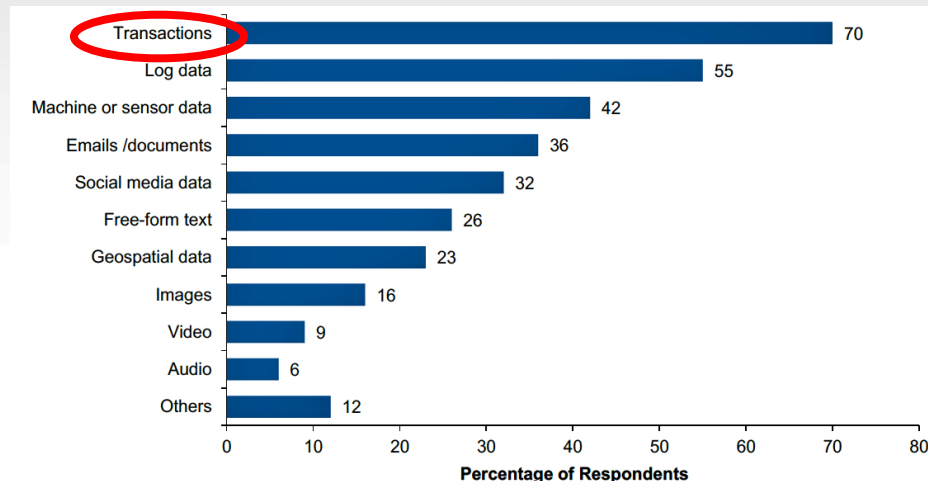
# Current analytics projects are heavily drawing on mainframe-based transactional data

What types of data/records are you planning to analyze using big data technologies?



Source: 2012 IBM Global Big Data Online Survey

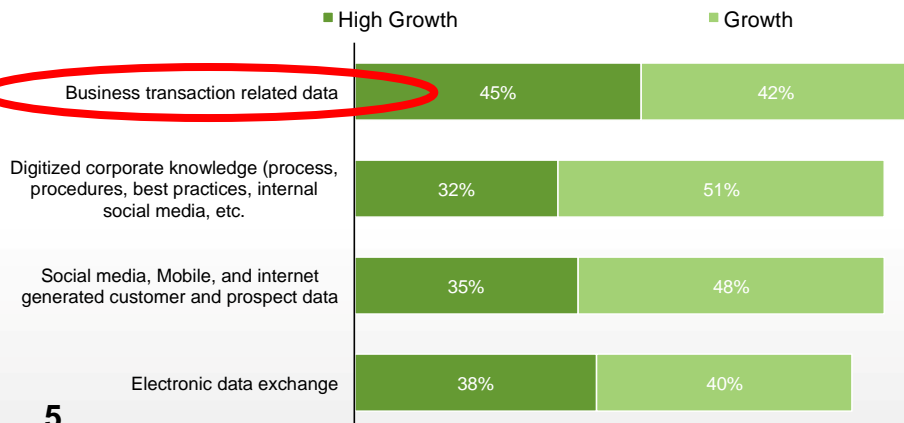
Types of data analyzed



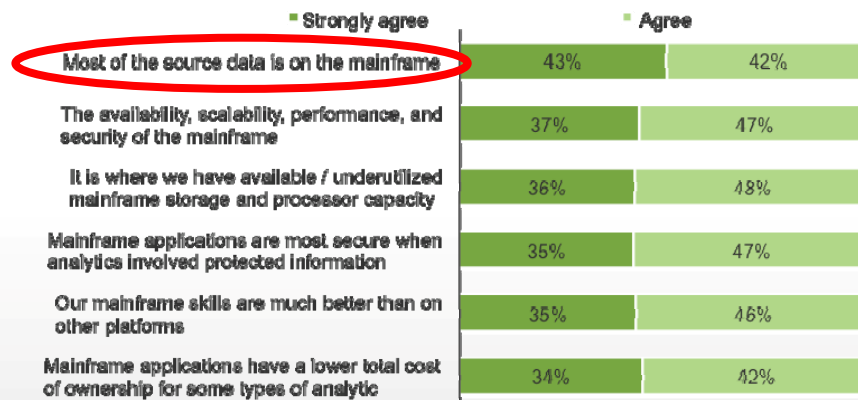
N =465 (multiple responses allowed)

Gartner Survey Analysis: Big Data Adoption in 2013 Shows Substance Behind the Hype  
Published: 12 September 2013

Please answer each statement as it best describes the growth of different types of data in your firm



How strongly do you agree with the following statements about why your firm runs analytics on the mainframe?



Source: A commissioned study conducted by Forrester Consulting on behalf of IBM, October, 2012



# Business Critical Analytics



# What is Business Critical Analytics?

- Integrated transactional and analytic applications critical to the optimal running of a business
- Analytics at the point of customer contact
- Enables the analytics-driven business for increased profitability and a competitive edge




**Point of Sale Marketing**



**Prevent Fraud**



**Operational Reporting**

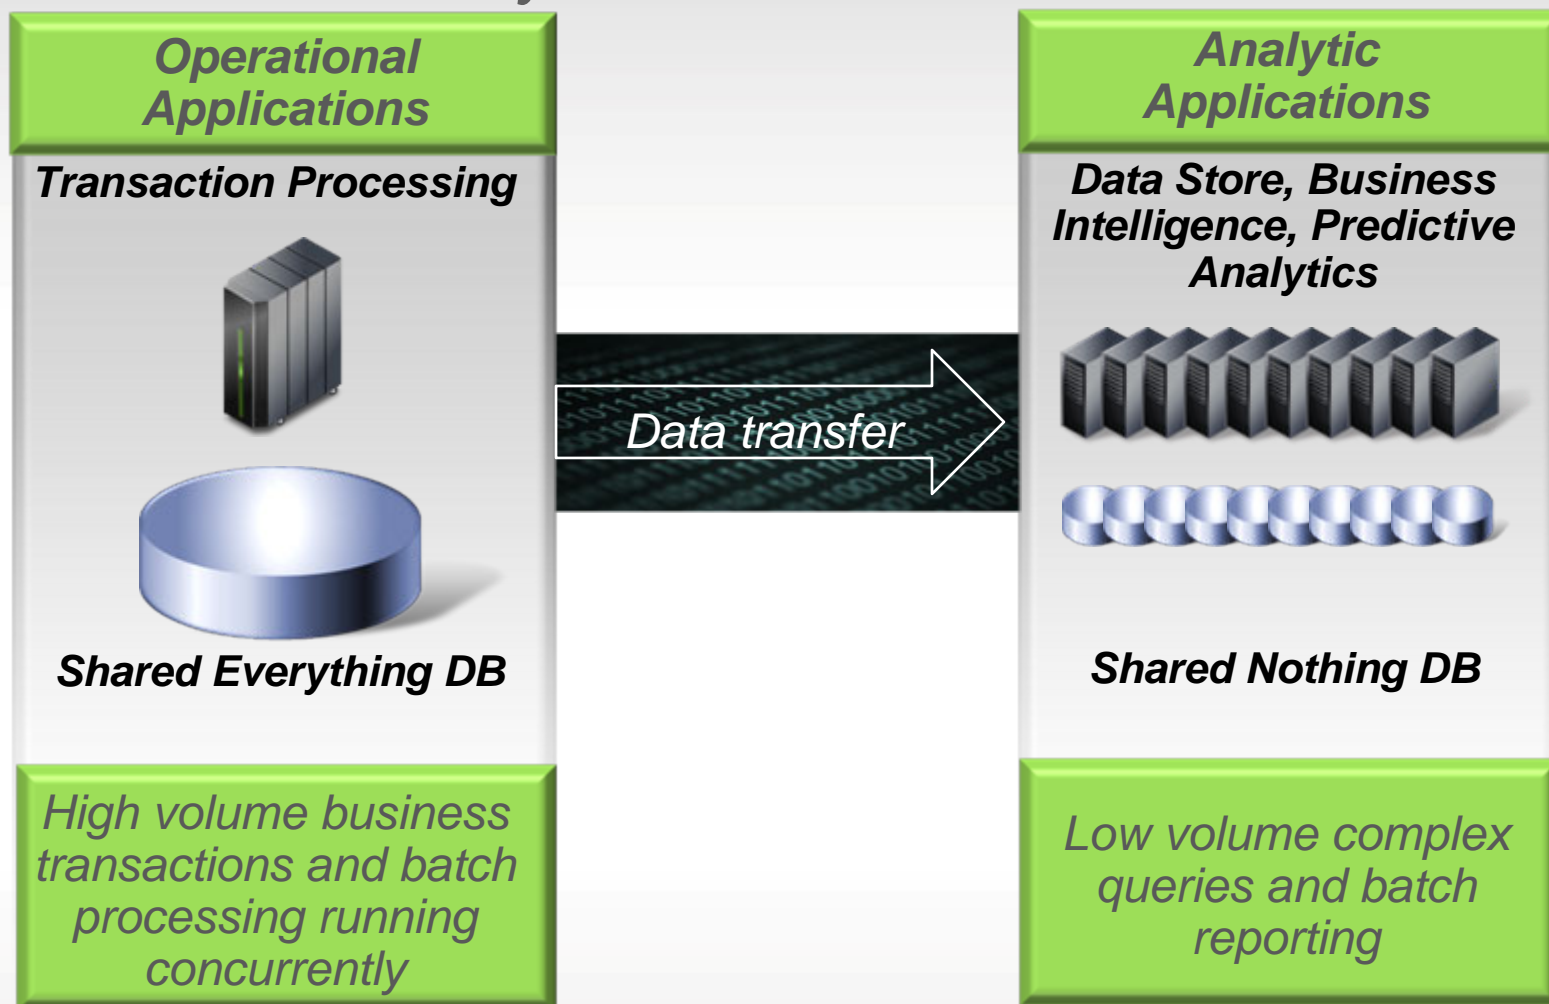


**Reduce Customer Churn**

***These applications require industry leading reliability and availability***

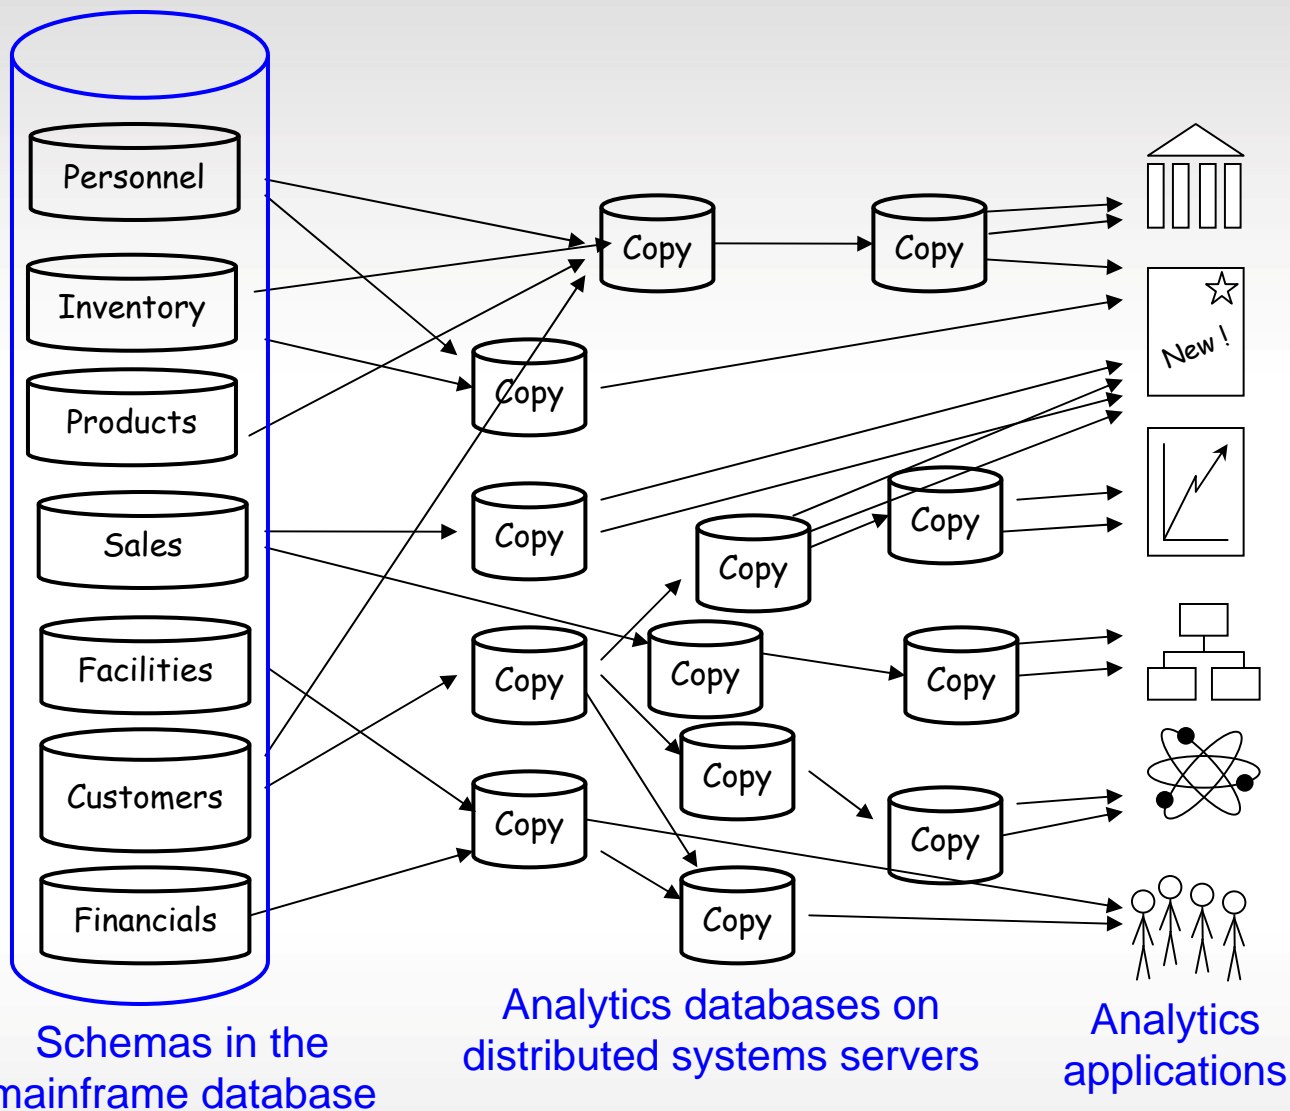
# The old conventional wisdom

*Move the data to the analytics*





... like this:



## Problems:

- Data latency: time between transaction and insight
- Expensive, resource-intensive data replication processes
- Greater risk of data security breaches
- Data governance issues: copies of data can become inconsistent – do users trust the data?
- Data currency challenges: copies of data can become out-of-date – users demand timely data
- Proliferation of data silos impedes integration, reduces value derived from data assets



# New Wisdom: The IBM Hybrid Approach

*Bring analytics to the data – on zEnterprise*

## *Combined Workloads*

*Transactional Processing, Traditional Analytics & Business Critical Analytics*



**Hybrid DB**

***Reduced Latency. Greater Security.  
Improved Data Governance. Reduced Complexity.***

*High volume business transactions and batch reporting running concurrently with complex queries*

***Delivering business critical analytics***



# IBM DB2 Analytics Accelerator and IBM zEnterprise

*Minimize latency. Improve performance. Drive innovation.*

## **Bring analytics to the data**

- Reduced latency
- Reduced complexity
- Reduced cost

## **Deliver business critical analytics**

- Timely, accurate, secure data
- Availability, scalability, performance
- Rapid deployment & expansion

## **Evolve with the business**

- Start with your top analytic requirement(s)
- Grow without changing your existing IT environment



A decorative graphic in the top left corner consists of several overlapping circles of various colors (yellow, orange, red, purple, blue) that are divided into segments, resembling a stylized sunburst or a cluster of data points.

# DB2 Analytics Accelerator Overview



# IBM DB2 Analytics Accelerator

*Do things you could never do before!*

## What is it?

The IBM DB2 Analytics Accelerator is 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 and exceptional business value

## What does it do?

- Accelerates complex queries, up to 2000x faster
- Lowers the cost of storing, managing and processing historical data
- Minimizes latency
- Reduces zEnterprise capacity requirements
- Improves security and governance
- Reduces operational costs and risk
- Complements existing investments





# IBM zEnterprise and DB2 Analytics Accelerator

*The Best Solution for Business Critical Analytics*

## ***How is it different?***

- **Integration:** DB2 applications can seamlessly combine OLTP and analytics on the same content
- **Performance:** Exceptional performance for both OLTP and analytic operations on the same platform and content
- **Transparency:** Accelerator is completely transparent to DB2 applications
- **Self-managed workloads:** Queries are automatically executed on the most efficient physical location
- **Rapid time to deployment:** 1-2 days from delivery to information
- **Simple administration:** Appliance hands-free operations, eliminating most database tuning tasks
- **Cost efficient:** Reduced cost through simplified administration and optimal use of enterprise computing assets

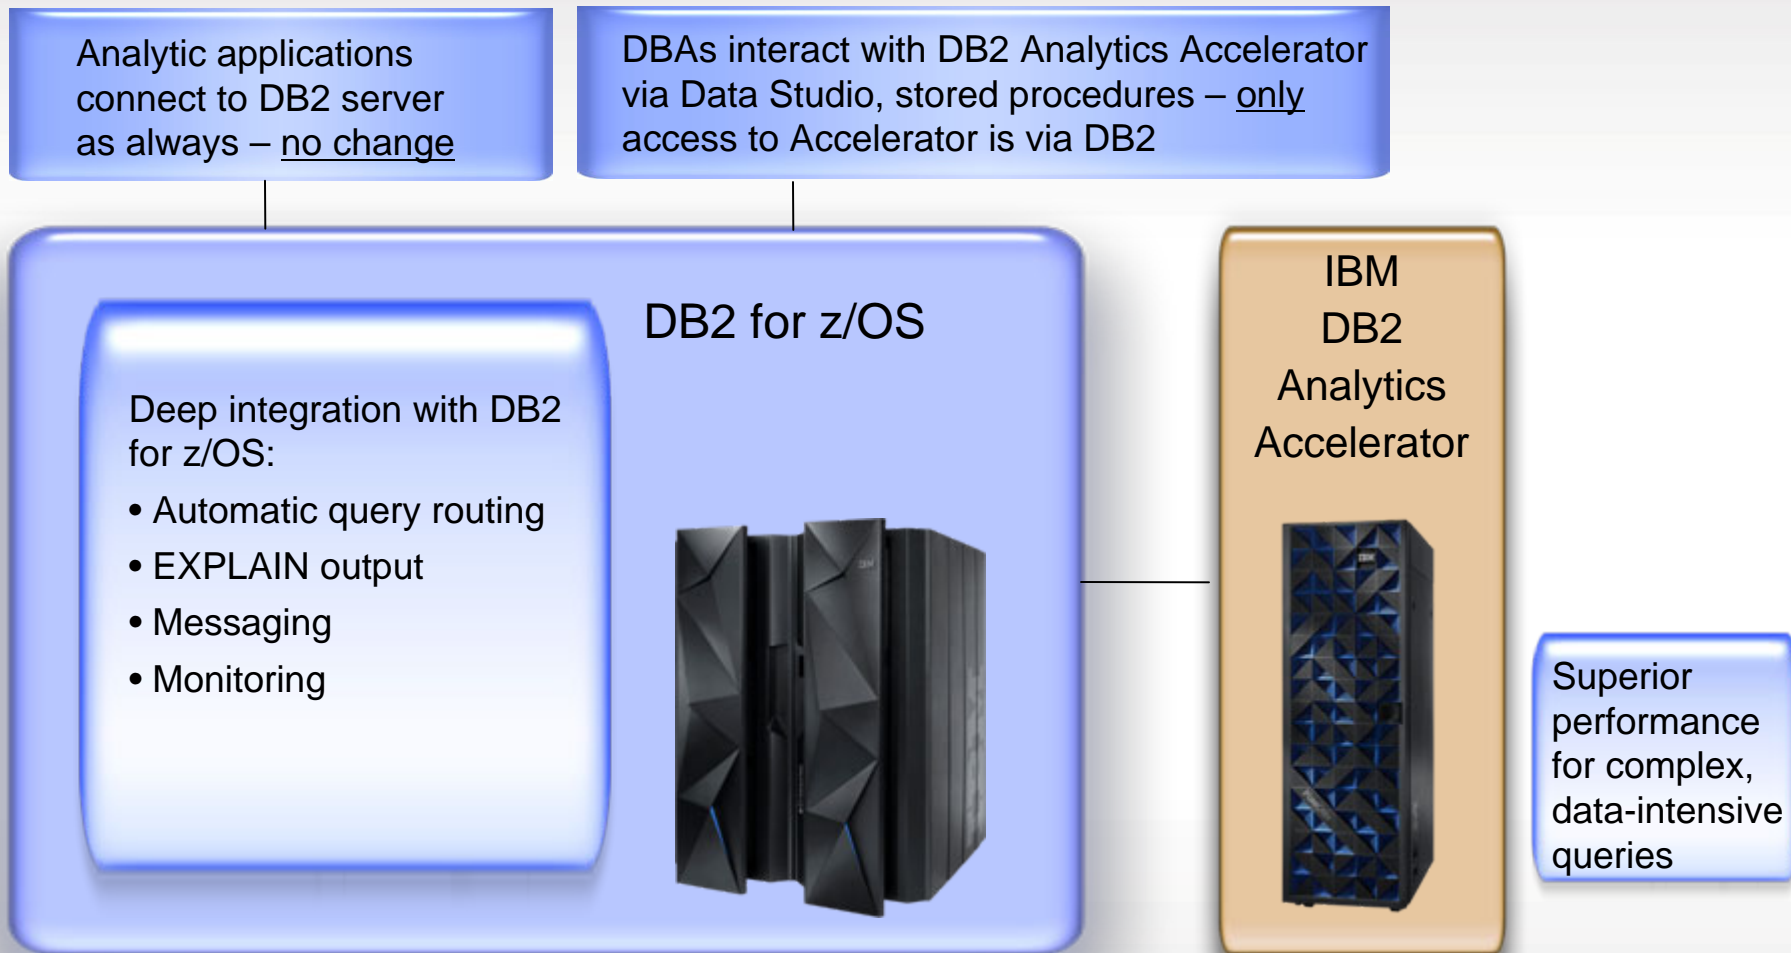




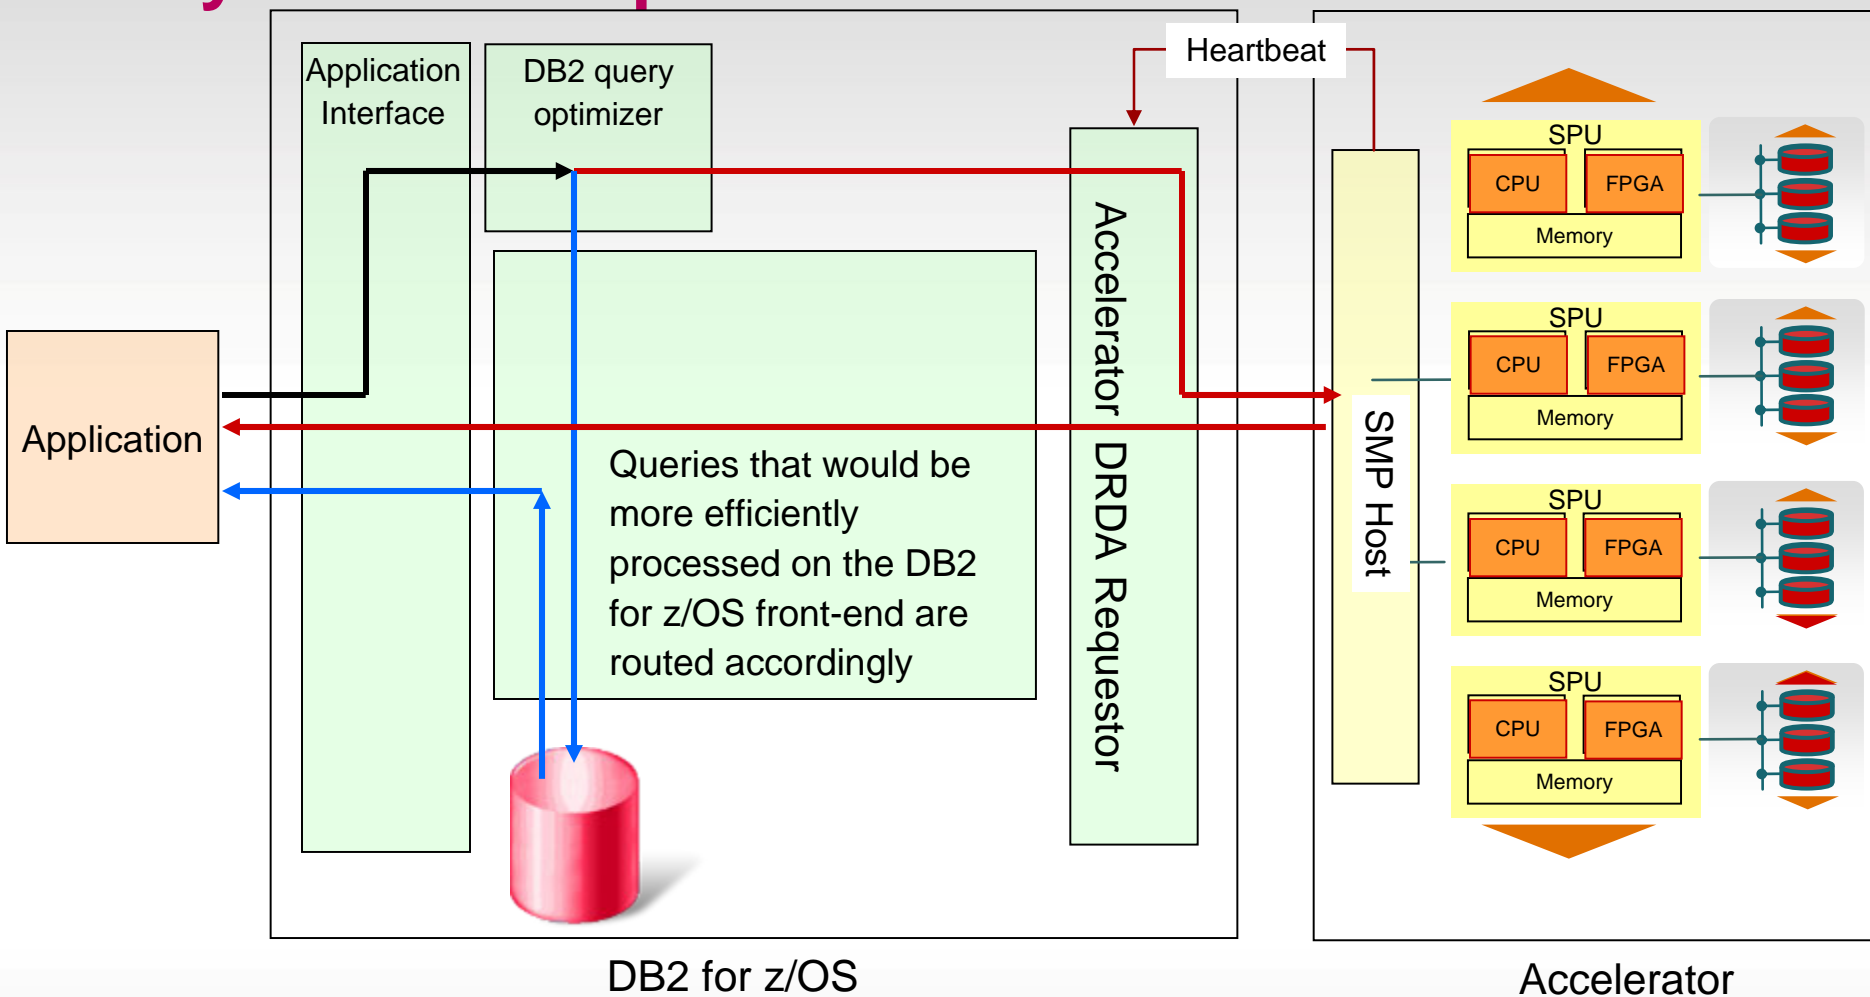


# DB2 Analytics Accelerator - not “just an appliance”

*An extension of the DB2 for z/OS system*



# Query execution process flow

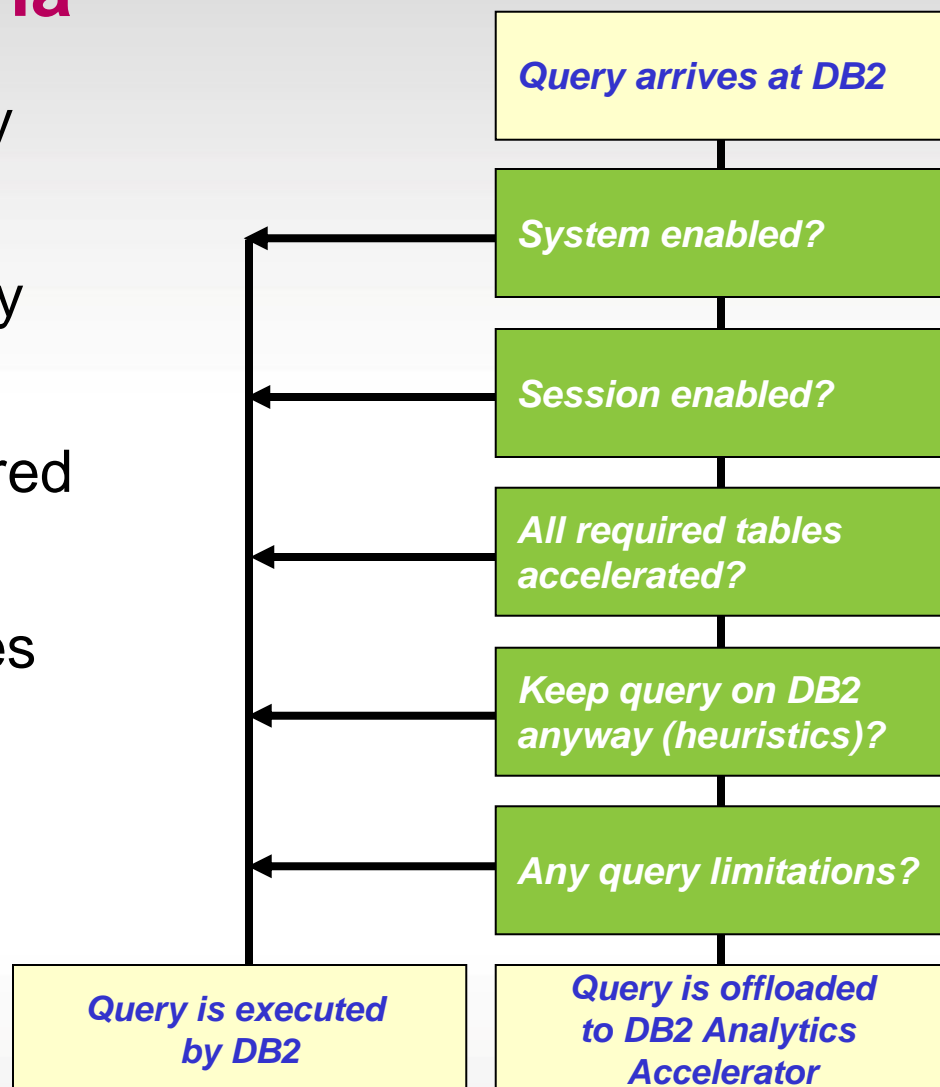


- Queries executed on the DB2 for z/OS front-end
- Queries executed on the Analytics Accelerator



# Automatic Routing Criteria

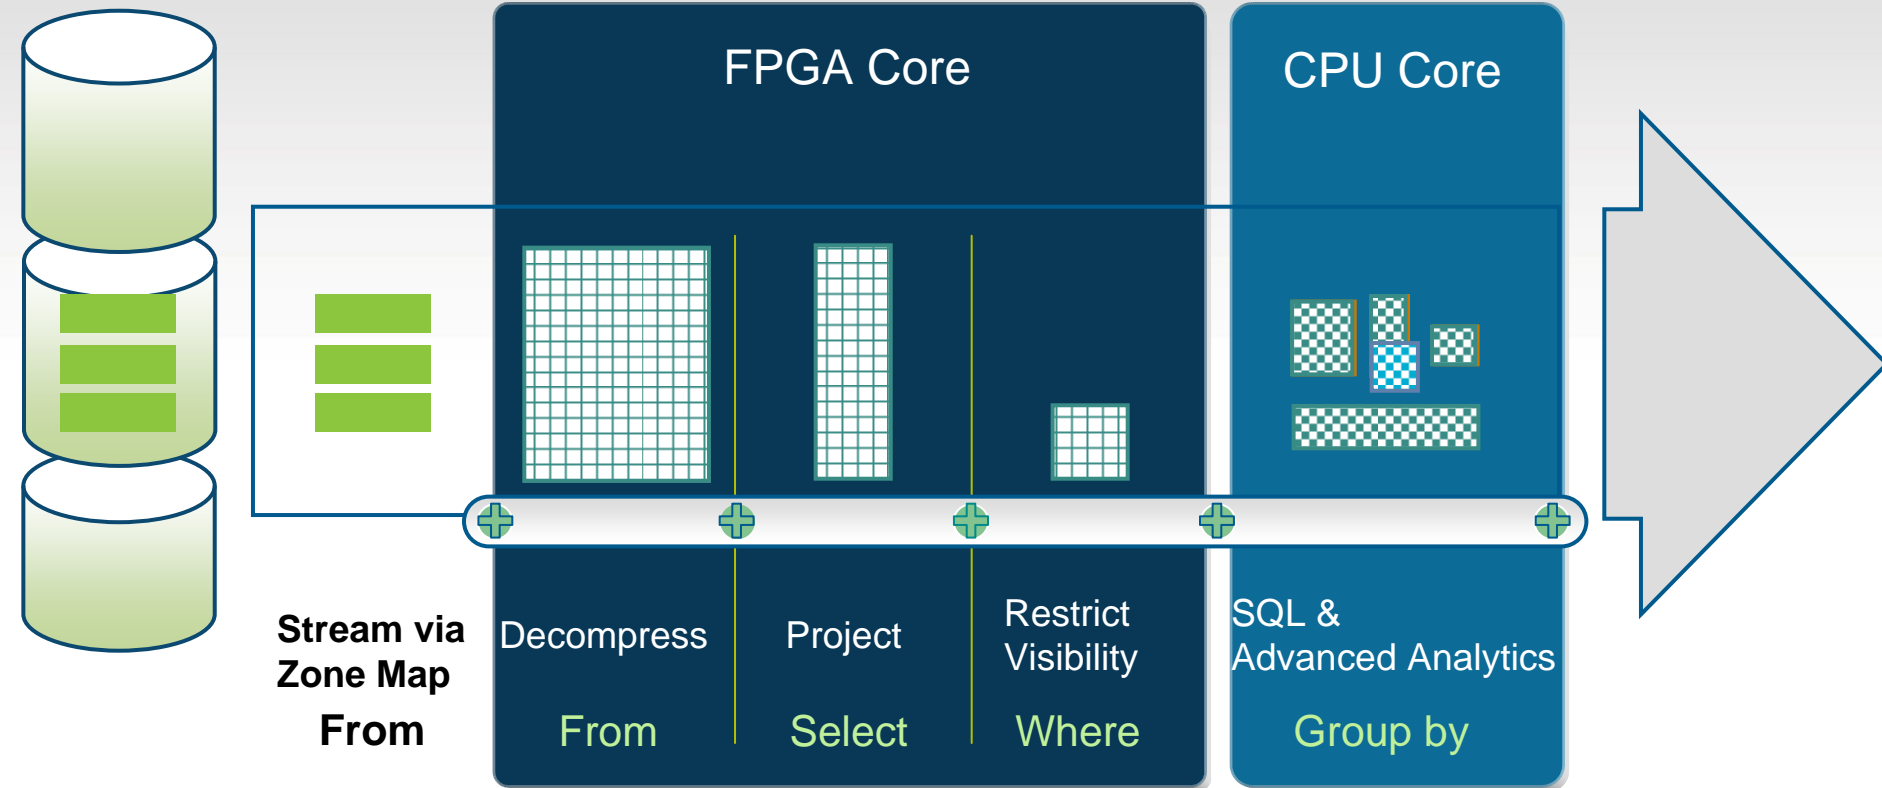
- DB2 Optimizer decides if query should be sent to accelerator
- Whole query, not parts of query are accelerated
- Only read queries are considered for acceleration
- Both static and dynamic queries can be accelerated\*



\* Routing for static queries determined at bind time

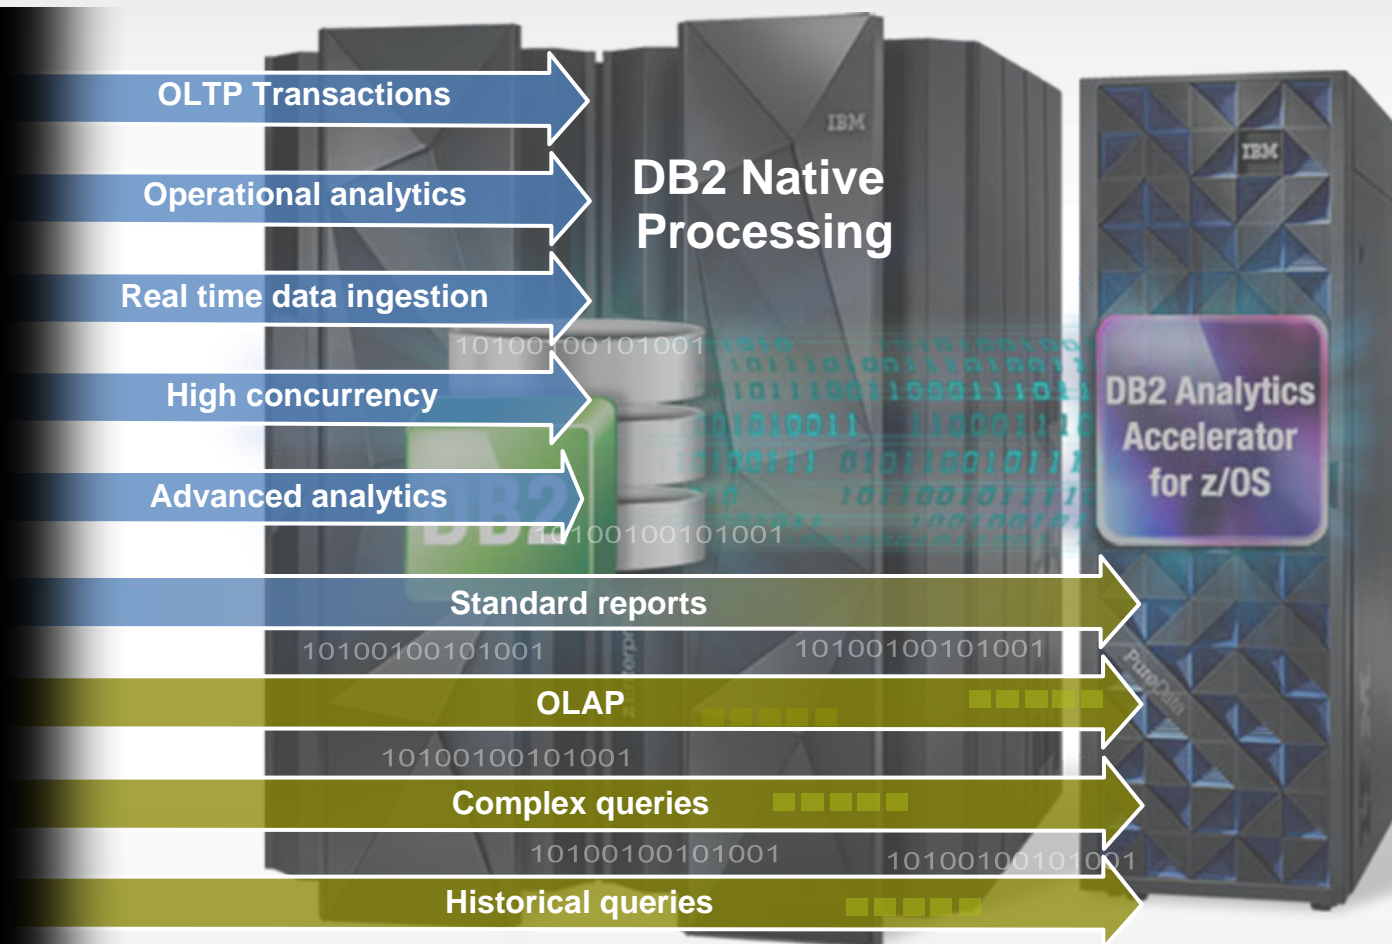


# The Key to the Speed



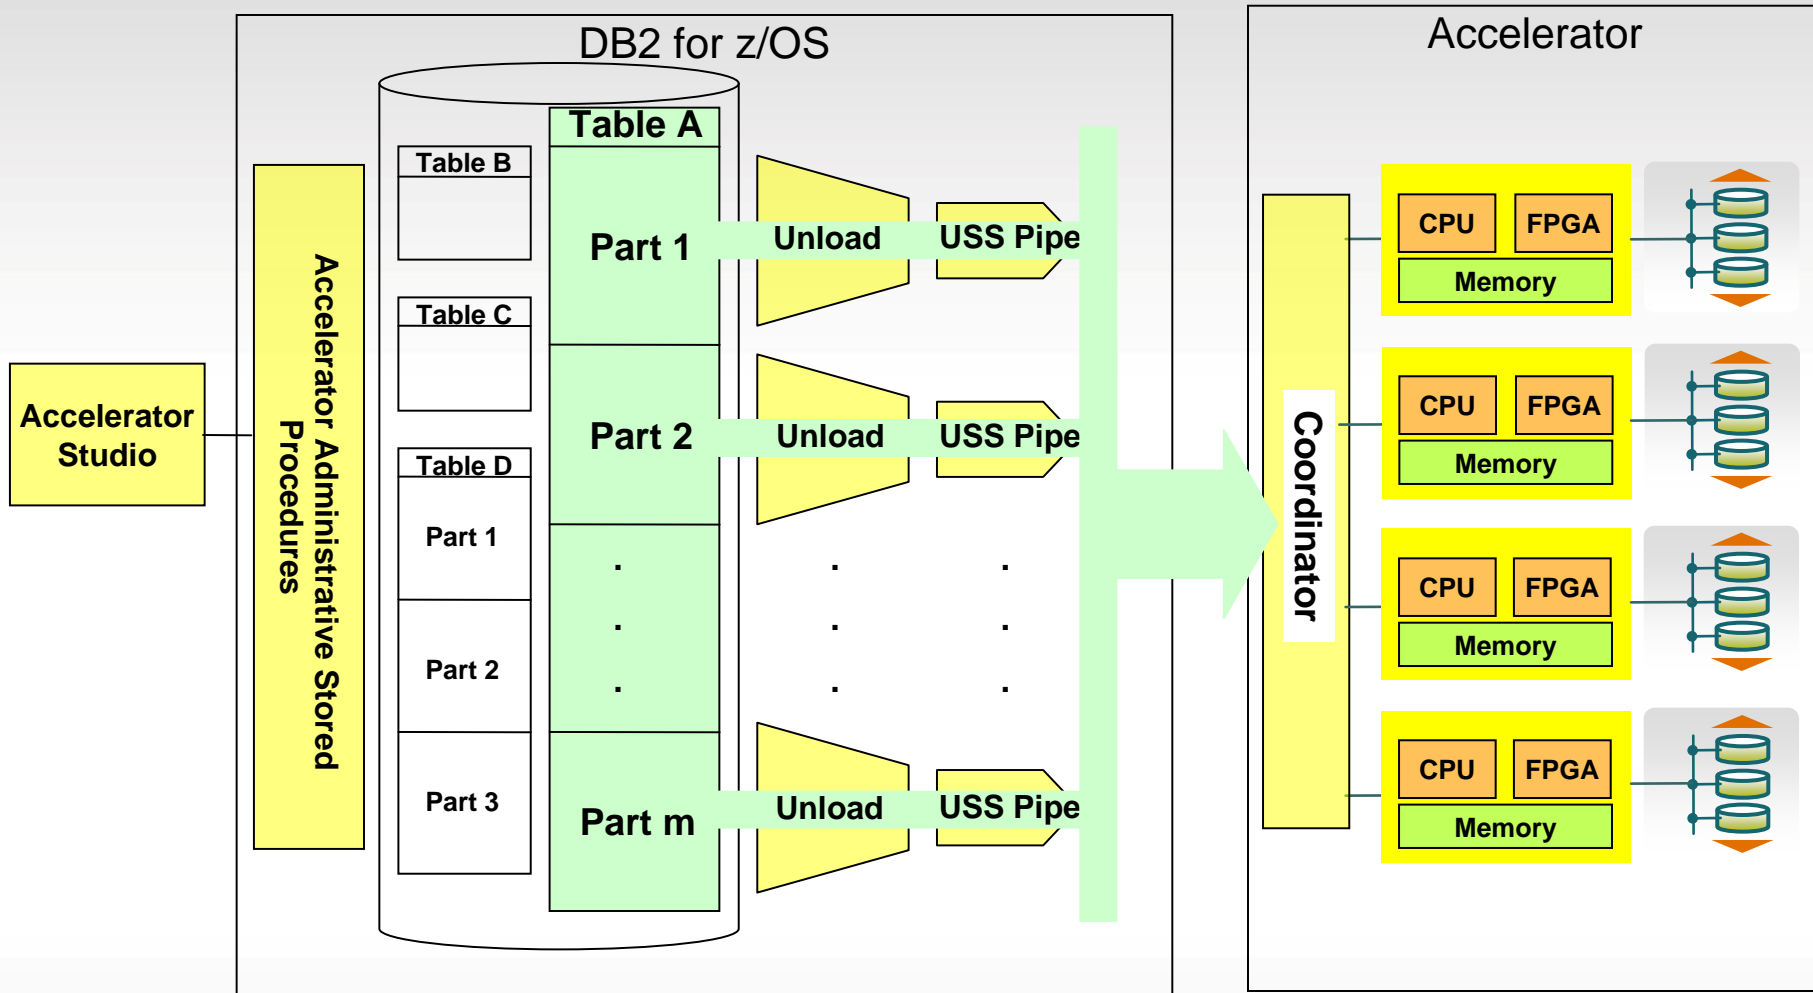
Select State, Age, Gender, count(\*) From Multi Billion Row Cost Center Table Where Date Birth < 1960  
 and State in ('FL', 'NC', 'CA', 'SC') Group by State, Age, Gender, State, Age, Gender, State, Age, Gender

# Workload-Optimized Query Execution



- Optimization for most efficient execution platform
- Special Register  
QUERY ACCELERATION
  - NONE
  - ENABLE
  - ENABLE WITH FAILBACK
  - ELIGIBLE
  - ALL
- Set via zPARM, Set Statement, Connection Properties / Driver, SQL Pre-Pend

# Accelerator Data Load



- Rate can vary, depending on CPU resources, table partitioning, columns...
- Update on table partition level, concurrent queries allowed during load
- Version 2.1 & Version 3.1 unload in DB2 internal format, single translation by Accelerator



# Synchronization Options with DB2 Analytics Accelerator

Synchronization options	Use cases, characteristics and requirements
<p><b><u>Full table refresh</u></b></p> <p>The entire content of a database table is refreshed for accelerator processing</p>	<ul style="list-style-type: none"><li>• Existing ETL process replaces entire table</li><li>• Multiple sources or complex transformations</li><li>• Smaller, un-partitioned tables</li><li>• Reporting based on consistent snapshot</li></ul>
<p><b><u>Table partition refresh</u></b></p> <p>For a partitioned database table, selected partitions can be refreshed for accelerator processing</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></ul>
<p><b><u>Incremental Update</u></b></p> <p>Log-based capturing of changes and propagation to DB2 Analytics Accelerator with low latency (typically few minutes)</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>

# Static SQL Support

*Available in DB2 Analytics Accelerator 4.1*

- Statically bound queries on active or archived data (discussed next) can be routed to the Accelerator
- New BIND options
  - ➔ QUERYACCELERATION
  - ➔ GETACCELARCHIVE
  - ➔ The possible values match the existing special register and zparm semantics
- Acceleration for static queries is determined and fixed at package bind time
  - ➔ Tables must be defined to an accelerator and enabled for acceleration prior to binding the package

Accelerator must be active and started when the static query runs

# High Performance Storage Saver

*Reducing the cost of high speed storage*

*Store historic data on the Accelerator only*

**Tables can be resident on:**

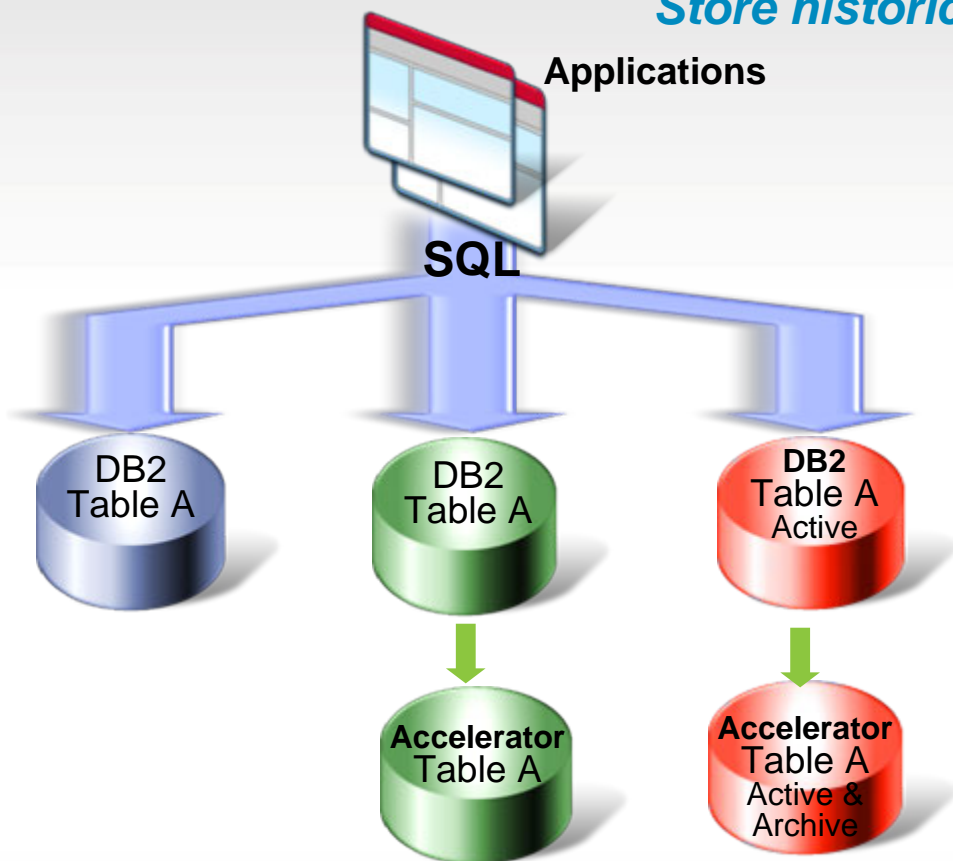
1. DB2 Only
2. DB2 and Accelerator
3. Archive to Accelerator

*When data no longer requires updating, reclaim the DB2 storage*

**Managed by zPARMs**

**Controlled by Special Registers:**

- CURRENT QUERY ACCELERATION
- CURRENT GET\_ACCEL\_ARCHIVE

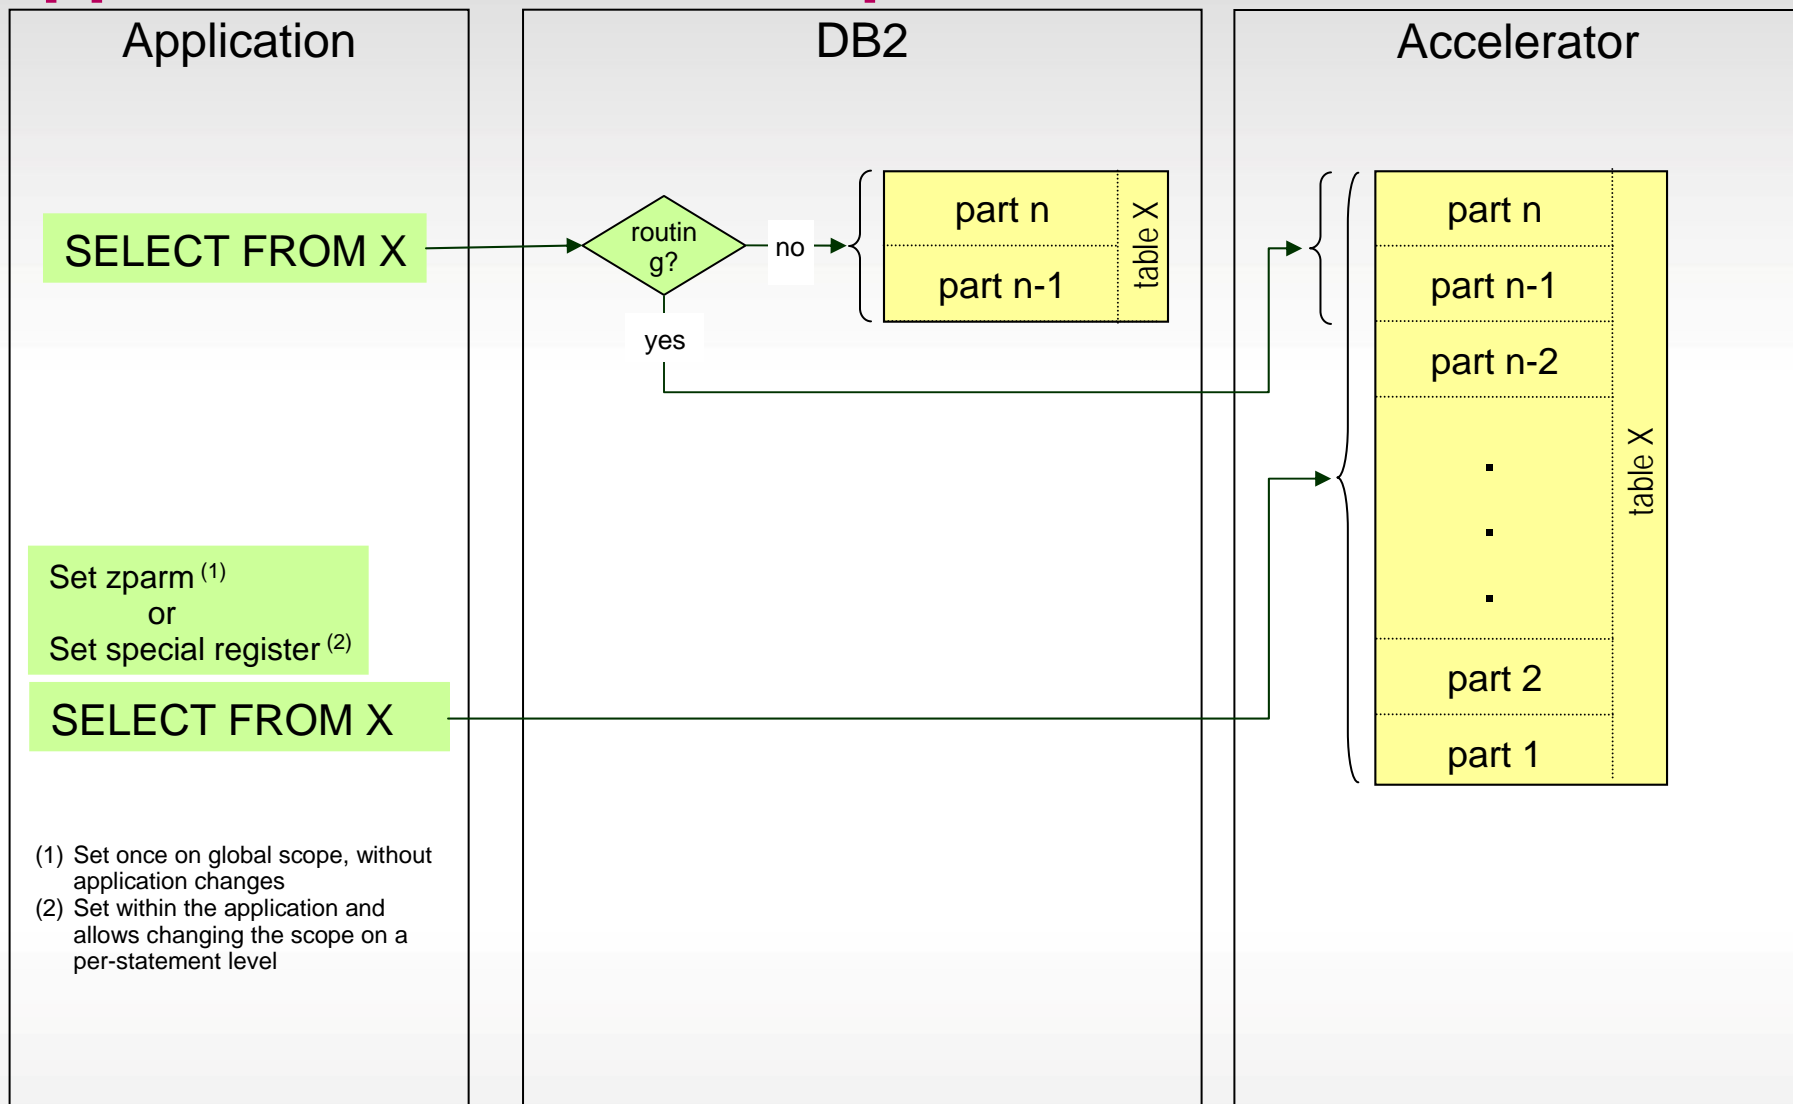


- *Best for OLTP*
- *High Speed Indexed queries*

- *Mixed Workload*

- *Active Only*
- *Archive Only*
- *Active & Archive*
- *Mixed Workload*

# Applications Have Transparent Access





A decorative graphic in the top left corner consists of several overlapping circles of various colors (yellow, orange, red, purple, blue) that are divided into segments, resembling stylized data points or a network diagram.

## **DB2 Analytics Accelerator Customer Experiences**



## What customers are saying...

- IBM says queries can run up to 2000x faster with the Accelerator, but we had one query run 4800x faster – from 4 hours to 3 seconds
- Our users call DB2 Analytics Accelerator the Magic Box
- Without acceleration, queries would take from several minutes to never returning – with DB2 Analytics Accelerator, queries return in less than 1 minute (usually 15 seconds)
- Whatever you paid for this, it was well worth it!
- It is unbelievable that there are still DB2 for zOS shops out there without IBM DB2 Analytics Accelerator



## A wide variety of use cases . . .

- Bringing an ODS back to z
- Analytics with the transactional detail data
- Web Log Analytics
- Accelerate scheduled reports
- Improve customer ad-hoc experience
- Lower the cost of SAS processing
- Meet ETL Service Level Agreements
- Accelerate operational reporting
- Accelerate DW queries
- Lower the cost of distributed DW and Mart extracts
- Lower the cost of batch billing cycles
- Warehouse Inventory Analysis
- Cognos and Business Objects reports
- Queries that were being killed by the Resource Limit Facility
- Call Center Analytics
- . . .



# Use Case One: Before & After the Accelerator



## Company Profile

- Over 1500 Branches
- Over 4 Million Customers
- 20,000 Employees
- Thousands of Business Users



## The Business Case

- Daily Transaction Log – Core Application
  - Application log of all the company's branch operations
  - Daily transactions include deposits, withdrawals, checks and adjustments for bank customer accounts
  - The head office needs access to this data (online and historical) to support bank operations

## Client Objectives

- Simplify existing Daily Transaction Log process
  - Lower costs and complexity
  - Eliminate unnecessary processes
  - Optimize use of compute resources
- Requirements
  - Strategic: Make more data accessible to end users (including operational and historical)
  - Operational: Simplify process for making data available to end users
  - Business: Make data accessible to improve customer satisfaction, managerial oversight and support regulatory requirements
- General:
  - Minimize ongoing costs associated with new technology
  - Improve performance of data accessible to end users
  - Maintain security of sensitive data
  - Provide high availability and reliability



## Before DB2 Analytics Accelerator

- Before Scenario
  - The Daily Log ran on a distributed Java platform using Hibernate framework
  - Each branch had a dedicated Daily Log in their DB2 Subsystem
  - Every night the daily data was copied from disk to tape via batch. The online data was then deleted; this was necessary to manage large volume of data
  - When data was needed by a user, a request was submitted to restore the data from tape to bring it back online. The data was restored at night via batch process
- The Business Need
  - The restore process caused a minimum 1 day of disruption to the end user productivity
  - Resource constraints required that the archived data be kept online for a maximum of 10 days

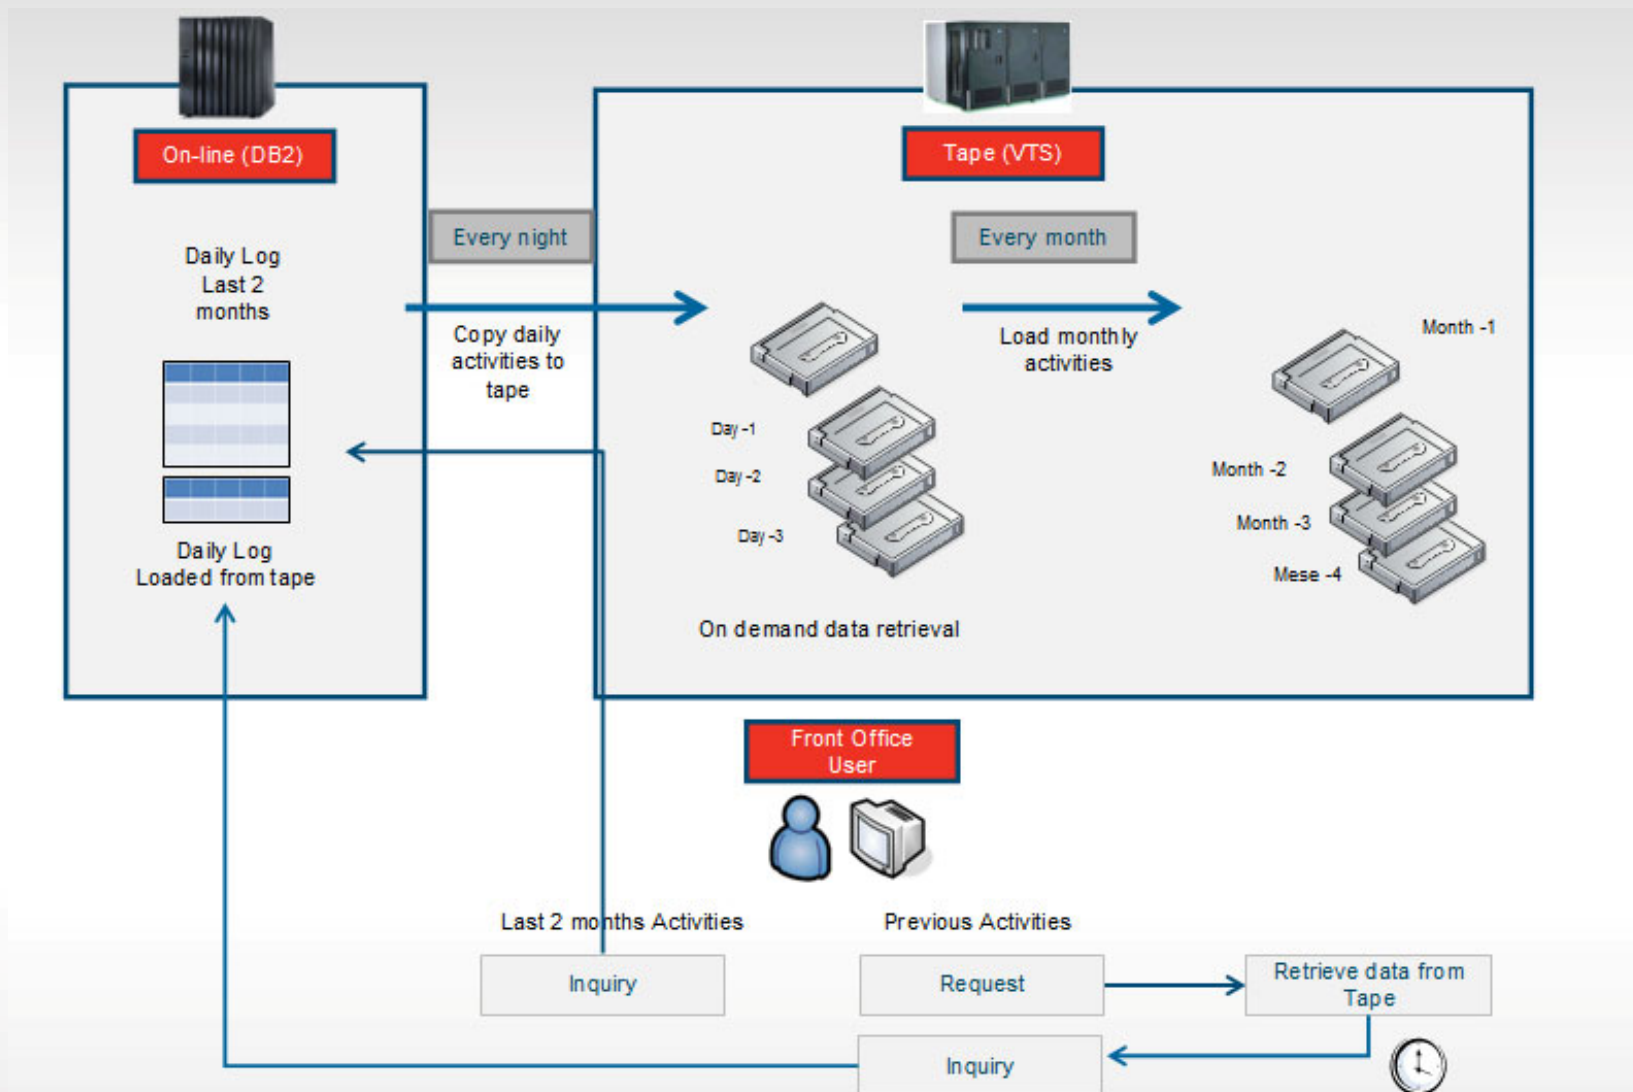




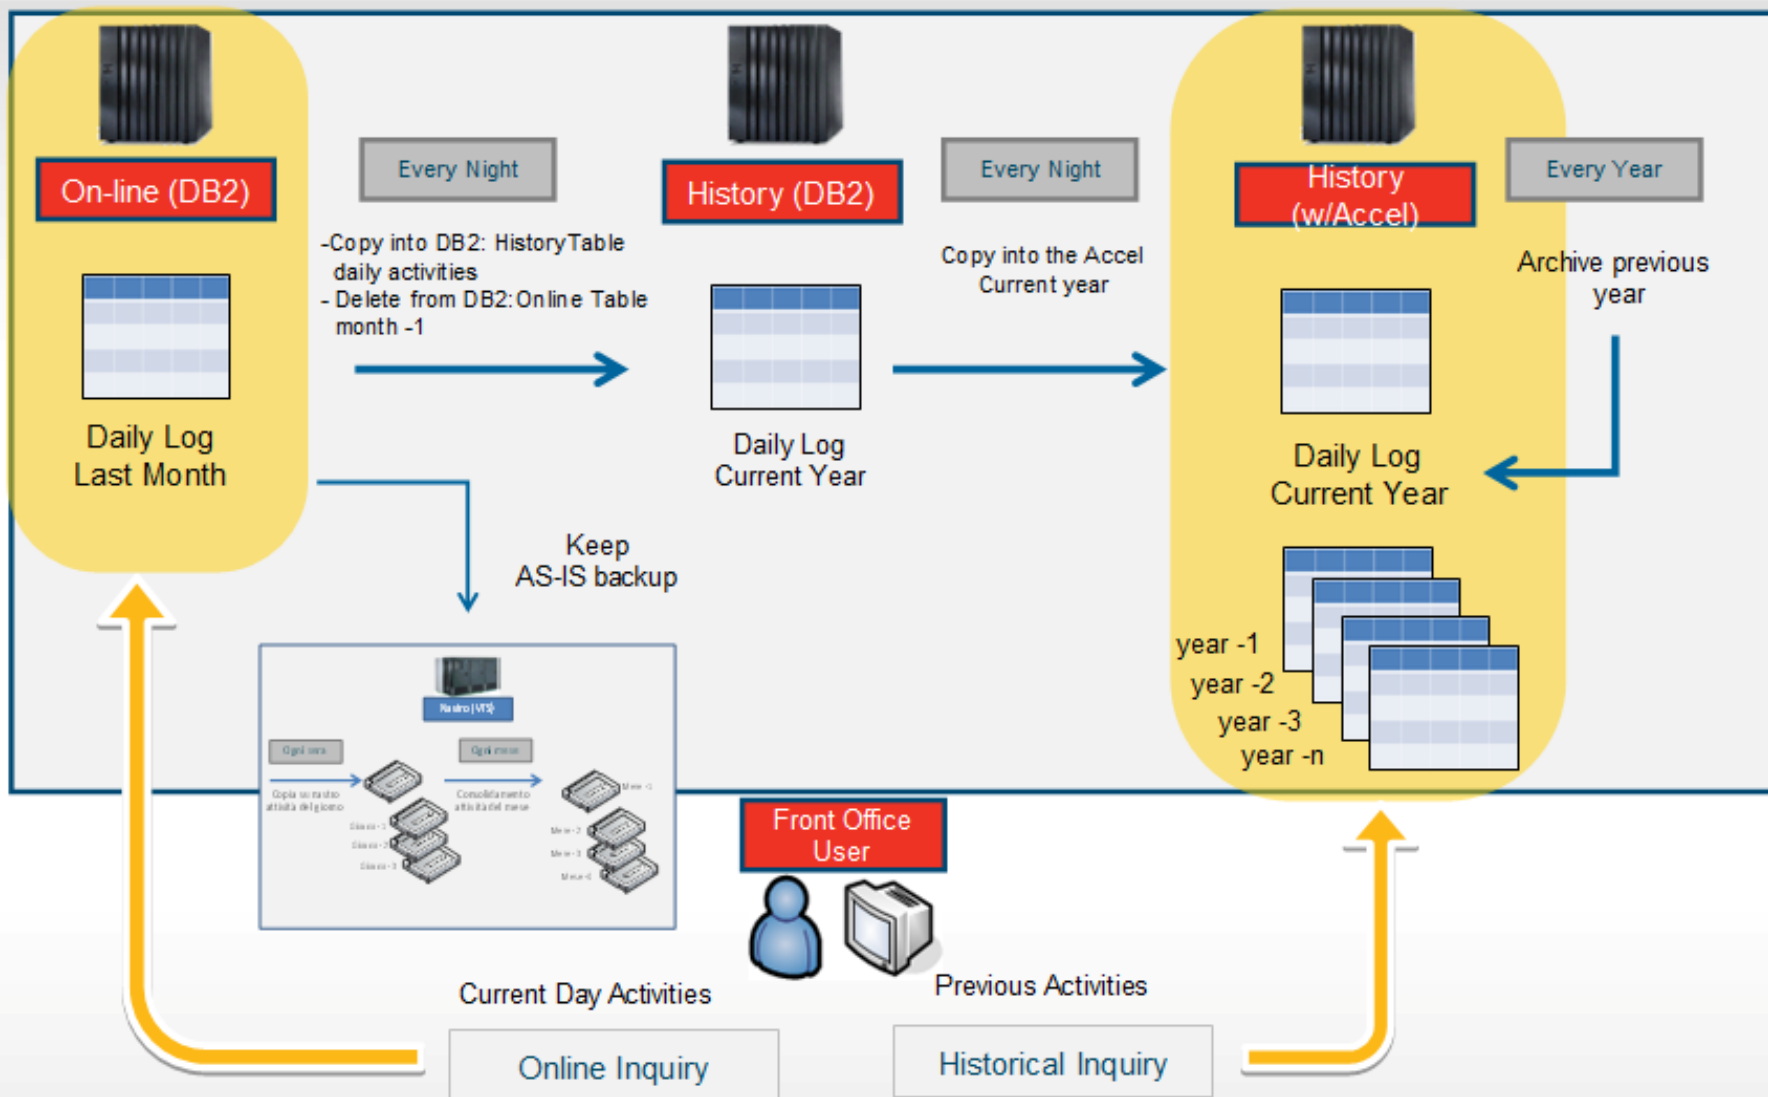
## After the DB2 Analytics Accelerator

- The IBM DB2 Analytics Accelerator allows:
  - All daily log data is available to the user - always
  - Historical data is accessible online as compared to a one day restoration process
  - Can now perform queries on large set of historical data
  - Eliminated the batch job that restored data from tape and to online tables
  - Eliminated the panel request process for historical data
  - Reduced online table sizes (keeping them separate from the historical data)
  - Ability to maintain only last month data in the online tables
  - Near zero CPU consumption to access historical data (the entire query is executed in the Accelerator)
  - Performance – Average response time is less than 1 second!

# Daily Log Process - Before



# Daily Log Process – After





## **Use Case Two: Before & After the Accelerator**



## Company Profile

- Over 1000 Branches
- Over 2 Million Customers
- Thousands of Business Users



## Client Objectives

- Modernize existing Analytics solution
  - Reduce time to deliver information and insight to business users
  - Utilize existing zEnterprise investments
- Requirements
  - Strategic: Deliver market analysis and intelligence to business users
  - Operational: Meet compliance reporting, regulation and business management requirements
  - Business: Improve time to market, satisfy customer requirements and deliver new customer specific services and offerings
- General:
  - Versatile, time and cost efficient solution that can leverage operational sources and provide near real-time analytics
  - Maintain security of sensitive data
  - Provide high availability and reliability



## Before DB2 Analytics Accelerator

- Initially industry specific applications were in place to fulfill the requirements
  - Resulted in a highly heterogeneous analytical environment
- An Enterprise Data Warehouse project was started adopting a distributed technology but a PoC failed due to the following problems:
  - Lack of integration with core business applications
  - Inability to scale to support production
  - Problems with hardware availability

**The top management decided to build a common platform for analytics applications based on cost/benefits and future trend/direction strategy**



## After the DB2 Analytics Accelerator

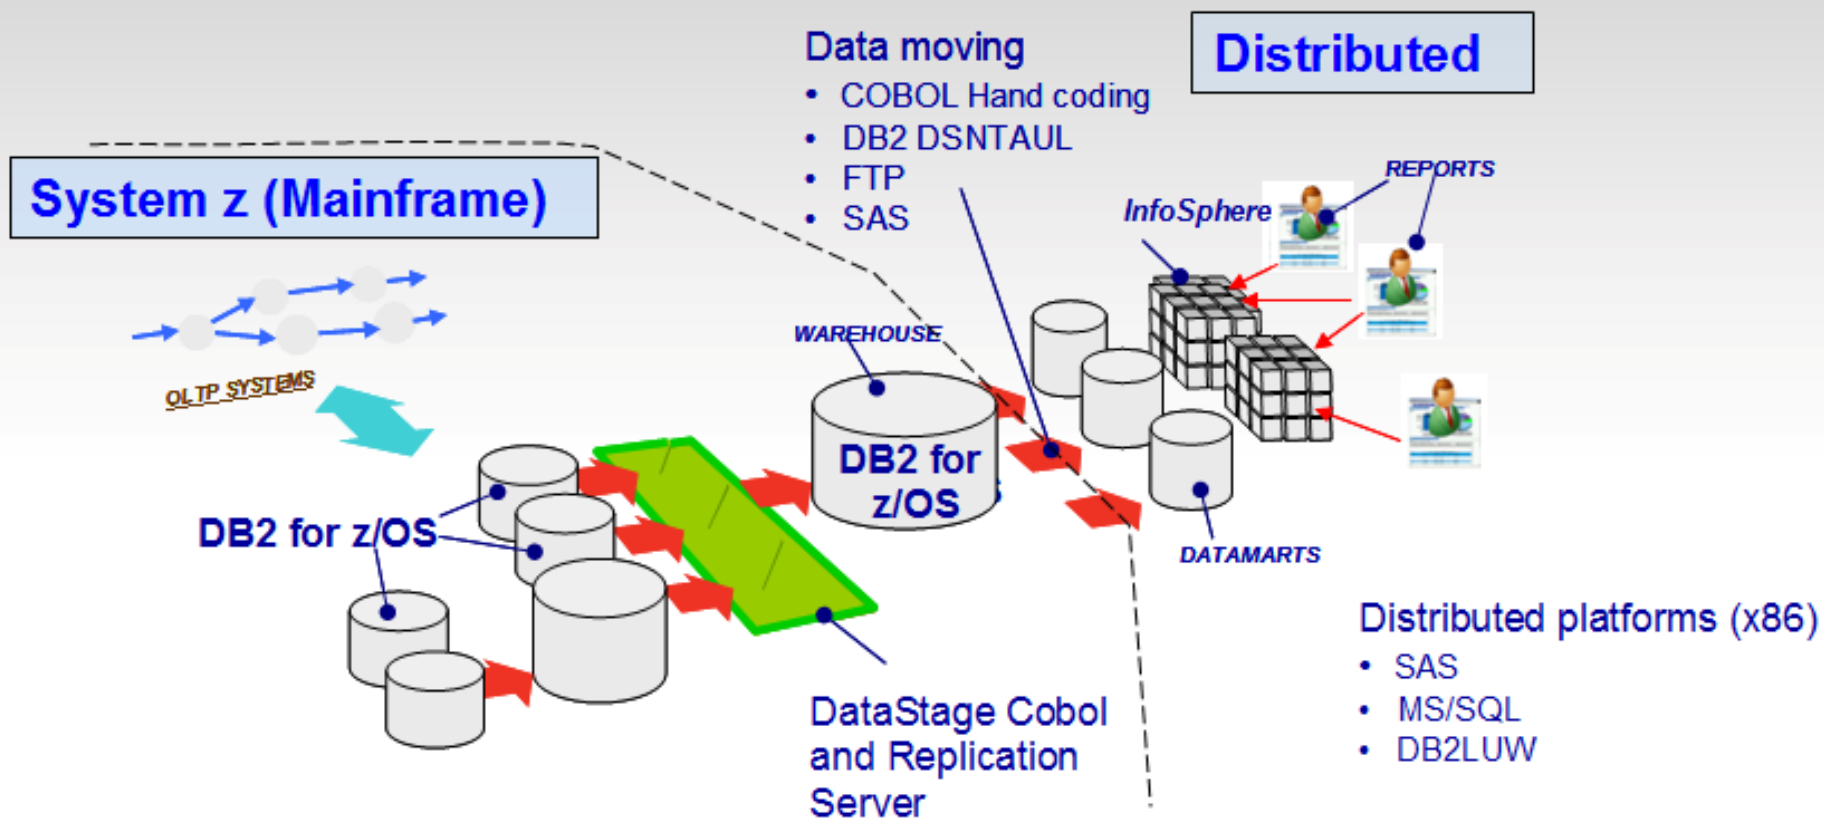
The IBM DB2 Analytics Accelerator and zEnterprise solution allows:

- Scalability to meet performance requirements, support for future solutions and an improved time-to-market (cut from months to weeks)
- Co-location of DWH and operational data provides optimized ETL process to load the most recent data to the DWH thus, marketing campaigns use accurate data
- Improved availability and user experience. Provide access to data anytime and in the required manner by exploiting the availability features of System z
- A DWH environment that ensures adoption of compliance regulations by exploiting the audit-ability and security features of System z
- Improved efficiency by changing from many thematic Data Warehouses to one scalable, versatile platform hence reducing the ETL overhead and utilizing one data source



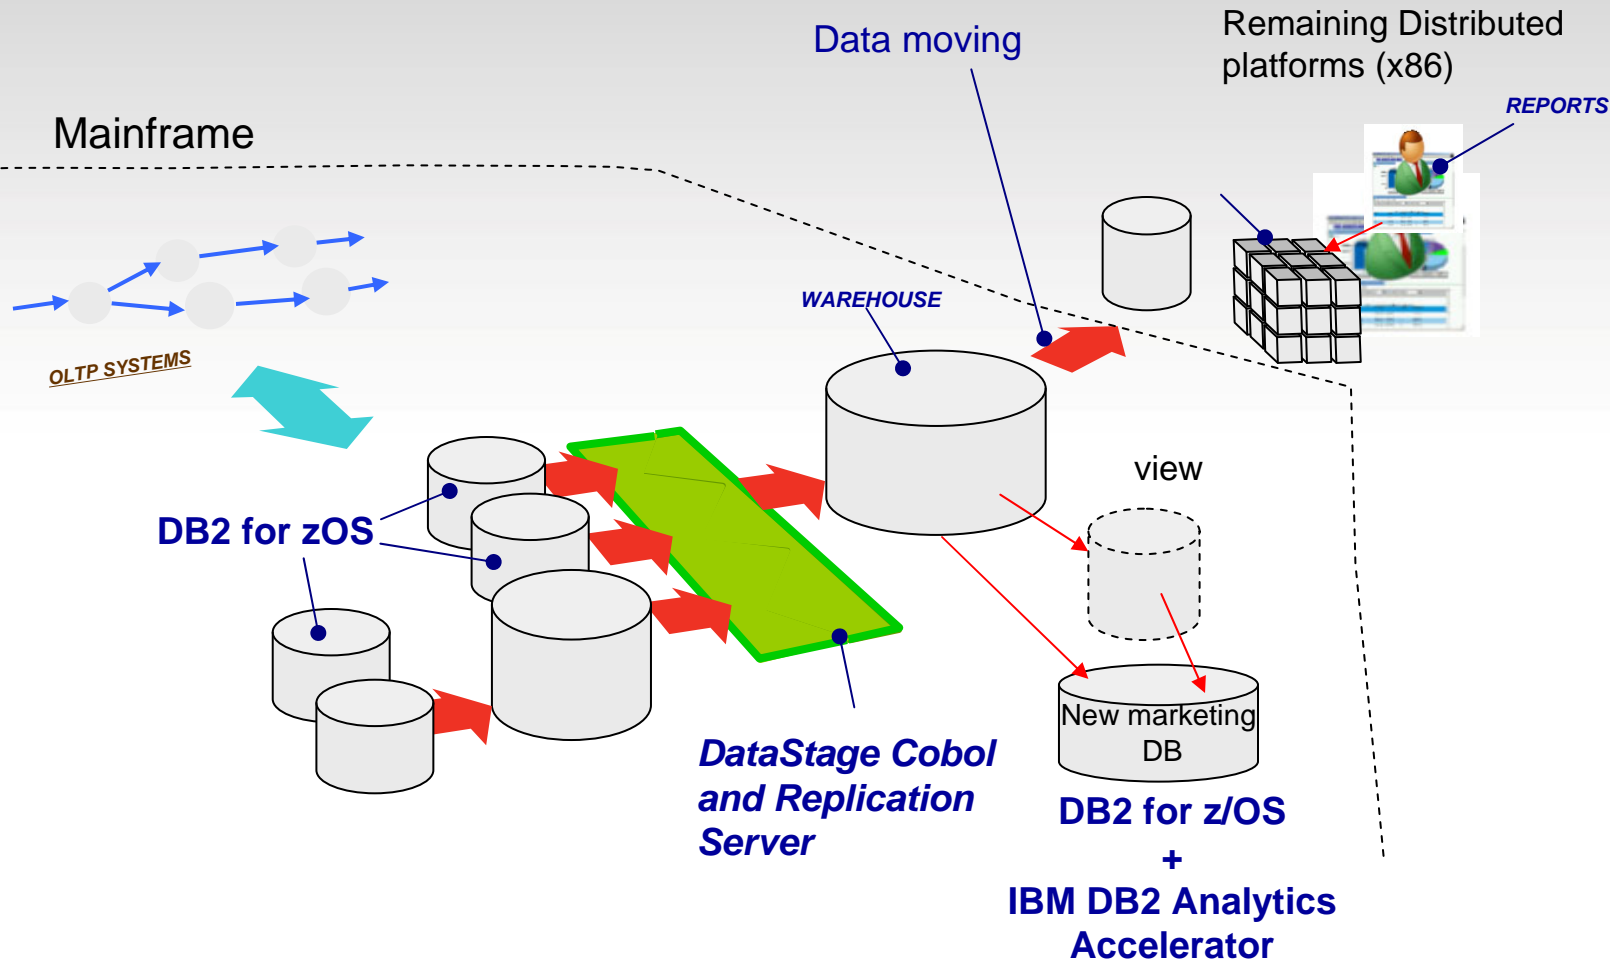


# Client's DWH solution - Before



- Domain focused data marts in support of industry-specific requirements
- Created siloed heterogeneous environment

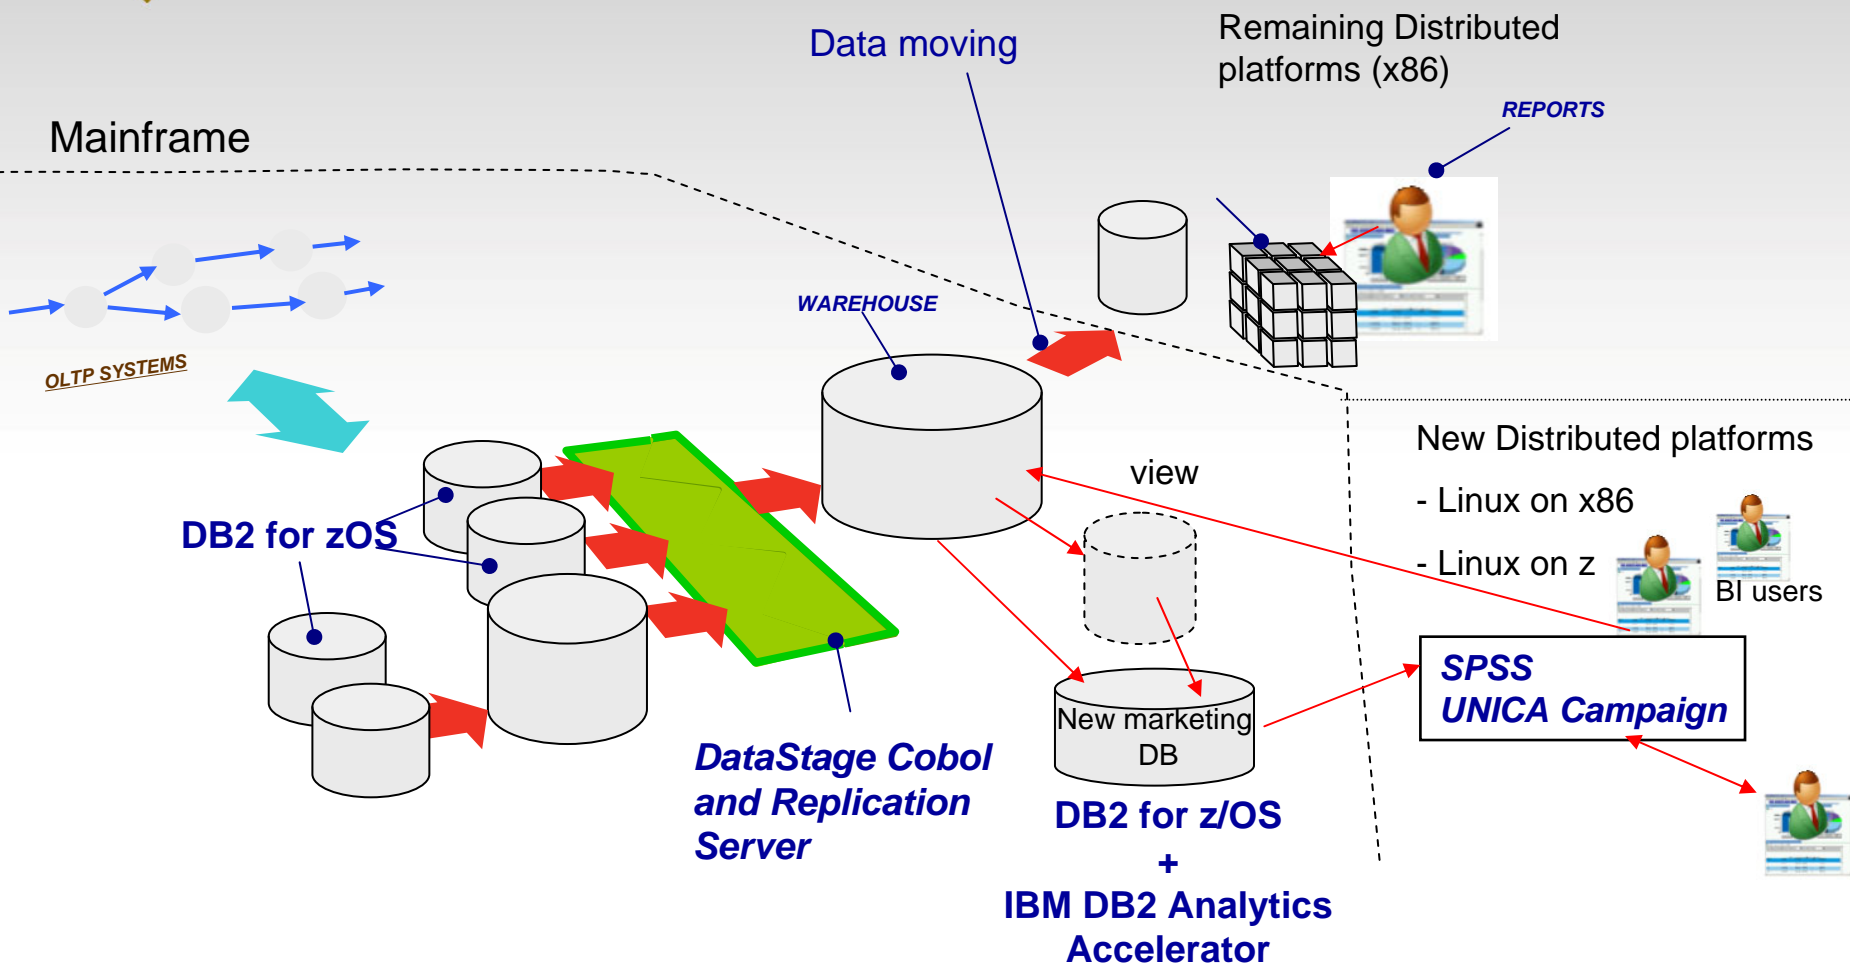
# Client's DWH solution - After



- The new environment utilizes the existing data infrastructure.
- The IBM DB2 Analytics Accelerator does not introduce additional efforts but provides quick response times to the user. New applications can utilize the environment with a more efficient use of time and resources.



# Client's DWH solution – After



- New Environment after UNICA Campaign Management



# **Best Practices for Assessing Workloads and Requirements**



# Snapshot Assessment and Package Check – Report



## IBM DB2 Analytics Accelerator -- Workload Analysis Results

Center of Excellence, Datawarehouse on System z, IBM Research & Development, Böblingen, Germany

### Summary for Dynamic Query Statements

This report analyzes SQL select statements from the DB2 dynamic statement cache to determine potential DB2 Analytics Accelerator benefits.

These queries are divided into 3 categories:

'With potential' (green) contain SQL syntax that may run on the accelerator.

'Uncertain' (yellow) contain SQL syntax that require further investigation.

'No potential' (red) contain SQL syntax that is ineligible for query acceleration.

The DB2 z/OS optimizer uses heuristics to further analyze which queries should continue running natively on DB2 z/OS, marked in blue. The number of blue queries can be influenced by acceleration settings, for example 'enable', 'eligible', or 'all'.


Note that this report is only meant as an indicator of the potential benefits, actual query acceleration may vary depending on your particular configuration.

Start trace time	End explain time	Min time stmt cached	Max time stmt cached
Aug 18, 2013 4:33 AM	Sep 1, 2013 7:19 PM	Aug 18, 2013 12:52 PM	Sep 1, 2013 7:19 PM

	Total	With potential	Uncertain	W/o potential
Queries	17500	17162 (98%)	153 (1%)	185 (1%)
Elapsed Time (s)	150653.49	138476.55 (92%)	10577.51 (7%)	1599.42 (1%)
CPU Time (s)	17440.15	15734.61 (90%)	1146.62 (7%)	558.91 (3%)

Reason breakdown for 1% queries with no potential	# Queries	% Queries	% Elapsed Time
Unsupported function	7	0%	0%
Not read only	4	0%	0%
Not select	19	0%	0%
Unsupported column type	2	0%	1%
Unsupported encoding	50	0%	0%
Common table expression	2	0%	0%
Unsupported special register	89	1%	0%
Data-change-table reference	36	0%	0%
Sequence expression	41	0%	0%
Quant_predicates (temp_restr.)	11	0%	0%
Unsupported cast	1	0%	0%

### Total Number of Queries: 17500



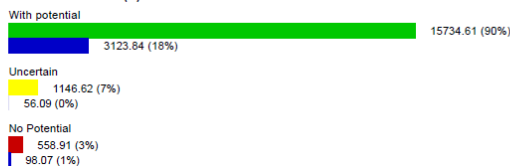
Other statements: 41686

### Total Elapsed Time (s): 150653.49



Other statements: 94192592.96

### Total CPU Time (s): 17440.15



Other statements: 98.07 (1%)

Elapsed time potential

CPU time potential

Query	Elapsed Time (s)	CPU Time (s)	Executed	Eligible / DB2 native	Rows exam./proc.	Table Sizes	IDAA Time	Comment	State
514365	3928.05	480.53	1	With potential	40 M / 5 M	69.41 GB	10-100x (N2001-010) 10-100x (N2001-005) 10-100x (N1001-010) 10-100x (N1001-005) 10-100x (N1001-002)	WHEN ORIG ENDA DFIDF DFIDF DFIDF LOAN DFIDF DB_CF FROM PROC LOAN, OUTEI WHER DFIDF INTFC DFIDR	(Prima

Query details

Queries by elapsed time

### Query Classification by Elapsed Time

Query classification	Unique Queries	Executed Queries	Total elapsed (s)	Total CPU (s)	CPU %	CPU eligible %
Total	17500	394253	150653.49	17440.15	100%	90%
> 60 min elapsed time	8	8	52385.17	3028.19	17%	9%
30-60 min elapsed time	9	9	23916.58	1901.01	11%	6%
10-30 min elapsed time	33	33	28675.61	4702.28	27%	21%
5-10 min elapsed time	41	43	19165.81	2754.02	16%	15%
1-5 min elapsed time	121	129	18207.67	3188.96	18%	16%
50-60 s elapsed time	19	28	1532.47	311.24	2%	1%
40-50 s elapsed time	24	27	1197.49	230.98	1%	1%
30-40 s elapsed time	33	34	1167.64	182.70	1%	1%
20-30 s elapsed time	31	35	827.03	163.49	1%	0%
10-20 s elapsed time	55	60	864.90	195.65	1%	0%
5-10 s elapsed time	64	71	515.66	150.85	1%	0%
0-5 s elapsed time	17033	393776	2197.44	630.78	4%	0%

# Accelerator Modeling

- Provides indicators for possible CPU and elapsed time savings if an IBM DB2 Analytics Accelerator was available
  - It does not require presence of the Accelerator
- DB2 11 or DB2 10
- Controlled by new zparm ACCELMODEL which can be set to YES or NO
  - If set to YES, DB2 accounting records (IFCIDs 3 and 148) include projected CPU and elapsed time savings
  - CURRENT QUERY ACCELERATION must be set to NONE
    - However, EXPLAIN will still indicate if the query is eligible for acceleration and, if not, the reason why in DSN\_STATEMENT\_TABLE.REASON
  - Like with any DB2 instrumentation, the new timers need to be formatted and reported by a monitor
- Functionality delivered via two DB2 10 APARs
  - PM90886: Dynamic SQL
  - PM95035: Static SQL
    - REBIND needed to enable acceleration modeling

# Accelerator Modeling Report

MEASURED/ELIG TIMES	APPL (CL1)	DB2 (CL2)
-----	-----	-----
ELAPSED TIME	4.830139	4.740227
ELIGIBLE FOR ACCEL	N/A	4.442327
CP CPU TIME	6.337894	6.336111
ELIGIBLE FOR SECP	4.990042	N/A
ELIGIBLE FOR ACCEL	N/A	6.329119
SE CPU TIME	0.000000	0.000000
ELIGIBLE FOR ACCEL	N/A	0.000000

1

2

3

- 1 Elapsed time that can be significantly reduced because the qualifying statements in the reported plan execution could be routed to the accelerator. If the statements are executed in parallel, the reduced elapsed time relates to the parent task only.
- 2 The part of CPU time spent on general purpose processors that can be saved to a large extent because the qualifying statements in the reported plan execution could be routed to the accelerator. If the statements are executed in parallel, the CPU saving includes the parent and all the subordinated parallel tasks.
- 3 The part of CPU time spent on specialty engine processors that can be saved to a large extent because the qualifying statements in the reported plan execution could be routed to the accelerator. If the statements are executed in parallel, the CPU saving includes the parent and all the subordinated parallel tasks.



# **Delivering Analytics Applications on zEnterprise**



# zEnterprise Analytics System 9700 and 9710

## *Flexible zEnterprise Deployment Models*

***A cost-competitive, integrated combination of hardware, software and services to deliver business reporting and business critical analytics***

- **Solution Priced** - for deployment as an additional logical partition (LPAR) on an existing system or as a new system
- **Preselected** – to deliver a comprehensive, yet flexible end-to-end solution
- **Pretested** – to meet business reporting and critical analytic demands

### *zEnterprise Analytics System 9700*



*zEnterprise EC12*

### *zEnterprise Analytics System 9710*



*zEnterprise BC12*

# Analytics on IBM zEnterprise



## An end-to-end solution

- IBM DB2 Analytics Accelerator
- IBM Cognos Business Intelligence
- IBM Capacity Management Analytics
- IBM Cognos TM1
- IBM Query Management Facility (QMF)
- IBM SPSS Predictive Analytics
- IBM zEnterprise Analytics System 9700/9710

## Delivers on business critical analytics

- Timely, accurate, secure data
- Availability, scalability, performance
- Rapid deployment & expansion
- Reduced cost and complexity

## Evolves with the business

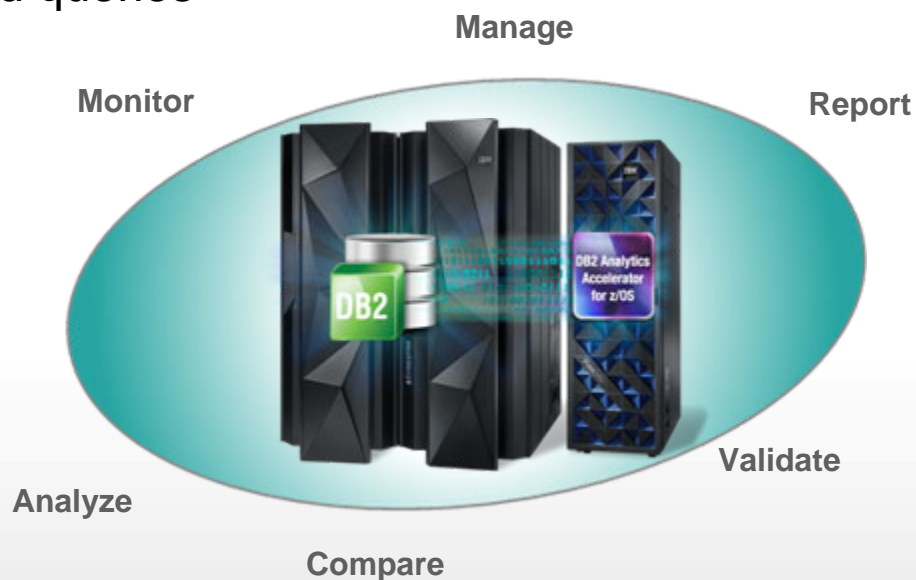
- Start where you want
- Grow without re-architecting

## Supports the Big Data Requirements

- Volume, Velocity, Variety and Veracity
- Simplified integration of Hadoop, streaming and transactional data

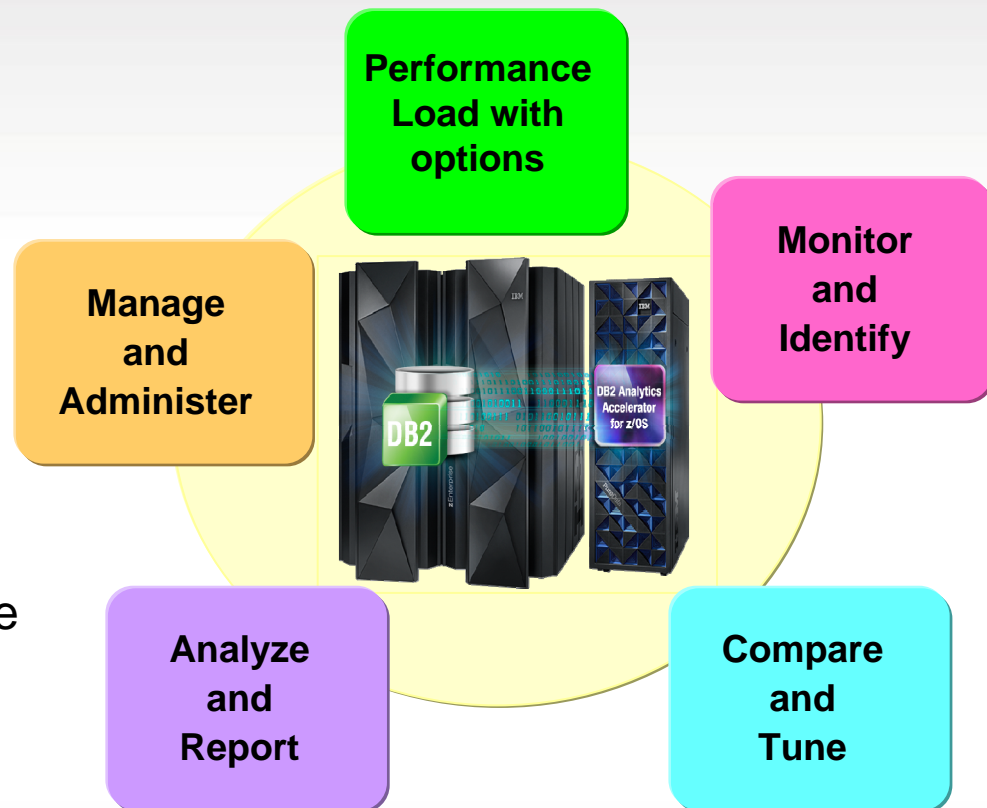
# Surround DB2 Analytics Accelerator with Powerful DB2 Tool Capabilities to Further Maximize Analytics Investments

- Analyze and report on Accelerator utilization
- Proactively monitor and manage your accelerated queries
- Accelerate load performance to Accelerator with more options to fit business needs
- Validate performance ROI of accelerated queries
- Streamline query candidate selection
- Determine best set of objects based on workload
- Filter out un-accelerated queries to maximize workload tuning efforts
- Administer and manage Accelerator



# IBM DB2 Tools: Monitor and Manage Capabilities

- Three separate areas where DB2 tools can add value to your Accelerator investment:
  - **Assessment**
    - Do I have a workload that would benefit from the Accelerator?
  - **Optimization**
    - How do I optimize a workload to take advantage of the Accelerator?
  - **Administration**
    - Can I manage the Accelerator more effectively?





# DB2 Tools and Accelerator: Optimize, analyze and report



- OMEGAMON XE for DB2 PE
  - Detailed performance measurements on Accelerator and accelerated query performance via online monitoring or batch reporting
  - Determine and assess workload benefits of acceleration
  - Analyze performance trends with Accelerator-specific performance metrics for historical usage and capacity planning
- Query Monitor for DB2
  - Provides measurable performance comparisons to show ROI justification for accelerated queries
  - Provides filtering to let you see accelerated and non-accelerated queries to aid in tuning efforts
  - Effective workload tuning with Query Workload Tuner integration
- Optim Query Workload Tuner – Workload advisor capabilities
  - Analyzes workloads for Accelerator offload eligibility and identifies eligible, ineligible and rewritable queries
  - Recommends which tables should be added or deleted
  - Quantify CPU savings and utilization



# DB2 Tools and Accelerator: Optimize, manage and load



- DB2 Administration Tool – Manage from ISPF interface
  - Management of Accelerators and accelerated objects
  - Change management support for reloading data after a change
  - Data currency information from Real-time stats to help determine if refresh of the data is needed
- DB2 Analytics Accelerator Loader
  - Flexible use cases to meet business needs
    - Simplify process and reduce cost when loading to both DB2 and the Accelerator
    - Minimize application impact on DB2 while loading data to the Accelerator
    - Enables data refresh on the Accelerator according to your point-in-time business needs
    - Enhance business analytics with greater use of System z and DB2 for z/OS investment
  - Built for performance
    - Can take advantage of zIIP support, FlashCopy, and DB2 Sort for z/OS



# Transaction to Business Action - Single zPlatform

*Delivering vital business insight better, faster and cost effectively*

- OLTP and analytics on the same platform



Reduce complexity and data latency, improve security and data governance, make near real-time business decisions

- Industry leading qualities of service and security



Confidently deliver business critical analytics

- Leverage existing zEnterprise infrastructure, processes and people



Reduce total cost of ownership

- Price competitive, easy deployment



- Deliver value quickly and cost effectively



# Thank You

Your feedback is important!