THE FILLMORE GROUP
Relational Database Solutions

Hybrid Transaction/Analytic Processing (HTAP)
The Fillmore Group – June 2015

A Premier IBM Business Partner

Premier Business Partner IBM

History

The Fillmore Group, Inc.

- Founded in the US in Maryland, 1987
- IBM Business Partner since 1989
- Delivering IBM Education since 1994
- DB2 Gold Consultant since 1998
- IBM Champions since 2009

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

The Fillmore Group, Inc.

- DB2 Technical Support and Consulting
- IBM Training Partner with Global Training Partner Arrow ECS
- IBM Information Management Software Reseller

IBM CHAMPION

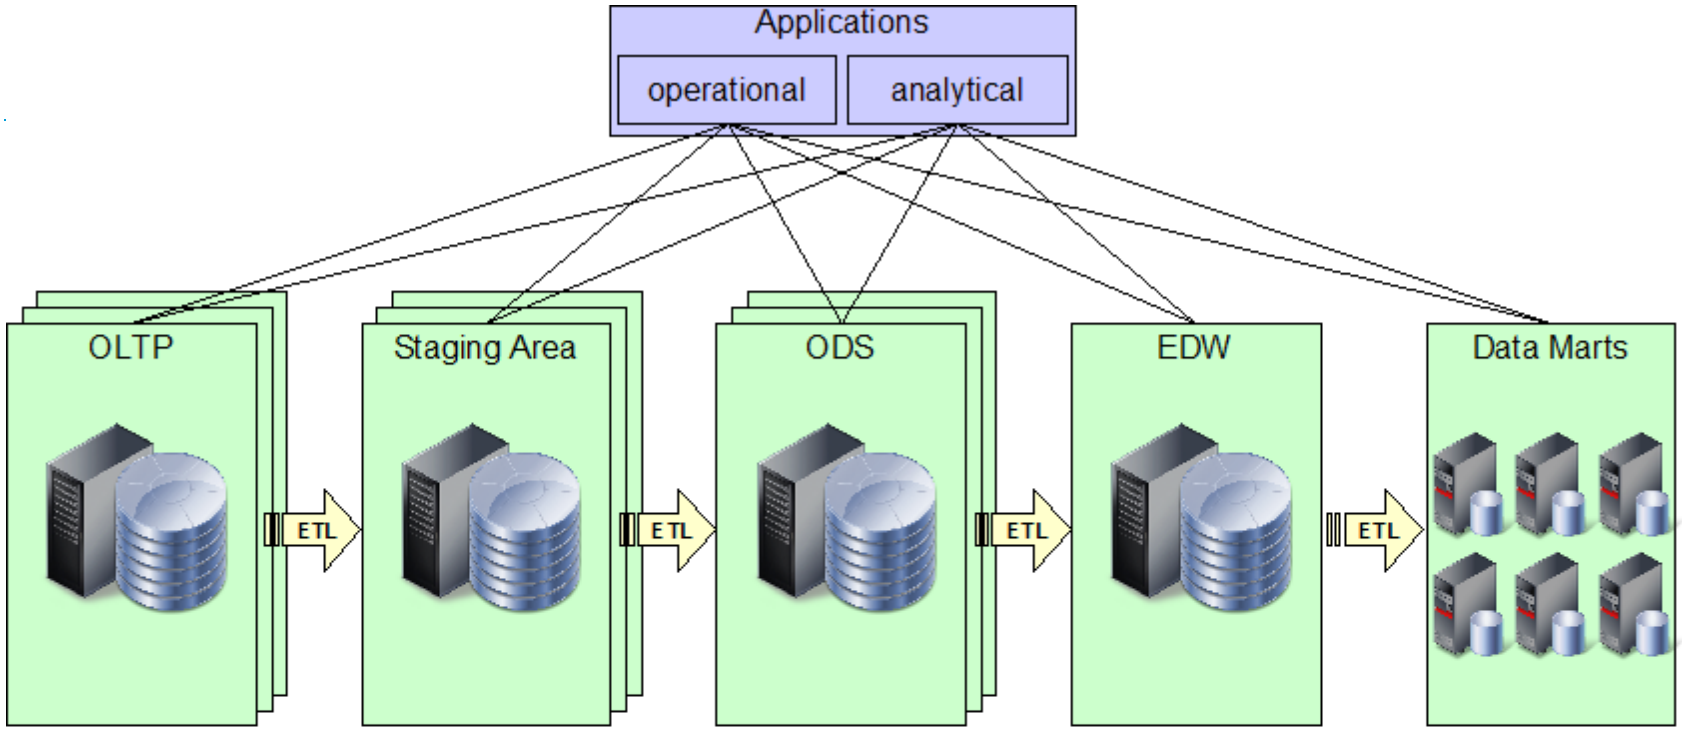

Negative ramifications:

- Complexity
  - both in systems management and in applications
- Difficulties in supporting real time analytics
- Inability to match ever more demanding SLA requirements
- High total cost of ownership

Historical reasons:

- Different access patterns
  - impact on performance
- EDW as the data integration hub
  - again, impact on performance
- Different life-cycle characteristics
  - and again, impact on performance
- Different Service Level Agreements (SLA)
  - Lack of broadly available workload management capabilities
  - Choice of lower cost-of-acquisition offerings

Hybrid Transaction/Analytic Processing

Benefits

- Uniform policies and procedures for security, HA, DR, monitoring, same tools, same skills, ...
- Efficient data movement within the system, often not involving network (ELT vs. ETL)
- Uniform access to any data for types of applications
- Opportunity to remove, i.e. consolidate some of the layers, ultimately leading to a single database

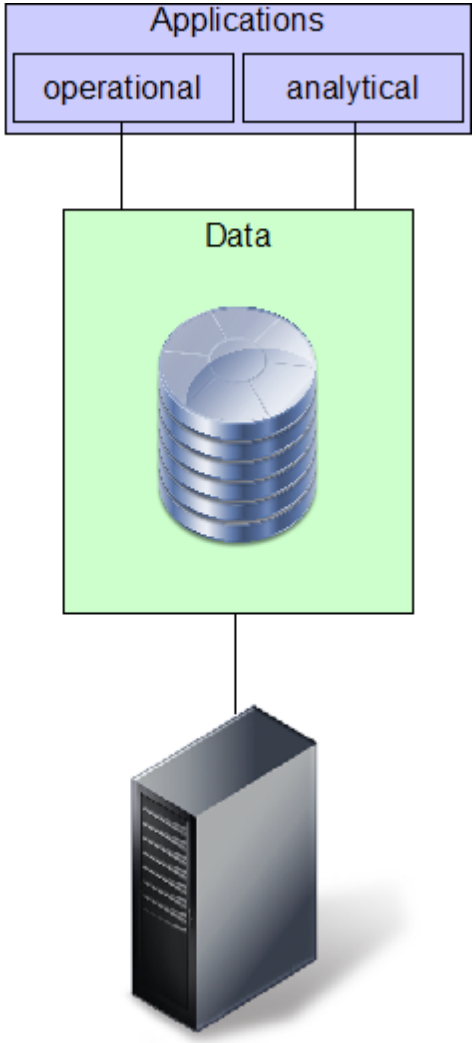

Prepayment analytics reduce cost

- Clients choose to detect <100% of transactions due to performance of sub-optimal analytics environment
- Organizations use rules-only approach to analytics, not deriving added business value of predictive
- Business-run analysis based on older, ETL data without benefit of most recent transaction data

Post-payment collection		Pre-payment predictive analytics	
Delta cost of collection per claim	**$25**	**$0**	Delta cost of collection per claim
Success rate	**50%**	**90%**	Success rate
Amount RECOVERED	**$118,750**	**$213,750**	Amount NOT PAID
Total cost to collect	**$59,375**	**$0**	Total cost to collect
Total amount RECOVERED	**$59,375**	**$213,750**	Total amount NOT PAID

Potential savings per day:	$154,375
Potential savings per year:	**$56,346,875**

Eliminating ETL reduces IT expense

Operational data

Analytical data

Multiple copies of data
Redundancy, inconsistency, complexity and cost

Significant compute power
Increases cost

Transaction and analytics isolation
Increases time to insight

- 1TB ETL per day
  (initial copy plus three derivatives)
  costs > *$8 million over 4 years*
- Clients report using 30% - 40% of System z cycles to move data

HTAP Infrastructure for DB2

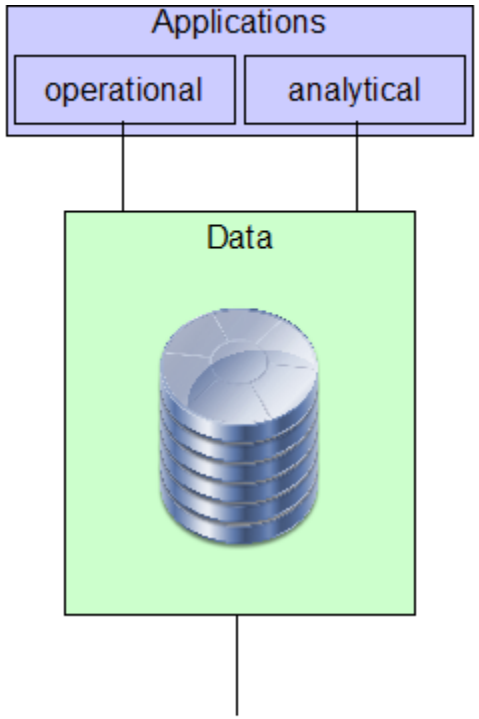

Benefits

- Uniform policies and procedures for security, HA, DR, monitoring, same tools, same skills, ...
- Efficient data movement within the system, often not involving network (ELT vs. ETL)
- Uniform access to any data for types of applications
- Opportunity to remove, i.e. consolidate some of the layers, ultimately leading to a single database

Challenges

- ✔ Mixed workload management capabilities
- ✔ Ensuring continuous availability, security and reliability
- ✔ Providing seamless scale-up and scale-out
- ? Providing universal processing capabilities to deliver best performance for both transactional and analytical workloads **without the need for excessive tuning**

Approaches

- Large RAM
  - 'In-memory' databases
- Massively parallel processing
  - Large number of sockets, cores, servers
  - Vector processing
- Hardware acceleration through special purpose processors
  - FPGA, GPU, ...
- Columnar stores
- Appliances

Building on proven technology base

- DB2 (both z/OS and LUW) already provide superior technology to address most of the challenges
- The remaining challenge is addressed by adding special purpose processing component for analytical workloads
  - DB2 for z/OS: **IBM DB2 Analytics Accelerator**
  - DB2 for zLinux: **BLU**

IBM DB2 Analytics Accelerator (IDAA): OLTP and Netezza *hybrid*

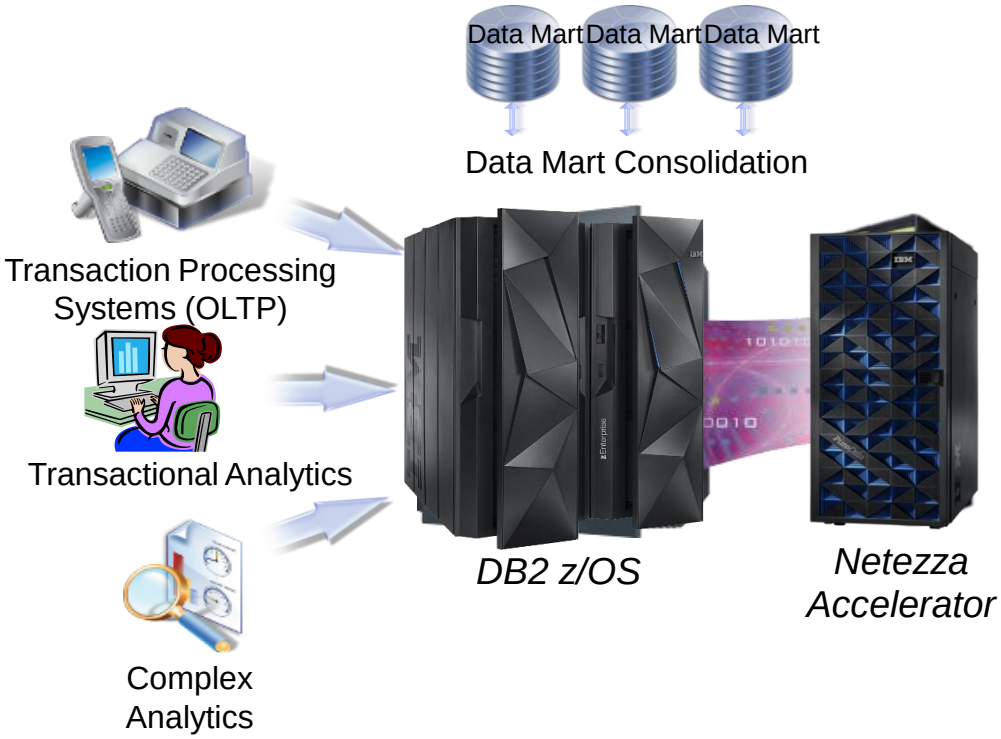

IBM DB2 Analytics Accelerator

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 business insights into operational processes to drive winning strategies. It automatically accelerates select queries, with unprecedented response times and negligible MIPS impact.

How is it different?

- Performance: unprecedented response times to enable 'train of thought' analyses frequently blocked by poor query performance.
- Integration: deep integration with DB2 for z/OS 10 and 11 provides transparency to all applications.
- Self-managed workloads: queries are automatically executed in the most efficient location
- Transparency: applications connected to DB2 are entirely unaware of the Accelerator
- Simplified administration: appliance hands-free operations, eliminating most database tuning tasks

Breakthrough Technology Enabling New Opportunities

THE FILLMORE GROUP
Relational Database Solutions

Loading dock to production ready in 2 days

				DB2 Only		DB2 with IDAA		Times Faster
Query	Total Rows Reviewed	Total Qualifying Rows	Total Rows Returned	Hours	Sec(s)	Hours	Sec(s)	
Query 1	591,941,065	2,813,571	853,320	2:39	9,540	0.0	5	1,908
Query 2	591,941,065	2,813,571	585,780	2:16	8,220	0.0	5	1,644
Query 3	813,343,052	8,260,214	274	1:16	4,560	0.0	6	760
Query 4	283,105,125	2,813,571	601,197	1:08	4,080	0.0	5	816
Query 5	591,941,089	3,422,765	508	0:57	4,080	0.0	70	58
Query 6	813,343,052	4,290,648	165	0:53	3,180	0.0	6	530
Query 7	591,941,065	361,521	58,236	0:51	3,120	0.0	4	780
Query 8	813,343,052	3,425,292	724	0:44	2,640	0.0	2	1,320
Query 9	813,343,052	4,130,107	137	0:42	2,520	0.1	193	13

- ***IBM DB2 Analytics Accelerator (N1001-010)***
  - ***Production ready - 1 person, 2 days***
- ***Table Acceleration Setup in 2 Hours***
  - ***DB2 "Add Accelerator"***
  - ***Choose a Table for "Acceleration"***
  - ***Load the Table (DB2 Loads Data to the Accelerator)***
  - ***Knowledge Transfer***
  - ***Query Comparisons***

- ***Initial Load Performance***
  - ***400 GB Loaded in 29 Minutes***
  - ***570 Million Rows (Actual: Loaded 800 GB to 1.3 TB per hour)***
- ***Extreme Query Acceleration - 1908x faster***
  - ***2 Hours 39 minutes to 5 Seconds***
- ***CPU Utilization Reduction***
  - ***Accelerated queries had negligible CP impact***

"We had this up and running in days with queries that ran over 1000 times faster"

Why do you care?

- Business critical analytic applications **demand low latency, high qualities of service and performance**
- The issue: spreading analytic components across multiple platforms can **increase data latency, cost, complexity and governance risk**
- Keeping analytic components closer to the source data **improves data governance while minimizing data latency, cost and complexity**

Reduce data latency
by up to 99%

Use cases

A large Brazilian bank delivers IT at the speed of business by eliminating critical reporting latency

The bank is using DB2 Analytics Accelerator to drive customer insight from operational data. Processes that previously took 24 hours for ETL and 11 hours more for reporting, now take 1 hour and 26 seconds.

THE FILLMORE GROUP
Relational Database Solutions

Run queries
up to 2000x
faster

Use cases

A large European convenience store chain is doing something they could never do before, increasing retail sales nearly 5% through reduced analytic query response times (99.8 % faster) on OLTP content

"The store employee enters what the customer is purchasing, and with the DB2 Analytics Accelerator appliance, the Cognos BI and SPSS tools deliver information on complementary products in seconds."
--A Chief Information officer--

Use cases

95% savings in host disk space

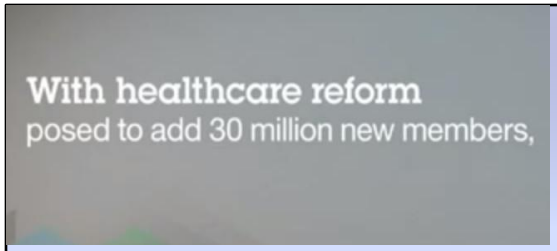

A large healthcare company is now **focused on business needs not technical constraints,** positioned to expand their membership and provide insight faster without impacting existing applications and infrastructure

*"…it means our **queries run dramatically faster"***

*"With the aging population, we expect a huge influx of data, so the cost of storing data is significant. By keeping data in the appliance, we expect **substantial storage cost savings**."*

Systems Engineering Manager

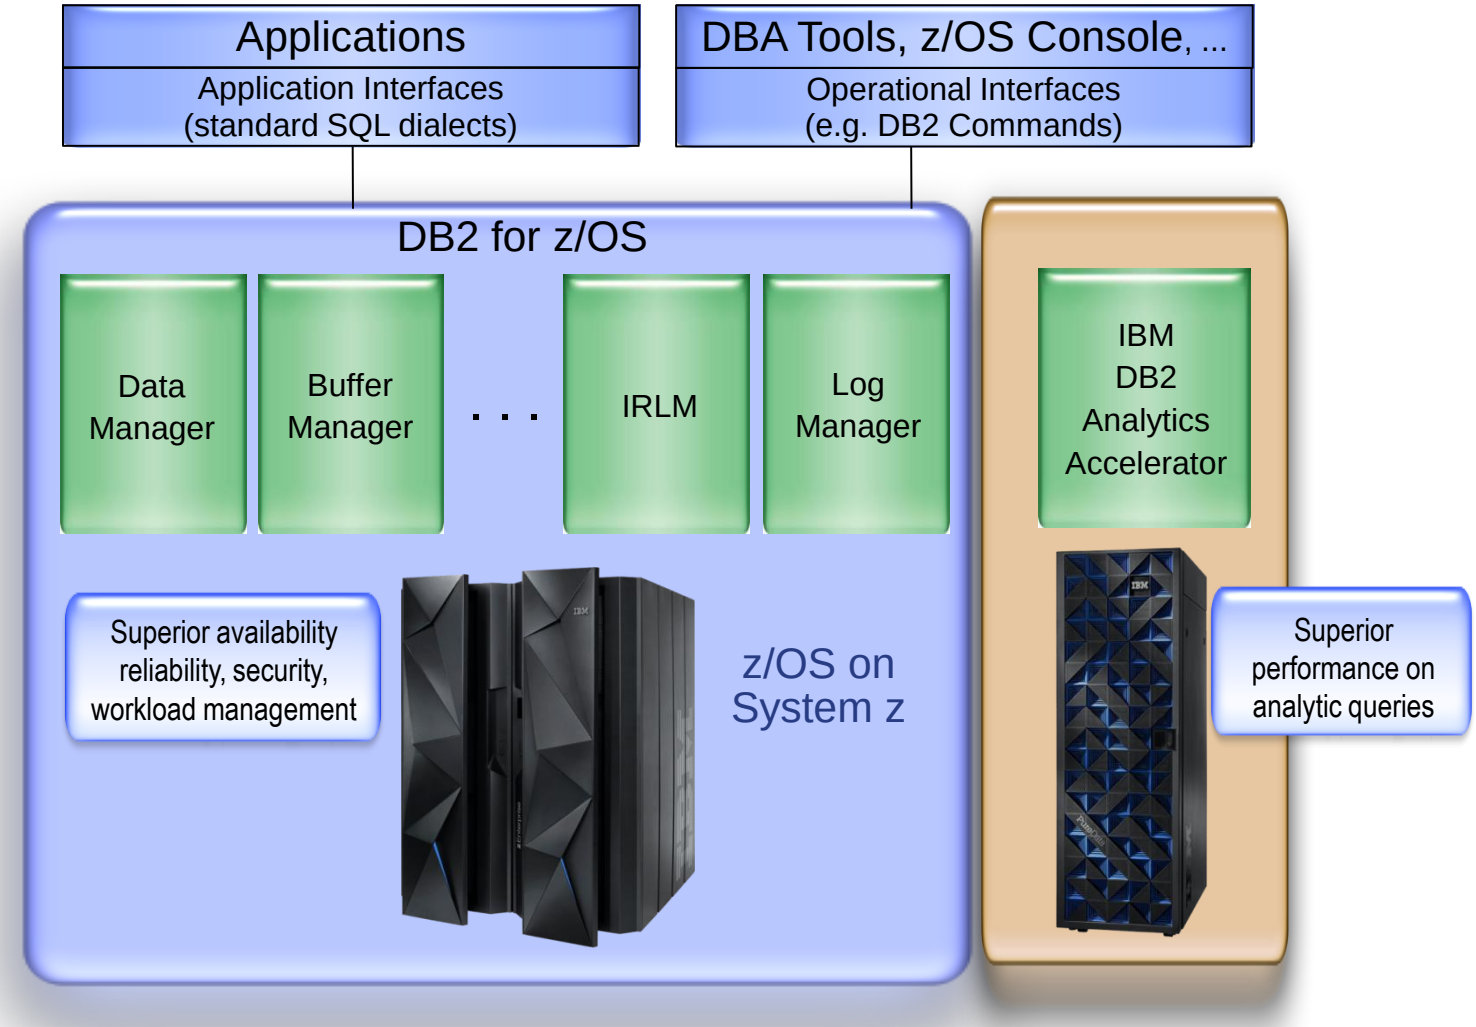
Applications
Application Interfaces
(standard SQL dialects)
DBA Tools, z/OS Console, ...
Operational Interfaces
(e.g. DB2 Commands)
DB2 for z/OS
Data Manager
Buffer Manager
. . .
IRLM
Log Manager
Superior availability reliability, security, workload management
z/OS on System z
IBM DB2 Analytics Accelerator
Superior performance on analytic queries

How it works

- Access to data in terms of authorization and privileges (security aspects) is controlled by DB2 and z/OS (Security Server)
- Uses DB2 for z/OS for updates, logging, fast single record look-ups
- DB2 for z/OS does backup and recovery
- DB2 for z/OS remains the system of record
- Management and monitoring of the Accelerator is via System z and DB2 for z/OS
- There is no external communication to the IBM DB2 Analytics Accelerator beyond DB2 for z/OS

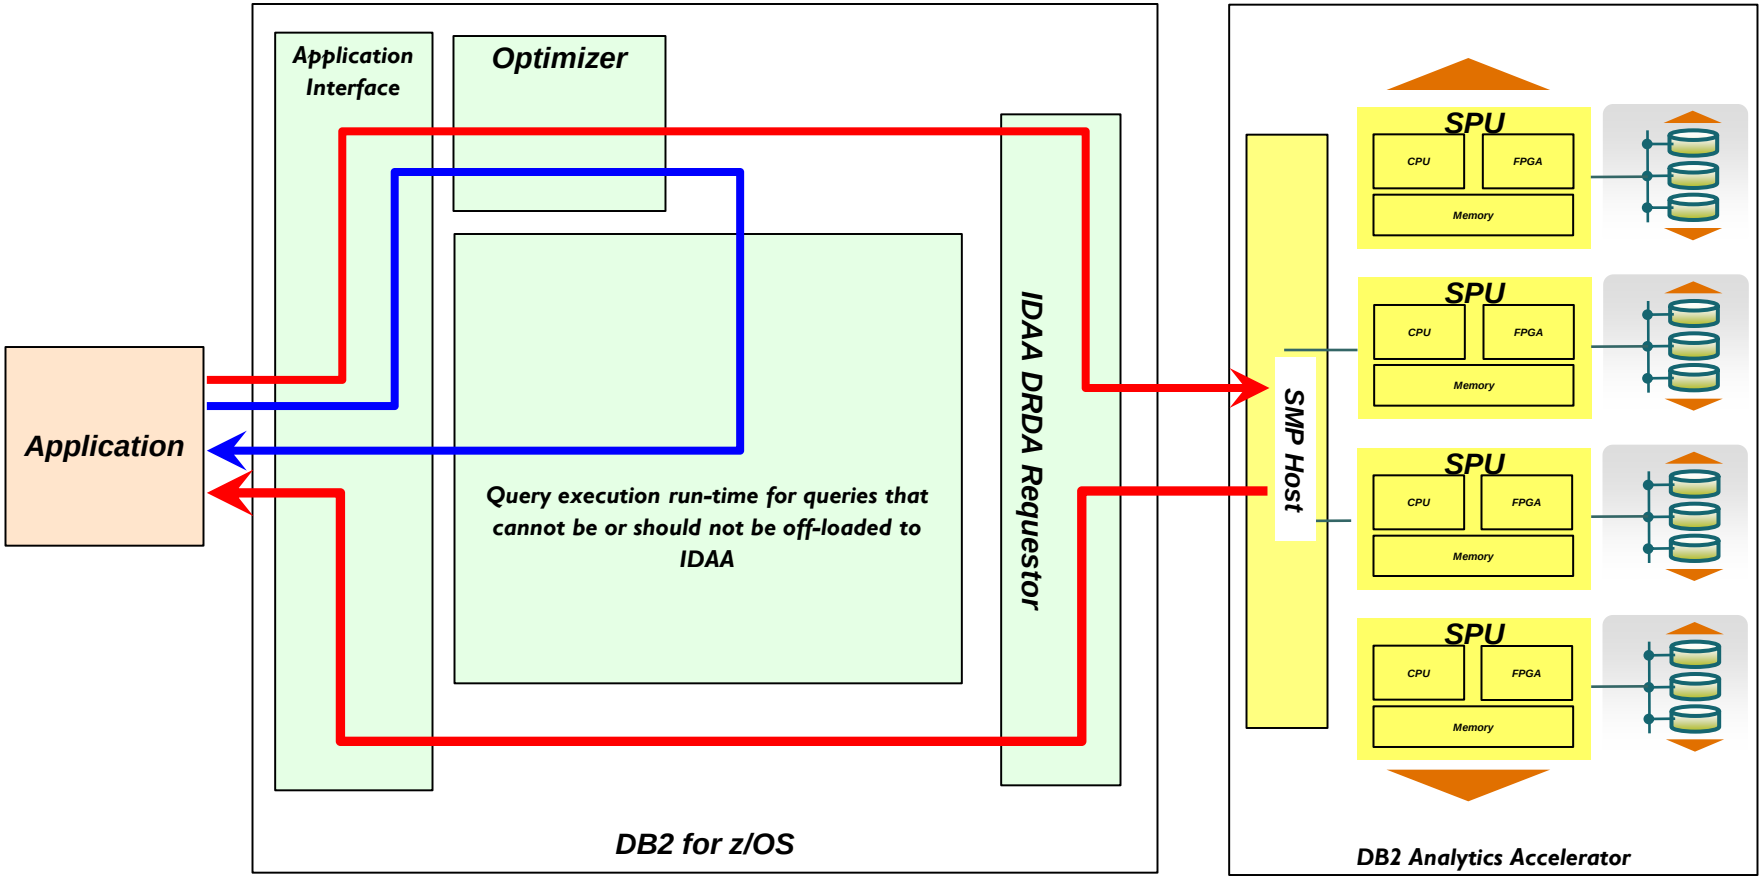

Queries executed without DB2 Analytics Accelerator

Queries executed with DB2 Analytics Accelerator

Premier Business Partner IBM

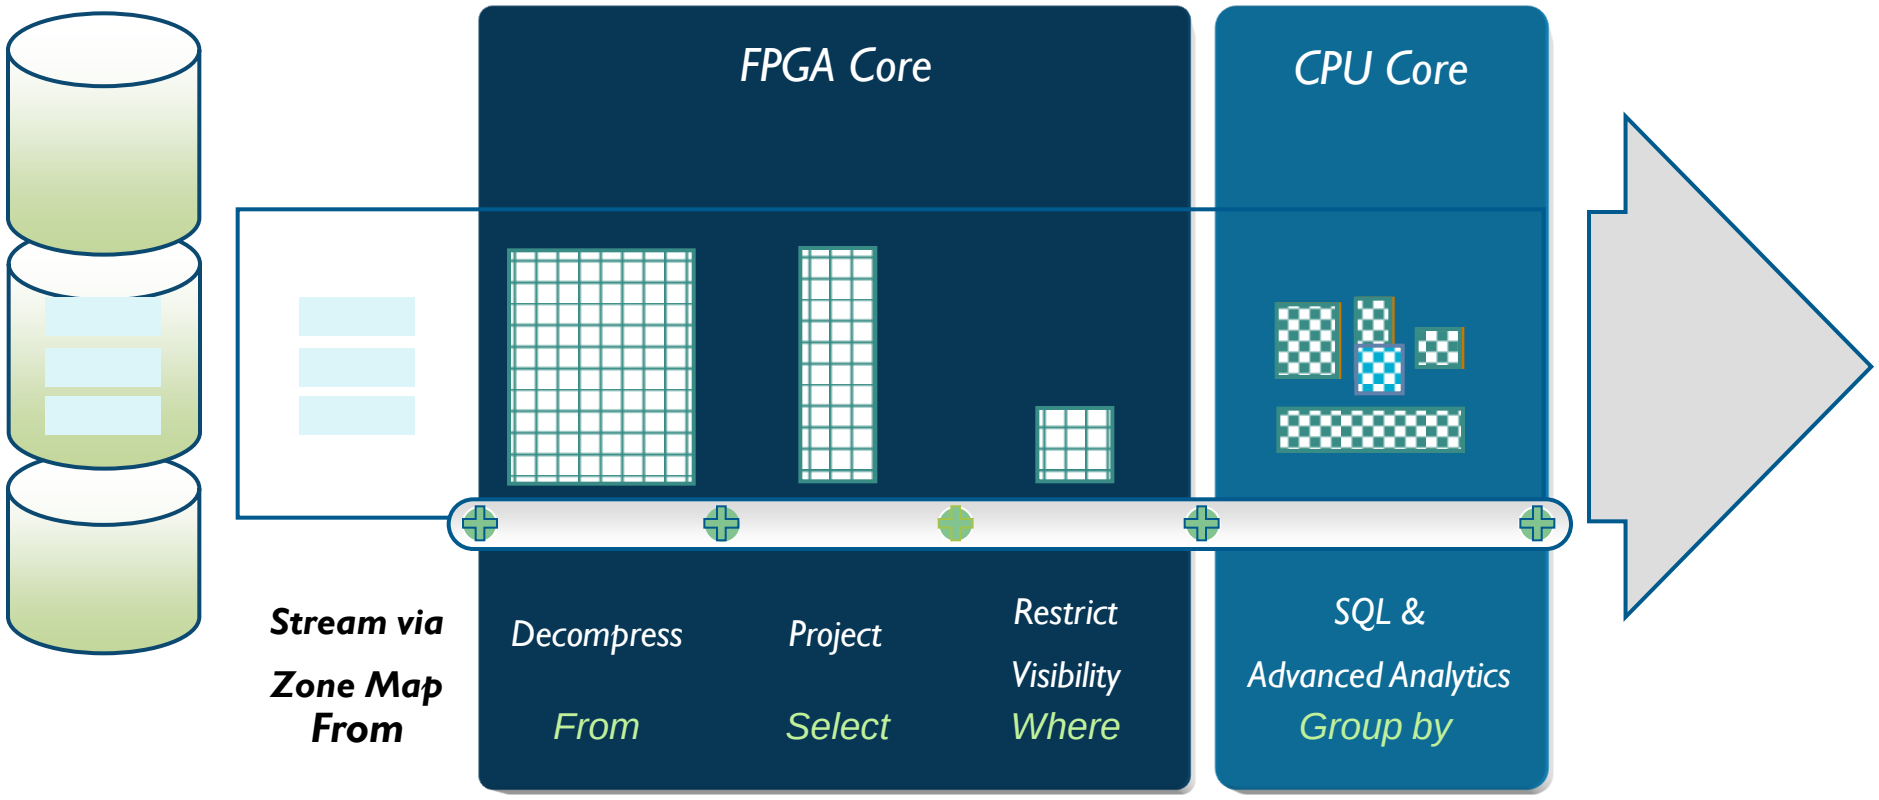

Select State, Age, Gender, count(*) From MultiBillionRowCustomerTable Where BirthDate < '01/01/1960' And State in ('FL', 'GA', 'SC', 'NC') Group by State, Age, Gender Order by State, Age, Gender

N3001 Model	001	002	005	010	020	040	080
Cabinets	n/a	¼	½	1	2	4	8
S-Blades	n/a	2	4	7	14	28	56
SPU CPU Cores	40	40	80	140	280	560	1120
SPU FPGA Cores	n/a	32	64	112	224	448	896
Capacity in TB (user data space)	4	8	24	48	96	192	384
Effective Capacity in TB (assuming 4x compression)	16	32	96	192	384	768	1536

N2002 Model	002	005	010	020	040
Cabinets	¼	½	1	2	4
S-Blades	2	4	7	14	28
SPU CPU Cores	32	64	112	224	448
SPU FPGA Cores	32	64	112	224	448
Capacity in TB (user data space)	8	24	48	96	192
Effective Capacity in TB (assuming 4x compression)	32	96	192	384	768

High Performance Storage Saver (HPSS)

- Historical and archival data need only reside on the Accelerator
- Saves DB2 for z/OS storage costs
- Provides cost-effective means to retain data "online" for search and analysis
- Supports auditing and compliance
- Special Register: GET_ACCEL_ARCHIVE

Premier Business Partner IBM

Synchronization options	Use cases, characteristics and requirements
Full table refresh **The entire content of a database table is refreshed for accelerator processing**	▪ **Existing ETL process replaces entire table** ▪ **Multiple sources or complex transformations** ▪ **Smaller, un-partitioned tables** ▪ **Reporting based on consistent snapshot**
Table partition refresh **For a partitioned database table, selected partitions can be refreshed for accelerator processing**	▪ **Optimization for partitioned warehouse tables, typically appending changes "at the end"** ▪ **More efficient than full table refresh for larger tables** ▪ **Reporting based on consistent snapshot**
Incremental Update **Log-based capturing of changes and propagation to IBM DB2 Analytics Accelerator with low latency (typically 1 minute)**	▪ **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**

Value Proposition

- Single platform, single API for OLTP and analytics
- Reduce
  - z/OS CPU utilization
  - Analytics latency
  - Complexity risk
  - Integration costs
  - Storage costs for archival and historical data
- Increase
  - Reliability, Availability, Serviceability

Flexible Deployment 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

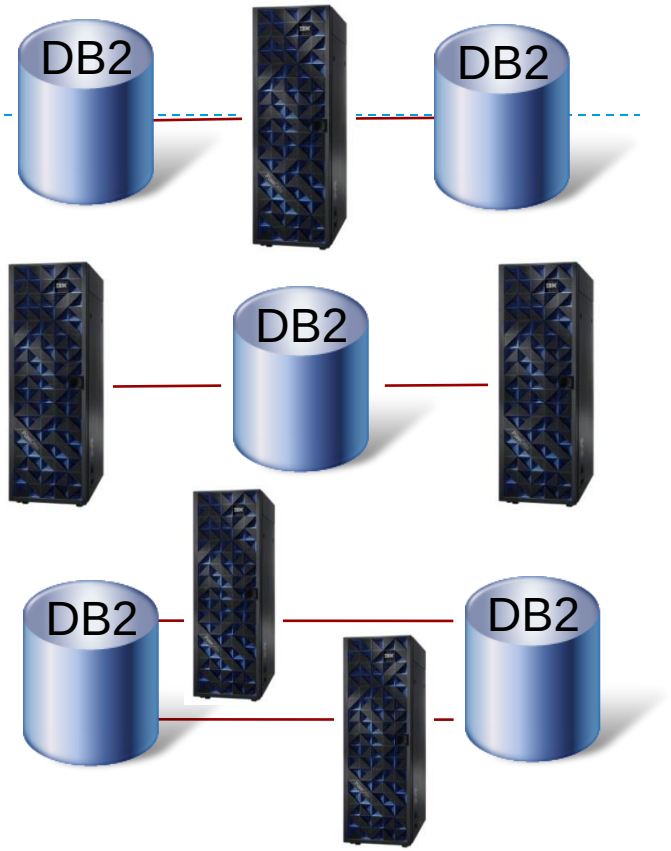

Better utilization of Accelerator resources
Scalability
High availability
Multiple options to deploy Dev/Test/QA

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

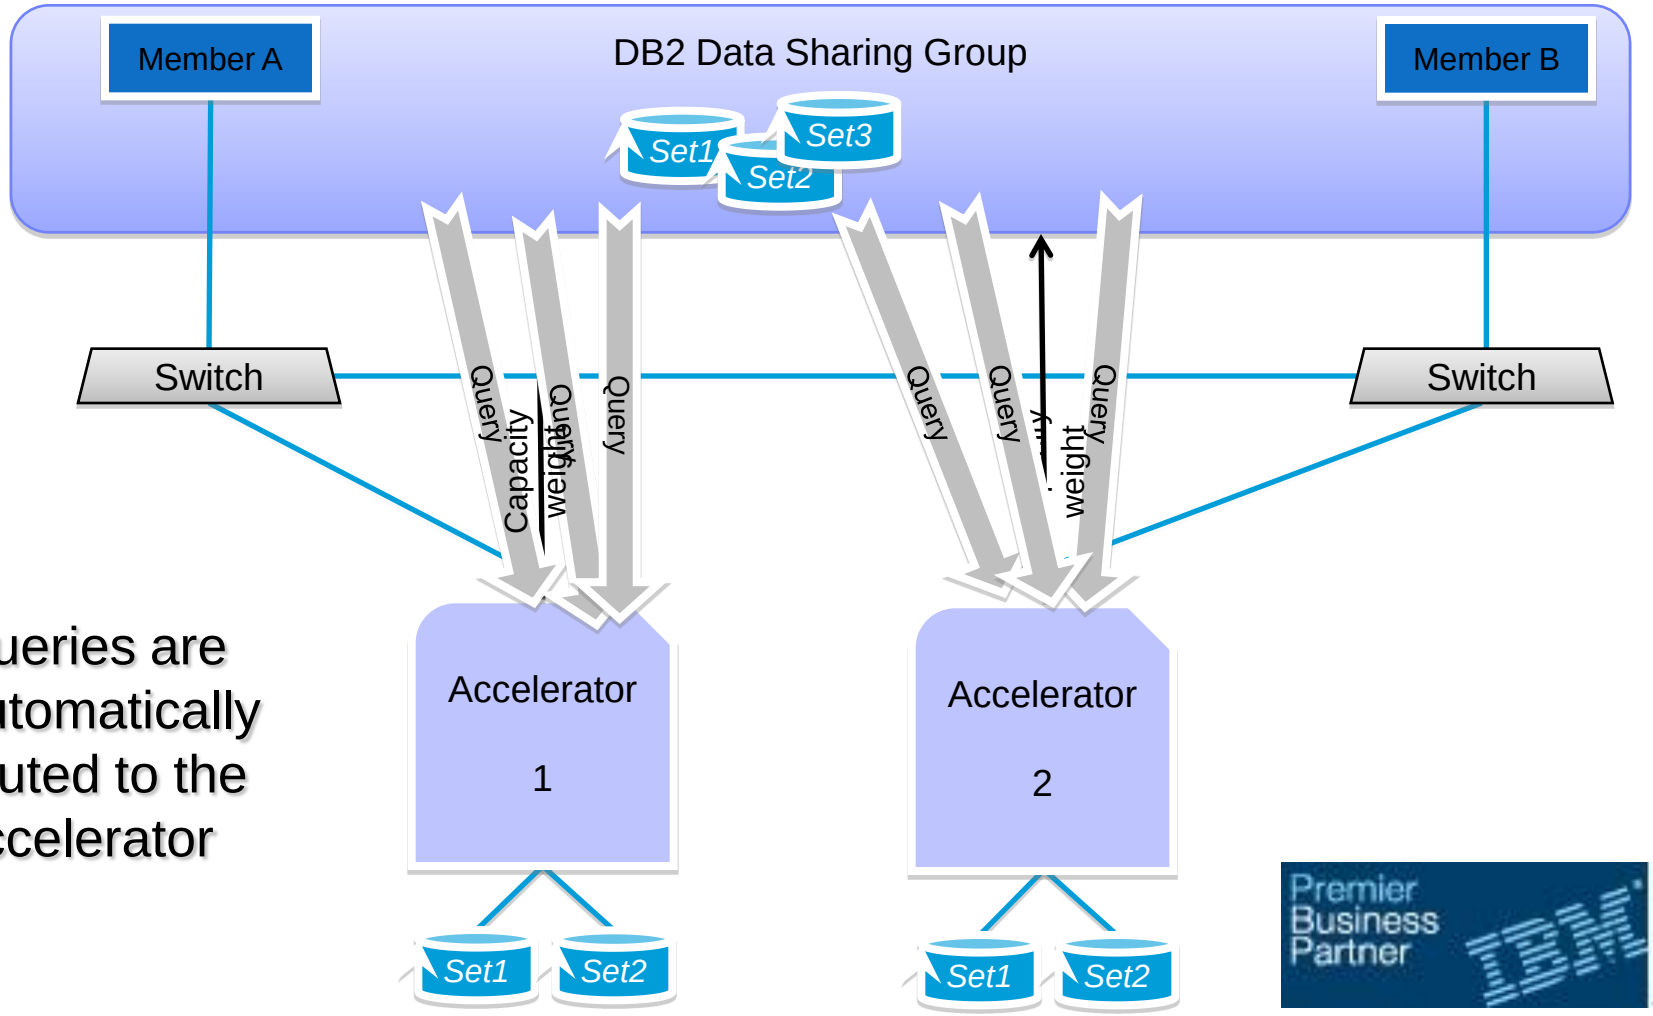

Queries are automatically routed to the accelerator

DB2 10.5 BLU Acceleration

- Memory optimized
  - In-memory ***columnar*** processing
  - Dynamic data movement from storage (no LRU)
- Actionable Compression
  - Patented ***compression*** technique that preserves order so that the data can be used without decompressing (column cardinality)
- Parallel Vector Processing
  - Multi-core and Single Instruction Multiple Data (SIMD) parallelism
- Data Skipping
  - Skips unnecessary processing of irrelevant data

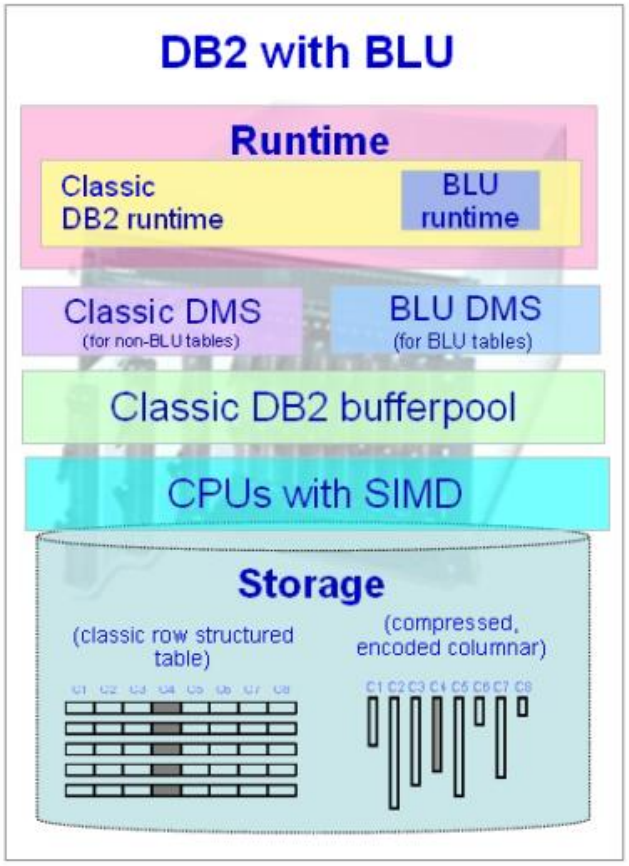

Premier Business Partner IBM

DB2 10.5 BLU Acceleration

- DB2 10.5 BLU Acceleration is a ***hybrid*** that supports mixed OLTP and ***analytic*** workloads
- Set DB2_WORKLOAD registry variable to ANALYTICS
- Column-organized tables will be the default table type
- Sets default page (32KB) and extent size (4) appropriate for analytics
- Data is always automatically compressed - no options
- For mixed table types can define tables as **`ORGANIZE BY COLUMN`** or **`ROW`**
- Utility to convert tables from row-organized to column-organized (**`db2convert`** utility)

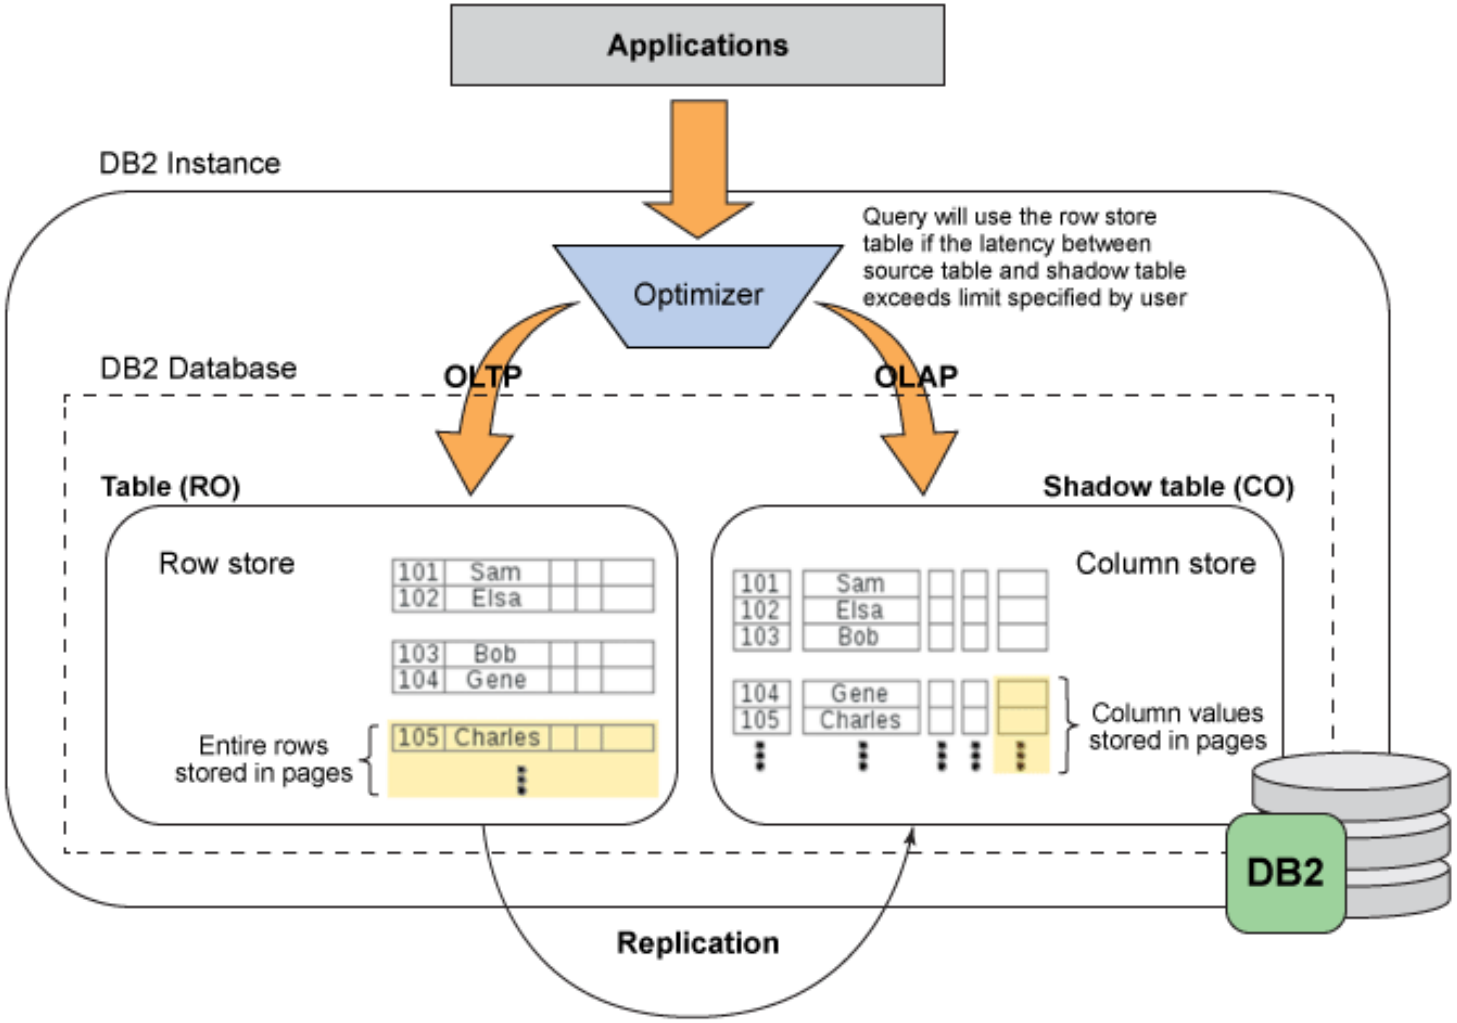
Applications
DB2 Instance
Optimizer
Query will use the row store table if the latency between source table and shadow table exceeds limit specified by user
DB2 Database
OLTP
OLAP
Table (RO)
Shadow table (CO)
Row store
101 Sam
102 Elsa
103 Bob
104 Gene
105 Charles
Entire rows stored in pages
Column store
101 Sam
102 Elsa
103 Bob
104 Gene
105 Charles
Column values stored in pages
DB2
Replication

Next Steps

- Hands-on Workshop
- Whiteboarding session
- Workload Assessment
  - Workload on DB2
  - Competitor workload (e.g. Teradata, MS SQL Server)
- Customer Value Engagement (CVE)
- Proof-of-concept (POC)

Business Use Case White Boarding Session

- Line of Business Sponsors
- Application Owners
- Information Architects
- 2-4 use cases

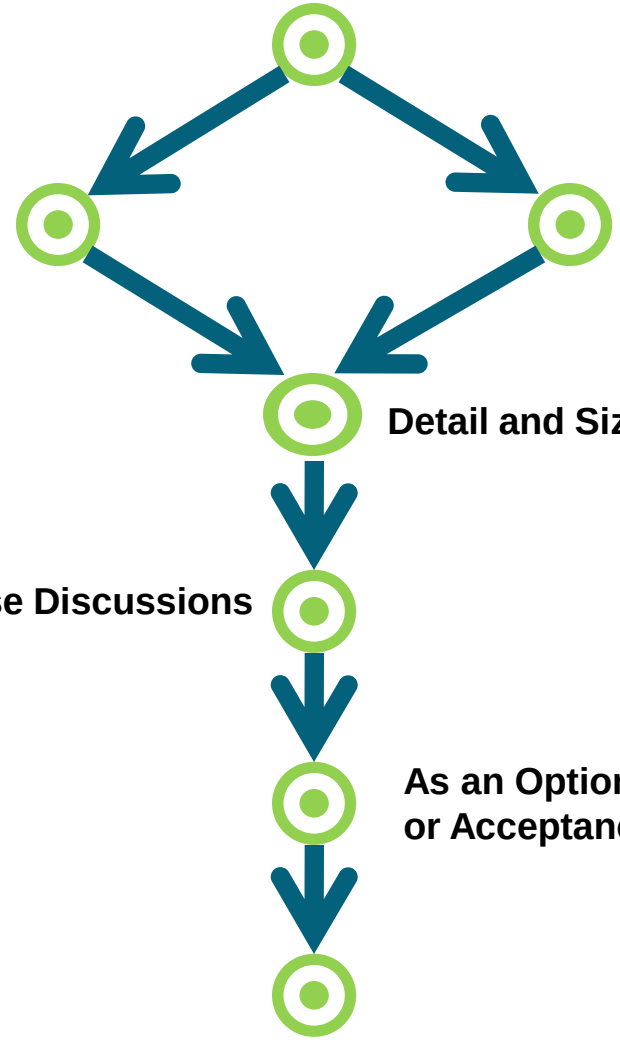

Hands on Workshop

- DBAs
- Developers
- Remote access to lab

Detail and Size Uses Cases

Begin Purchase Discussions

As an Optional Closing Tool, Introduce WLA or Acceptance-Based POC (TIBI)

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

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

With potential
17162 (98%)
16664 (95%)

Uncertain
153 (1%)
142 (1%)

No Potential
185 (1%)
179 (1%)

Other statements: 41686

Total Elapsed Time (s): 150653.49

With potential
138476.55 (92%)
15181.23 (10%)

Uncertain
10577.51 (7%)
4027.70 (3%)

No Potential
1599.42 (1%)
657.06 (0%)

Other statements: 94192592.96

Total CPU Time (s): 17440.15

With potential
15734.61 (90%)
3123.84 (18%)

Uncertain
1146.62 (7%)
56.09 (0%)

No Potential
558.91 (3%)
98.07 (1%)

Elapsed time potential

CPU time potential

Query	Elapsed Time (s)	CPU Time (s)	Executed	Eligibe / DB2 native	Rows exam./proc.	Table Sizes	IDAA Time	Comment	State
514365	3928.05	480.53	1		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)		(Prima SELEC DFIDF DTDIS ED_AN WHEN ORIG_ END_A DFIDF DFIDF DFIDF LOAN DFIDF DB_CF FROM PROD LOAN. OUTEI OUTEI WHER DFIDF INTFC DFIDR

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%

Proof-of-concept - Goals

- Manageability – Understand the tools and processes required to define, deploy and administer performance objects in the IDAA
- Functionality - Understand and witness the ability of IDAA solution to redirect queries to a workload optimized, appliance-like query accelerator based on IBM Netezza technology
- Performance and ease of migrating distributed databases
- Performance of accelerated queries
- A 2-3 week POC executed according to mutually defined plan

Attributions

- Dwaine Snow, IBM
- Jeff Feinsmith, IBM
- Patric Becker, IBM Boeblingen Lab
- Knut Stolze, IBM Boeblingen Lab
- Namik Hrle, IBM Fellow
- Ayesha Zaka, IBM Toronto Lab

Premier Business Partner IBM

Resources

- Redbooks
  - "Optimizing DB2 Queries with IBM DB2 Analytics Accelerator for z/OS" SG24-8005
  - "Hybrid Analytics Solution using IBM DB2 Analytics Accelerator for z/OS V3.1" SG24-8151
  - "Reliability and Performance with IBM DB2 Analytics Accelerator Version 4.1" SG24-8213
- www.thefillmoregroup.com/blog

Premier Business Partner IBM

Contacts

- Kim May
  - kim.may@thefillmoregroup.com
  - twitter.com/KimMayTFG
  - www.linkedin.com/pub/kim-may/4/462/84
- Frank Fillmore
  - frank.fillmore@thefillmoregroup.com
  - twitter.com/ffillmorejr
  - www.linkedin.com/pub/frank-fillmore/6/597/9a6/
  - tinyurl.com/ChannelDB2
  - Flipboard for iPad, iPhone, Android: "BigData"