



The DB2 Analytics Accelerator: new capabilities and use cases



Agenda

- Overview of the DB2 Analytics Accelerator for z/OS
- New capabilities
- New use cases

Overview of the DB2 Analytics Accelerator for z/OS

The IBM DB2 Analytics Accelerator (1)

- What is it?
 - An integration of IBM PureData technology (formerly Netezza) with zEnterprise technology that delivers dramatically faster business analytics
- What does it do?
 - Accelerates complex queries, up to 2000x faster
 - Improves access to and lowers the cost of storing, managing and processing historical data
 - Minimizes latency
 - Reduces zEnterprise capacity requirements
 - Improves security and reduces risk



* Via the IBM DB2 Analytics Accelerator Loader

IBM DB2 Analytics Accelerator (2)



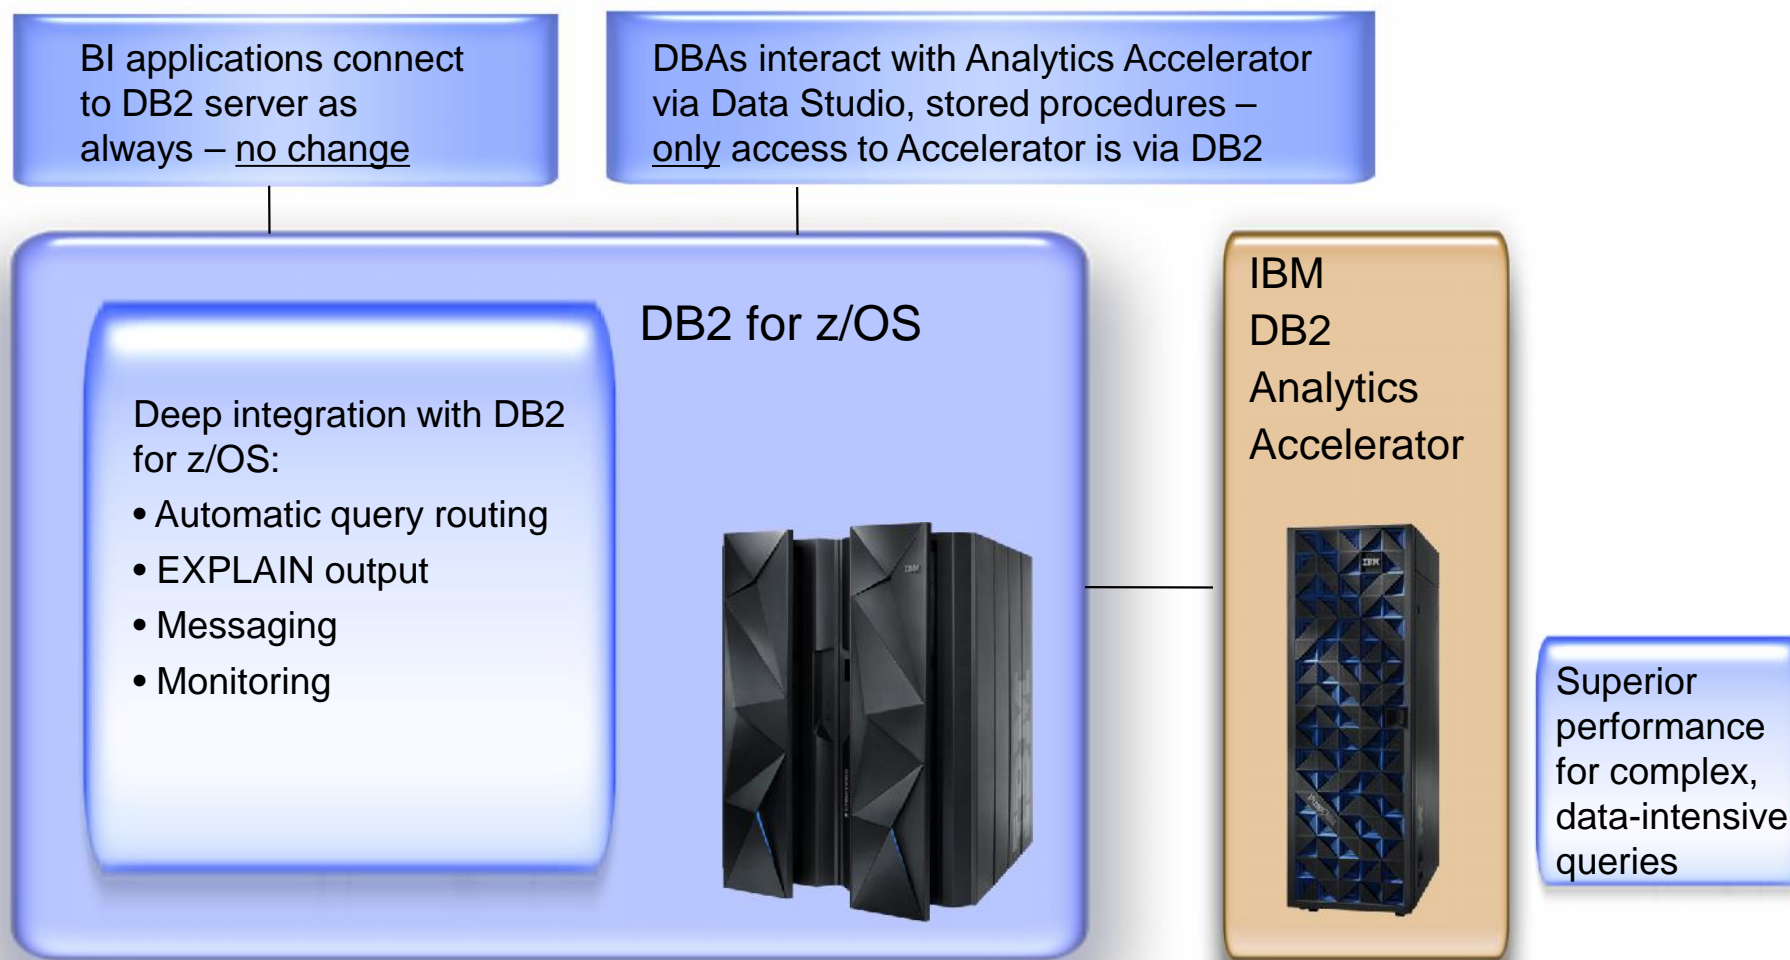
* Via the IBM DB2 Analytics Accelerator Loader

Enabling data-driven insight

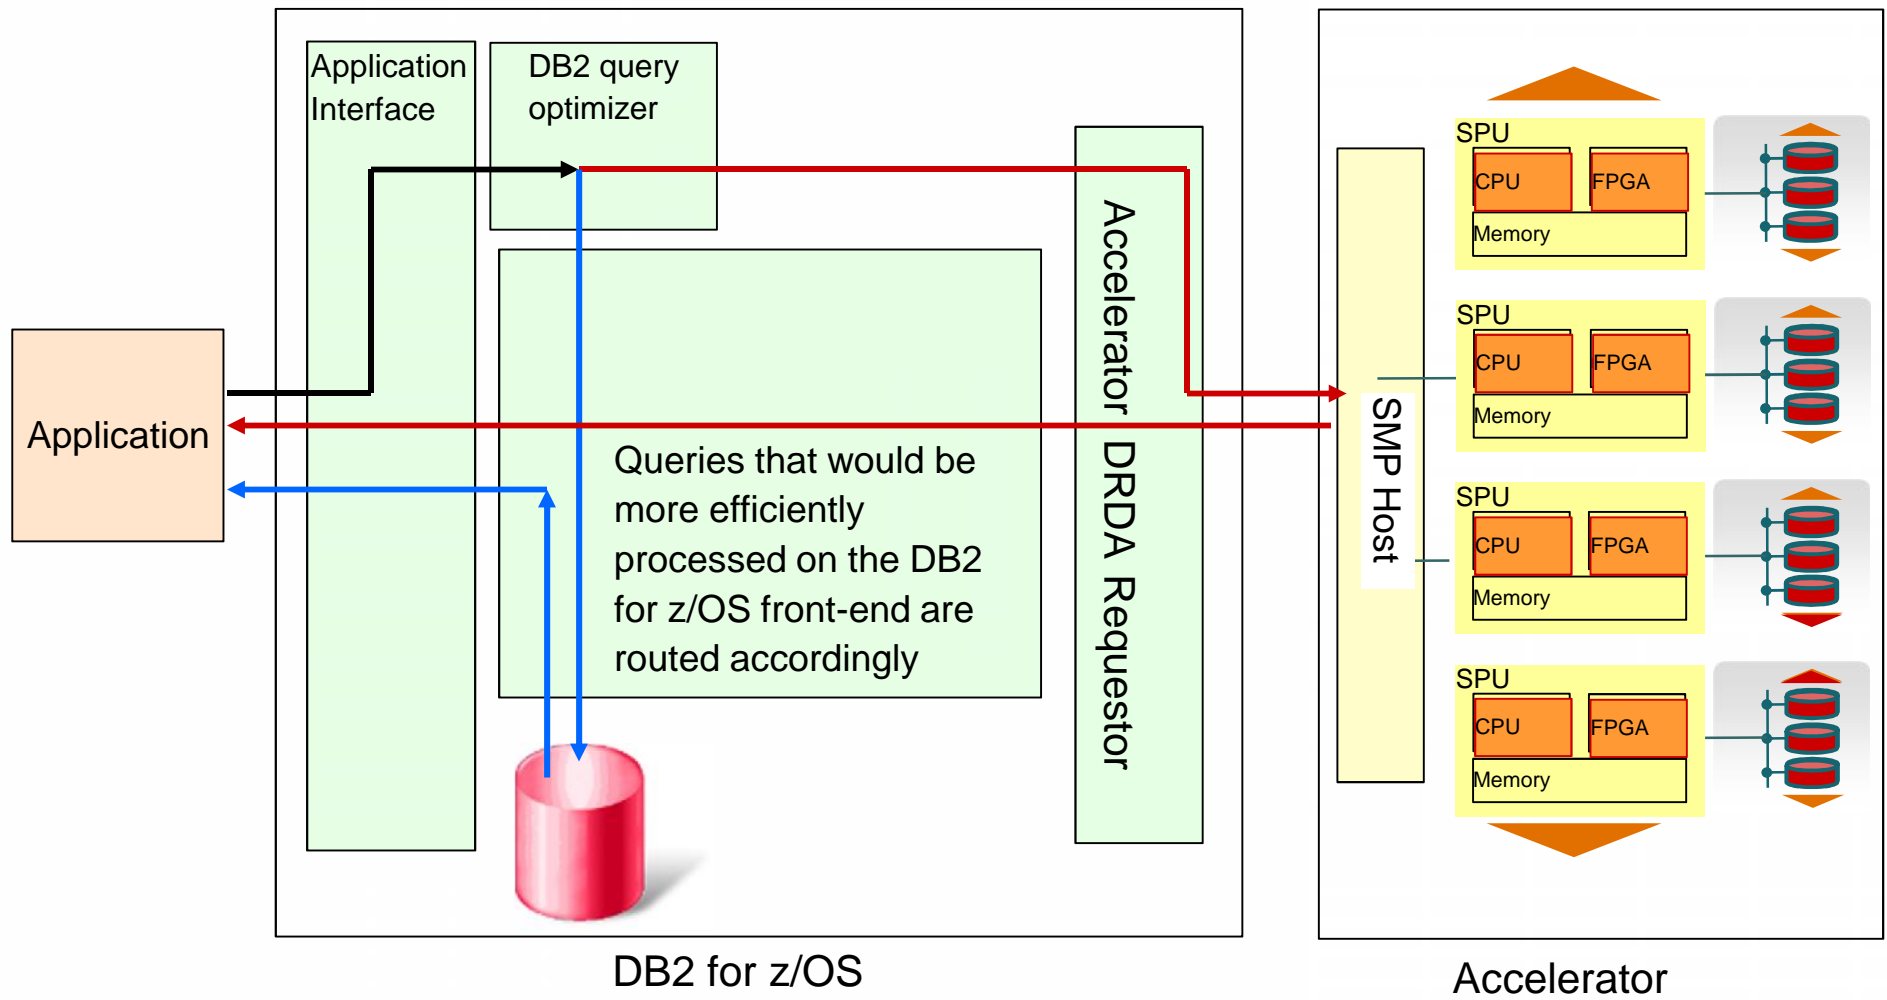
- Transparent to applications and reporting tools that access DB2 data
- Inherits DB2 for z/OS data security attributes
- Simplified database design – no need for indexes
- Eliminates need for tuning of analytic queries
- Fast deployment and time-to-value
- Join data from IMS, DB2 and other sources, with superior performance

Why it's transparent

A DB2 Analytics Accelerator is an extension of a DB2 for z/OS system



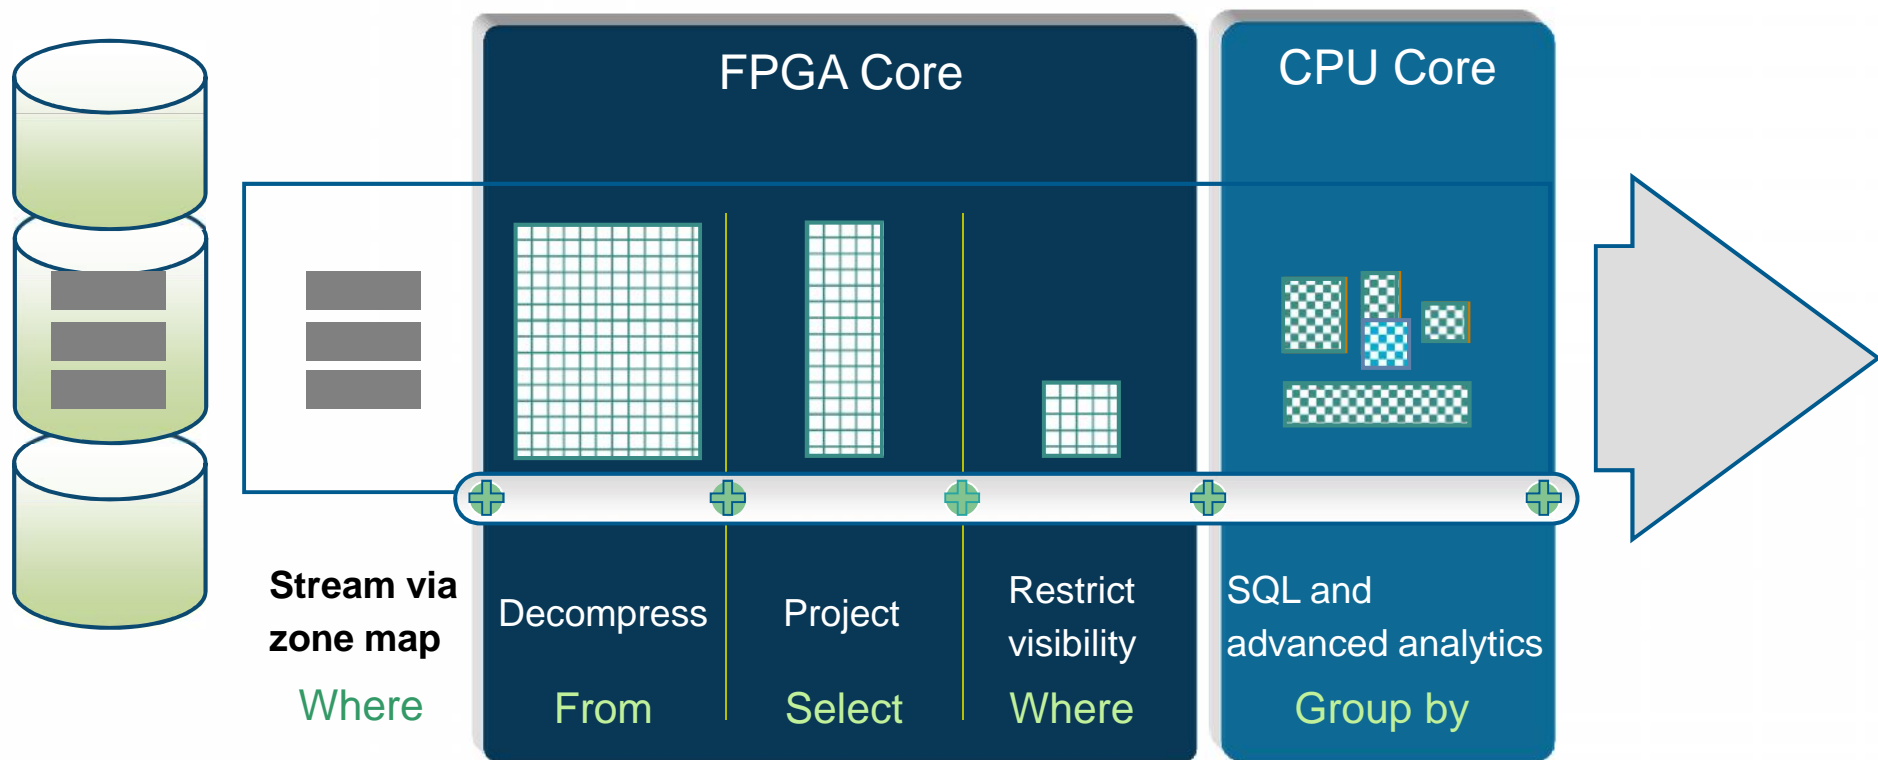
Query execution process flow



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

Why it's so fast

- More than just highly parallelized query execution (plenty of systems do that)
- What's really different: result set row and column filtering as *data is streaming from disk into memory*




```
Select State, Age, Gender, count(*) From MultiBillionRowCustomerTable Where BirthDate < '01/01/1960'
And State in ('FL', 'GA', 'SC', 'NC') Group by State, Age, Gender Order by State, Age, Gender
```


DB2 Analytics Accelerator data change propagation options

Synchronization options	Use cases, characteristics and requirements
<p>Full table refresh</p> <ul style="list-style-type: none"> The contents of a table that has been “accelerated” (i.e., a copy resides on the Accelerator) are refreshed in their entirety 	<ul style="list-style-type: none"> Existing ETL process replaces entire table Multiple sources or complex transformations Smaller, un-partitioned tables Reporting based on consistent snapshot
<p>Table partition refresh</p> <ul style="list-style-type: none"> For a partitioned table, selected partitions can be refreshed for accelerator processing 	<ul style="list-style-type: none"> I deal for tables partitioned on a date/time basis (data inserts and updates tend to be limited to “last” partition) More efficient than full table refresh for larger tables Reporting based on consistent snapshot
<p>Incremental update (“trickle feed”)</p> <p>Log-based capture of data changes, with immediate propagation to DB2 Analytics Accelerator for reduced latency (typically 1-2 minutes)</p>	<ul style="list-style-type: none"> Scattered updates after “bulk” load Reporting on continuously updated data (e.g., an ODS), considering most recent changes More efficient for smaller updates than full table refresh

High Performance Storage Saver feature

Storing historical data in Accelerator only



- Very well suited to situation in which:
 - Table is partitioned on a date/time basis
 - Only the most current partitions are used in a transactional context (frequent data changes, short-running queries)
 - The entire table is used for analytics (data intensive, complex queries).

New capabilities

DB2 Analytics Accelerator Version 4.1 – enhancements

- Version 4.1 became generally available in November of 2013
- Enhancements provided with this release include the following:
 - Support for static SQL queries
 - Improved performance for incremental data update
 - Enhanced usability and robustness for the High-Performance Storage Saver
 - Extended integration with the z/OS Workload Manager
 - Rowset cursor and multi-row fetch support for local-to-DB2 applications
 - Support for a mix of EBCDIC and Unicode tables in an Accelerator

Static SQL support

- The most requested feature since the Accelerator's first release
 - The first several releases of the Analytics Accelerator supported routing of only dynamic SQL statements to the Accelerator
 - Dynamic SQL tends to be dominant in query and reporting applications
 - Many organizations, however, run large reporting workloads – often batch jobs written in languages such as COBOL – involving execution of static SQL statements
- Version 4.1 of the DB2 Analytics Accelerator extended support to static SQL statements

More on DB2 Analytics Accelerator static SQL support

- New BIND and REBIND PACKAGE options give you control
 - QUERYACCELERATION
 - NONE – no acceleration of static queries
 - ENABLE – static query will be accelerated if that would optimize performance, and then must execute on Accelerator
 - ENABLE WITH FALLBACK – like ENABLE, but query can run in DB2 front-end if execution on Accelerator fails or returns error
 - ELIGIBLE – like ENABLE, but query will be accelerated if possible, even if that would not optimize performance (e.g., for a simple, quick query)
 - ALL – all static queries are to be accelerated; bind fails if a static query in the package cannot be accelerated
 - GETACCELARCHIVE (NO | YES)
 - Refers to data that has been “archived” to the Accelerator via the High-Performance Storage Saver feature (slide 10)

Enhanced performance for incremental data update

- Referring here to “trickle feed” versus bulk data refresh
- Performance improvement gained through leveraging of a DB2 11 log-based data change capture enhancement (retrofitted to DB2 10 via the fix for APAR PM90568)
 - Often, not all front-end DB2 tables defined with DATA CAPTURE CHANGES are targets for incremental update in Accelerator
 - Previously, data change capture task (using DB2 IFI 306 to interact with DB2 log manager) had to filter out data changes that are not to be propagated to target system (e.g., an Accelerator)
 - Now, DB2 will do that filtering (based on object IDs provided with IFI 306 read request), and will provide to capture task only pertinent changes
 - Result: significant reduction in CPU cost of data change capture task
 - Especially when data changes associated with compressed table spaces

High Performance Storage Saver enhancements

- Archived partitions are placed into a new Partition Read Only (PRO) state that prevents data modifications
- Several image copy enhancements:
 - No new image copies can be created for partitions in the PRO status
 - Up to 4 image copies per partition can be created
 - Naming schema based on templates
- Restore of archived partitions encapsulated in an administrative stored procedure
- Table can be archived in multiple accelerators
 - Image copy used as the source for subsequent accelerators

Some other DB2 Analytics Accelerate Version 4.1 enhancements

- z/OS Workload Manager application priority passed to Accelerator for local as well as network-connected DB2 applications
 - Complements static SQL support
- Rowset query offload and multi-row fetch support for local applications
 - Improves performance of data retrieval for queries with large result sets issued by local-to-DB2 applications
- Multiple code page support
 - A mix of EBCDIC and Unicode tables from the same DB2 subsystem are now allowed on the Accelerator
 - Note: a query that accesses both EBCDIC and Unicode tables cannot be offloaded to the Accelerator

Also new: Accelerator modeling

- Provides indicators for estimated CPU and elapsed time savings *if queries were directed to an Accelerator for execution*
 - Does not require presence of an Accelerator
- Supported by DB2 10 (via APARs PM90886 and PM95035) and DB2 11 (via APAR PM96478)
- Controlled by new ZPARM parameter ACCELMODEL(NO | YES)
 - If YES, DB2 accounting records include projected CPU (on front-end DB2 system) and elapsed time savings that would result from routing of queries to an Accelerator
- Monitor support provided by IBM Tivoli OMEGAMON XE for DB2 Performance Expert (or Performance Monitor), via APAR PM94202
- REBIND needed to enable acceleration modeling for static queries in a package

Accelerator modeling as reported by OMEGAMON for DB2

MEASURED/ELIG TIMES	APPL (CL1)	DB2 (CL2)	
-----	-----	-----	
ELAPSED TIME	4.830139	4.740227	1
ELIGIBLE FOR ACCEL	N/A	4.442327	
CP CPU TIME	6.337894	6.336111	2
ELIGIBLE FOR SECP	4.990042	N/A	
ELIGIBLE FOR ACCEL	N/A	6.329119	
SE CPU TIME	0.000000	0.000000	3
ELIGIBLE FOR ACCEL	N/A	0.000000	

- 1** Elapsed time that can be significantly reduced because the qualifying statements in the reported program execution could be routed to an Accelerator
- 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 program execution could be routed to an Accelerator
- 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 program execution could be routed to an Accelerator

Also new on the scene: the DB2 Analytics Accelerator Loader

- A separately licensed utility that supports loading data into:
 - The DB2 for z/OS front-end and the DB2 Analytics Accelerator, or
 - *Only the Accelerator*
- Allows loading of data from:
 - DB2 image copy files
 - Also: DB2 image copy + log files (i.e., “play an image copy forward”)
 - Data from other sources
 - IMS
 - VSAM
 - DB2 for LUW
 - Oracle
 - SQL Server
 - And more – if you can get the data into a sequential file, the Loader can get it into the Accelerator

New use cases

Once upon a time...

- At first, the DB2 Analytics Accelerator was thought of as a great solution for vastly improving query performance for existing DB2 for z/OS-based data warehouse applications
- The Accelerator is still an excellent fit for that situation, but over time more and more use cases for the technology have emerged – a few of these will be highlighted on the following slides

Intelligent, real-time cross-sell suggestions



- A gasoline retailer uses the DB2 Analytics Accelerator to provide employees in the mini-marts at stations with cross-sell suggestions at the point of sale
 - Items entered into a point-of-sale terminal at check-out time are input to the cross-sell application, and the Accelerator provides suggested complementary products in seconds (in time to make a difference)
 - The result: a better shopper experience, and increased retail sales
 - The cross-sell application, by the way, was built using IBM Cognos BI and IBM SPSS predictive analytics software

Extending the power of analytics to more users

- Using the DB2 Analytics Accelerator, a company in the banking industry extended analytics capabilities to many more employees in their enterprise
 - An example of “operational analytics” – users getting rapid-response query access to operational data records
- Over 1,000 of the bank’s business users utilize the Accelerator-based analytics system to get fast access to vital insights
 - The result: more effective development of new products, services and strategies that are growing the business



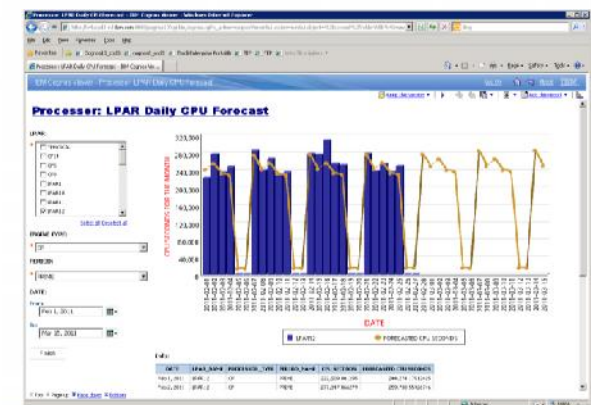
Providing online access to formerly offline, archived data

- An insurance organization was archiving historical data pertaining to claim events on tape
- The DB2 Analytics Accelerator provided a means whereby this data could be made available for online query access in a cost-effective manner
 - Little load on the front-end z Systems server
 - Low-cost data storage on the Analytics Accelerator
- Result: not only online query access to vast amounts of historical data for more informed decision-making, but with consistently good response time for both small-scale and data-intensive queries



High-performance, dynamic, predictive IT capacity management

- An insurance company needed to improve their enterprise IT capacity management capabilities
- The solution: IBM Capacity Management Analytics, combining IBM's Tivoli Decision Support for z/OS, Cognos BI, and SPSS Modeler to provide:
 - Graphical representation of z Systems capacity and utilization information
 - Flexible, dynamic reporting options
 - Predictive analysis of z Systems capacity requirements
- Building the solution on the DB2 Analytics Accelerator provided much-improved capacity management capabilities, with major z Systems CPU savings versus their old application
 - CPU times on the front-end DB2 system for capacity analysis queries decreased 70-99%



Predictive fraud analysis

- A credit card issuer previously detected fraud only at the micro level (i.e., at the individual card level)
- A new DB2 Analytics Accelerator-based application enabled analysis of much larger sets of data, providing the capability to anticipate fraud at the merchant and store level
- Result: not just detecting, but preventing incidents of credit card fraud
- Millions of dollars in savings:
 - Reduced fraud losses
 - Reduced card deactivation costs
 - Reduced call center costs
 - Savings, as well, for the company's partners (processors and merchants)



Value-add from a dynamic, customer-facing reporting system

- A company that manages 401k retirement plans had been providing participation data to plan sponsors via “canned” reports – an approach that threatened to undermine the firm’s competitive standing
- The solution: a new DB2 Analytics Accelerator-based application that provides sponsors with ad-hoc query access to plan participation metrics
- Results:
 - Improved sponsor satisfaction
 - Enhanced value for plan participants
 - A simplified technology stack for the company



Fast, customized access to SAP ERP data

- An oil and gas company is utilizing a DB2 Analytics Accelerator to provide users with dynamic, high-performance access to data generated by their SAP ERP application
- Results:
 - Significantly reduced query run times versus their former approach to accessing SAP ERP data
 - IT staff time formerly consumed in tuning queries targeting SAP ERP data now freed up for higher-value work



Leverage the DB2 Analytics Accelerator with IMS data

Make better decisions faster



Large-volume reporting of combined IMS and DB2 assets

Better understand your customers



Leverage full breadth of transactional data for analytics

Trust your data



Ensure consistency of data relationships between IMS and DB2



Get Started Today!

Technical Whitepaper and “how-to” guide available [here](#)

Additional Resources for IMS clients

- **White paper** on how to use the IBM DB2 Analytics Accelerator with your IMS data:
 - http://ibm.biz/accelerate_insights_ims_transactional_data
- Technical Implementation Guide (“cook book”)
 - https://ibm.biz/ims_idaa_technical_implementation
- YouTube – video on the IMS and IBM DB2 Analytics Accelerator solution
 - https://ibm.biz/demo_ims_idaa

And this is just scratching the surface...

- Static as well as dynamic query support; low-cost, online access to historical data; better-performing incremental data update; the ability to load data from non-DB2 sources directly into an Accelerator; etc.
 - It all adds up to create new opportunities to exploit DB2 Analytics Accelerator technology to lower costs, grow revenues, improve customer experiences, and enhance organizational agility
 - How could your organization put the Accelerator to work?



Thanks for your time!