THE FILLMORE GROUP
Relational Database Solutions

IBM Db2 Analytics Accelerator

The Fillmore Group – July 2018

A Premier IBM Business Partner

Premier Business Partner IBM

Agenda Day 1:

- Introductions
- The Fillmore Group
- IDAA Overview/Orientation
  - Use cases
  - How it works
  - What's new in IDAA v7?
    - "Brain transplant"
    - "Heart transplant"
  - Value proposition

Agenda Day 1 (cont):

- Granularity and level of control
  - zParms
  - Special Registers
- Lab 1 – The Basics
  - Data Studio GUI
  - Managing the IDAA
    - Adding, Loading, Enabling Tables
    - Running queries without/with acceleration
    - EXPLAIN Access Paths

Agenda Day 1/2:

- What did we learn? Questions? Problems?
- Lab 2 – Advanced
  - Manage IDAA through XML and JCL
  - Incremental Update
  - Static SQL
  - What did we learn? Questions? Problems?

Premier Business Partner IBM

Agenda Day 2:

- Lab 3 – Accelerator-only Tables (AoTs)
  - What did we learn? Questions? Problems?
- IDAA Strategy and Roadmap
- Next steps
  - Whiteboarding session
  - Workload Assessment (WLA)
  - Cost Benefit Analysis (CBA)
  - Proof-of-concept (POC)
- Survey and wrap-up

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

Services:

- Db2 for z/OS
  - Installation, development, healthchecks, performance tuning, backup & recovery
  - Version-to-version migrations (V9 – V10 – V11 – V12)
  - IBM Db2 Analytics Accelerator
  - Support CA (Broadcom), BMC, IBM administration tools
  - Db2 Connect, Replication, Federation

Premier Business Partner IBM

Is the mainframe dead?

"Mainframe dead? IT still sees it as robust, long-term solution, shows study" ComputerWeekly 2013

"Help!! My Manager Said We're Moving Off the Mainframe!"

Machine Learning Introduced on System z

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 v10, v11, and v12 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

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**

THE FILLMORE GROUP
Relational Database Solutions

Reduce data latency
by up to 99%

Use cases

A large South American 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.

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

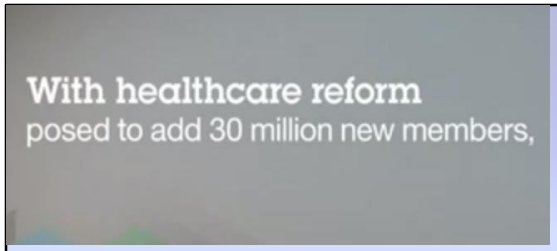

A large healthcare company is now **focused on business needs not technical constraints,** positioned to expand their membership and provide insight faster without impacting existing applications and infrastructure

*"…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

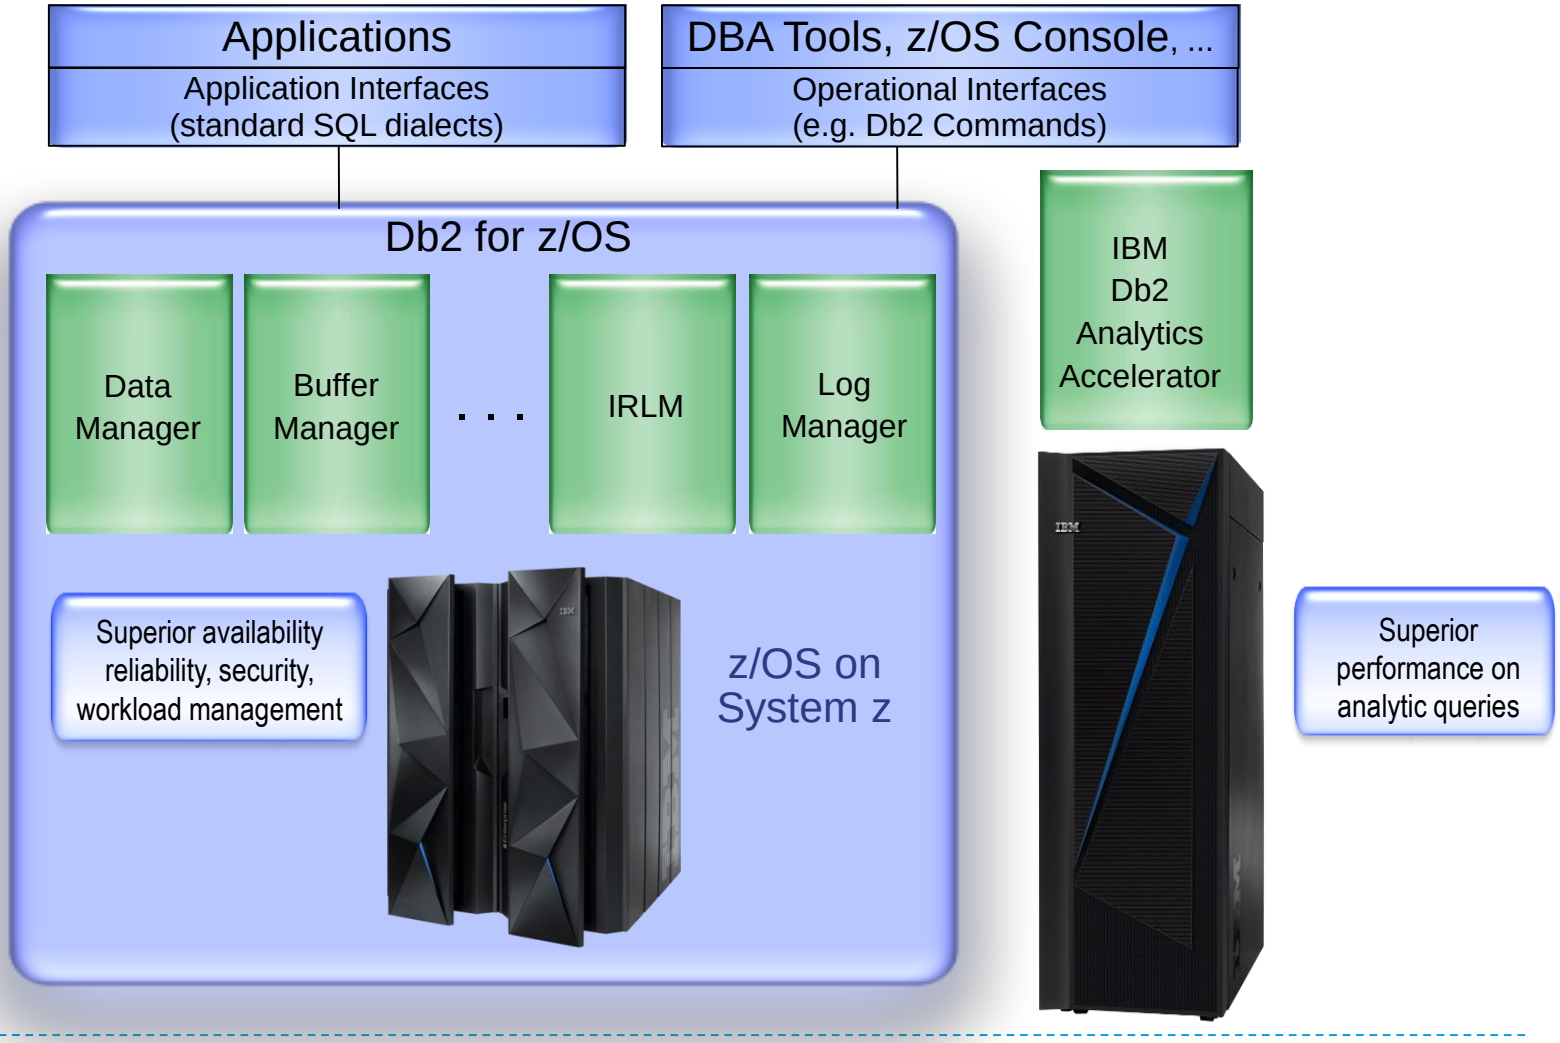
Applications
Application Interfaces
(standard SQL dialects)
DBA Tools, z/OS Console, ...
Operational Interfaces
(e.g. Db2 Commands)
Db2 for z/OS
Data Manager
Buffer Manager
. . .
IRLM
Log Manager
Superior availability reliability, security, workload management
z/OS on System z
IBM Db2 Analytics Accelerator
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

The hybrid computing platform on System z

- *Supports transaction processing and analytics workloads concurrently, efficiently and cost-effectively*
- *Delivers industry leading performance for mixed workloads*
- *The unique heterogeneous scale-out platform*
- *Superior availability, reliability and security*

Db2 Accelerator personalities

- *Turbo charged access path with hardware assisted early filtering*
- *Full-width index*
- *Specialty engine*
- *Archive*
- *Tablespace*
- *ETL/ELT and in-database analytics acceleration*
- *Integration hub: fast federated joins across heterogeneous sources*
- *Hybrid cloud*

What's new in IDAA v7? – "Brain transplant"

- IDAA v5: PostgreSQL
- IDAA v7: Common SQL Engine
  - Db2 Warehouse
  - Db2 for LUW with BLU Acceleration

Premier Business Partner IBM

THE FILLMORE GROUP
Relational Database Solutions

Columnar Everywhere

- Reduce I/O
- Increase data density in RAM
- Increase CPU efficiency

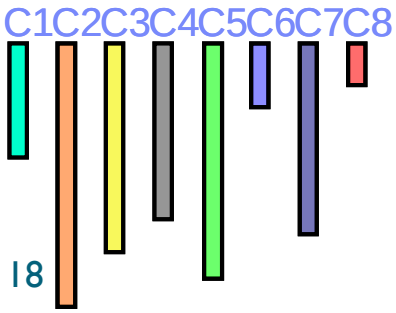

Skip Boring Data

- Queries skip uninteresting data
- Synopses on every column, automatically.
- *"Data Skipping"*

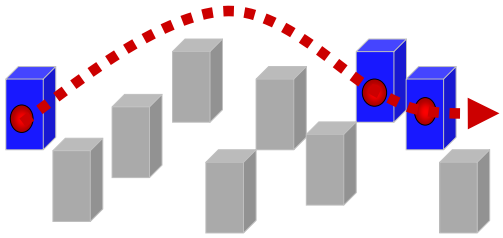

Rethink Memory

- Cache intelligently for analytics
- Predictive I/O with *"Dynamic List Prefetching"*
- Massive I/O reduction

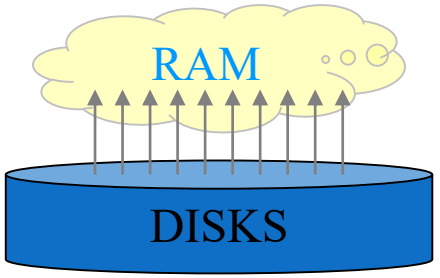

What's new in IDAA v7? – "Heart transplant"

- IDAA v5
  - Field Programmable Gate Arrays (FPGAs)
  - Commodity disks
- IDAA v7
  - POWER Chips
  - Solid State Disks

What's new?

- Common Db2 in-memory, columnar SQL engine with BLU acceleration and SQL compatibility improvements, for example:
  - Native support for the EBCDIC MBCS and GRAPHIC data types (instead of a conversion to UTF-8 as in earlier versions)
  - Acceleration of all types of correlated sub-queries (only small subset was offloaded in version 5), including table expressions with sideway references
  - Support for timestamps of precision 12 (these were truncated to precision 6 in earlier versions
  - Native support for the FOR BIT DATA subtype, for all types of table encoding (EBCDIC, UNICODE, ASCII). This was available for EBCDIC only in earlier versions).
  - Native support for the TIMESTAMP value 24:00:00 (this was mapped to 23:59:59 in earlier versions)

Premier Business Partner IBM

What's new (cont)?

- Improved support for tables in mixed encoding
- EBCDIC tables can be added to accelerator even if UNICODE tables are already present
- Better time synchronization through use of system heartbeat
- Improved accuracy for CURRENT_TIME, CURRENT_TIMESTAMP, and CURRENT_DATE special registers

- Smart Load, meaning that the SYSPROC.ACCEL_LOAD_TABLES stored procedures optimizes the workload and the degree of parallelism for load jobs.

The core of future Accelerator Deployments

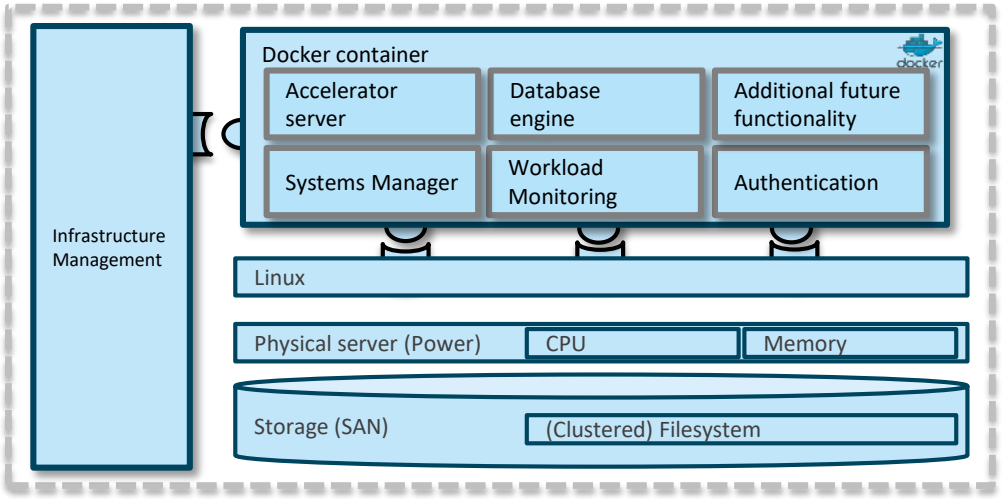

Docker is a widely adopted, open-source project that automates the deployment of applications inside software containers.

It allows to bundle and preinstall components in a docker image and then to launch the container from the image.

The accelerator docker container includes

- the accelerator server that establishes the connection between the Db2 for z/OS subsystem and manages all accelerator tasks
- a database engine
- other components required for high availability, monitoring and security

Outside of the container physical compute resources are needed

- a server that includes multi-core CPUs
- large memory
- shared filesystem to persist the data

On top of the physical hardware there is a docker supported Linux operating system that is just used to launch the docker container and manage the HW resources.

IBM Integrated Analytics System Configurations

IBM Power 8 S822L 24 core server 3.02GHz

	M4001-003 1/3 Rack	M4001-006 2/3 Rack	M4001-010 Full Rack
Servers	3	5	7
Cores	72	120	168
Memory	1.5 TB	2.5 TB	3.5 TB
Flash Storage Capacity	16 TB	32 TB	48 TB

	M4002-003 1/3 Rack	M4002-006 2/3 Rack	M4002-010 Full Rack	M4002-020 Double Rack	M4002-020 Quad Rack
Servers	3	5	7	13	25
Cores	72	120	168	312	600
Memory	1.5 TB	2.5 TB	3.5 TB	6.5 TB	12.5 TB
Flash Storage Capacity	27 TB	54 TB	81 TB	162 TB	324 TB

Comparison

PDA N3001 (Netezza Technology)

Integrated Analytics System M4001

Accelerator on IBM Z

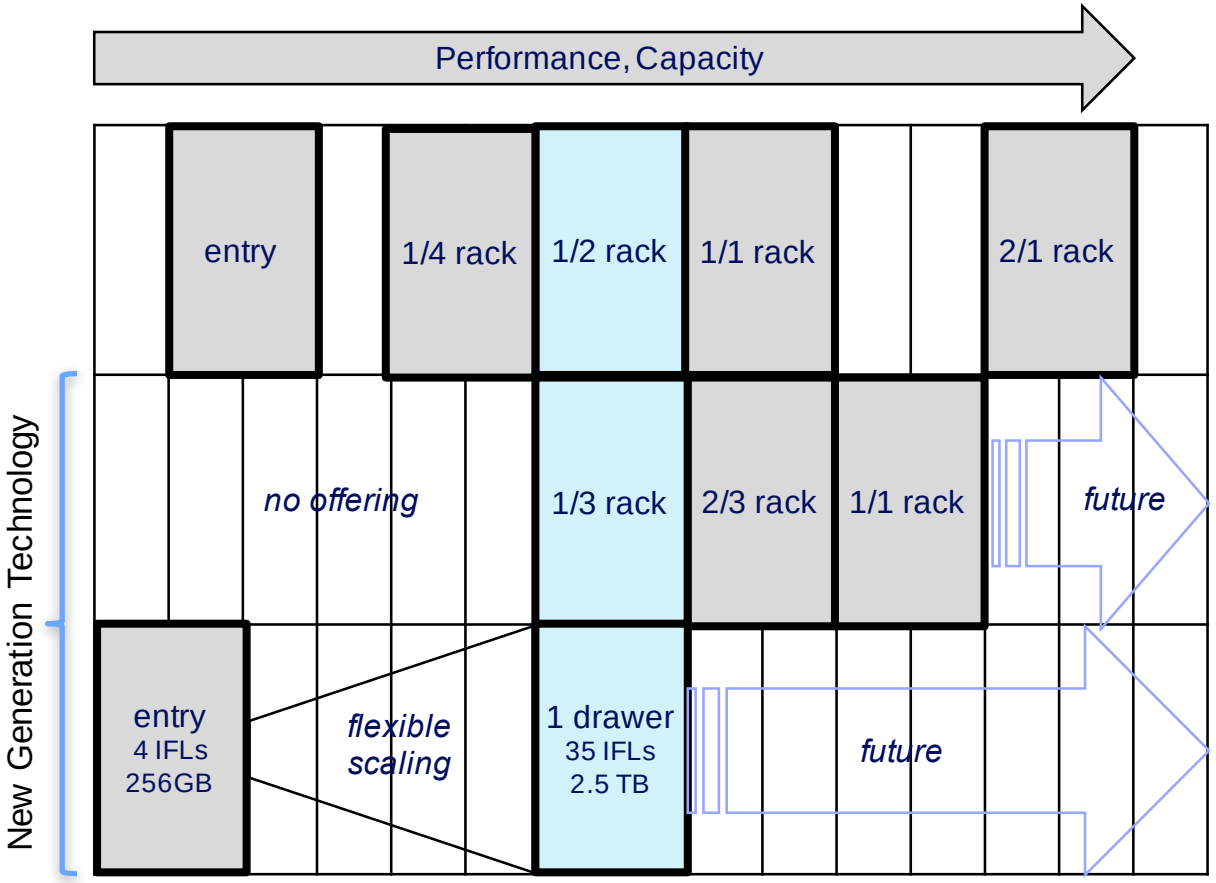

High Availability – Hardware Elements

Robust Hardware Elements

- Power8
- **IBM Flash System 900**
  - Completely resilient storage subsystem for the appliance
  - Storage array is protected with two load sharing power supplies, redundant fans, and two separate storage controllers
  - RAID5 layout across Flash elements within each Flash module, and then again RAID5 layout across the Flash modules

Redundant Hardware Elements

- **Data Fabric, Management Network and Storage Fibre Channel Network –** pairs of switches are used to provide complete failover redundancy
- **Processing Nodes** – Organized into Highly Available clusters to provide continuous operation in the event of failure of one of the nodes
  - Partitions of failed node are reassigned evenly to the surviving nodes within the same rack
  - System is designed with excess processing capacity, as to provide good performance in the event of a node failure
  - System can operate with as few as one half of the original nodes + 1

HA/DR recommendations with Accelerator

- Db2 Analytics Accelerator supports and complements existing HA/DR environments
- If applications leveraging the accelerator rely on it's availability similarly to Db2: consider additional accelerators to follow existing HA/DR concepts
  - Add an accelerator per HA site
  - Add an accelerator per DR site
  - If bandwidth between sites permits, keep all accelerators up to date and use them
- Data needs to be maintained on all active accelerators
- IBM Redpaper on HA/DR with the Accelerator (PDA based, but concepts still apply):
  - http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/redp5104.html?Open

HA and DR with GDPS PPRC and GDPS XRC

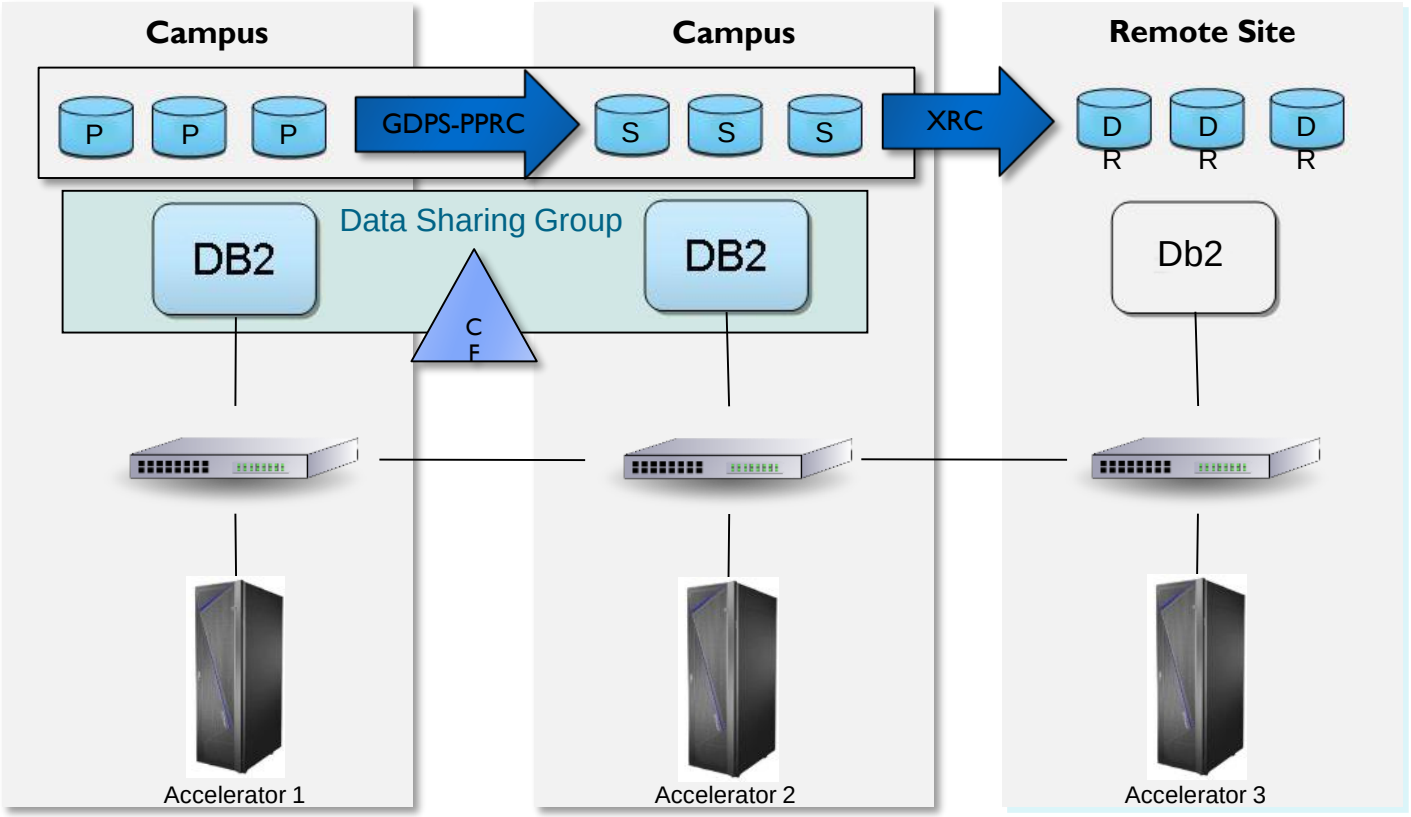

Workload Balancing

- If multiple accelerators are defined to a Db2 subsystem and based on the accelerated tables a query could be routed to multiple accelerators, Db2 uses accelerator utilization information to route the query to an accelerator
- Ensures that all defined accelerators are well balanced utilized

High availability concepts – Workload Balancing

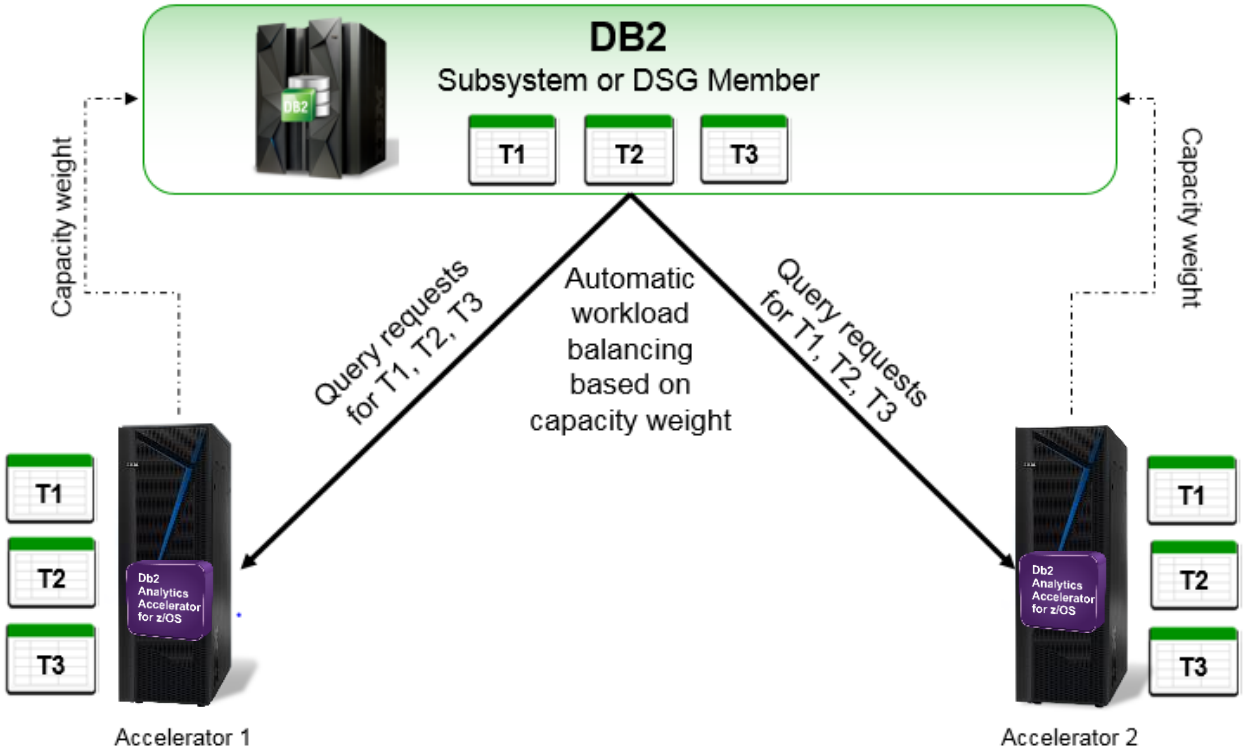

Flexible Deployment

- 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

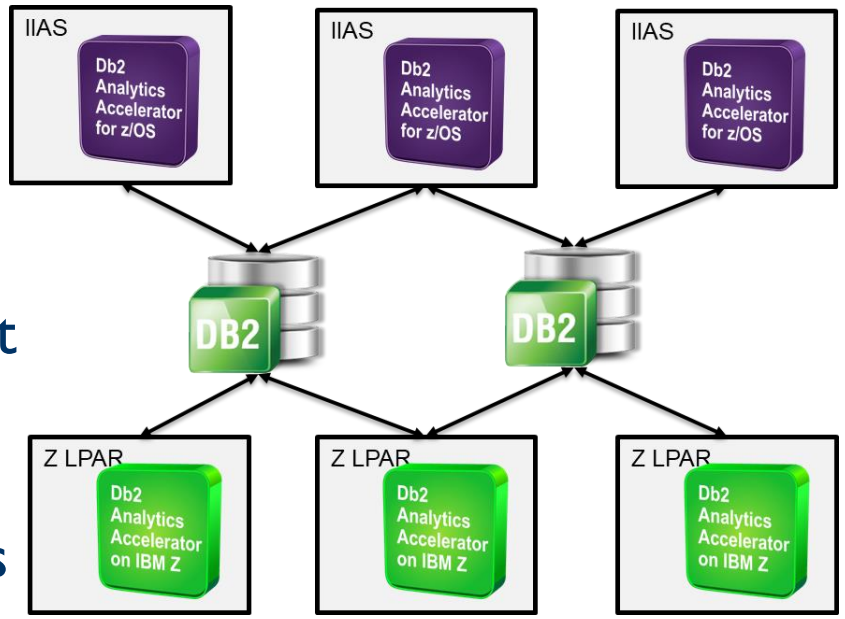

Better utilization of Accelerator resources
Scalability
High availability
Multiple options to deploy Dev/Test/QA

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

Considerations – IIAS

- Migration
  - Data movement
  - Benchmarking
- Connectivity
  - Enables external connectivity to ALL compute nodes
  - Unlike Netezza where access is only to Host 1 and Host 2
- Expansion
  - Add resources to exiting appliance

Value Proposition

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

Granularity and level of control

- zParms
- Special Registers
- Heuristic rules

Premier Business Partner IBM

zParms

- ACCEL=COMMAND | AUTO | NO
  - COMMAND to start the accelerator by manually
  - AUTO to automatically start the accelerator when the Db2 for z/OS starts
  - NO to specify that the accelerator cannot be used with this Db2 subsystem
- *GET_ACCEL_ARCHIVE=NO | YES*

  Specifies whether data that has been moved by the High-Performance Storage Saver (HPSS) is searched when a query is executed
  - *YES moved data is searched*
  - *NO moved data is not searched*

Premier Business Partner IBM

Query Acceleration - zParm and Special Register

- NONE:
  - The query always runs in Db2
- ENABLE <with FAILBACK>
  - The query runs in the Accelerator if it is validate by the heuristic and costing Db2 Optimizer check <will route to Db2 on failure on PREPARE or first OPEN>
- ELIGIBLE:
  - The query always runs in the Accelerator (skips heuristics & costing in the ENABLE option). The Query runs in Db2 only if it contains functions not supported by the Accelerator
- ALL:
  - The query always runs in the Accelerator. If the query contains functions not supported by the Accelerator it fails with return code -4762

Heuristics to evaluate acceleration

- In general, typical OLTP access path patterns are *not* routed to the accelerator,
  - Equal unique access
  - One row fetch access
- If *none* of these: WHERE, GROUP BY, ORDER BY, aggregate functions is specified (i.e. all rows are to be returned), the query is not routed
  - If all the tables referred in the query are "small", the query is not routed
  - SYSIBM.DSN_PROFILE_ATTRIBUTES (Db2 v11)
    - ACCEL_TABLE_THRESHOLD (default: 1,000,000 rows)
    - ACCEL_RESULTSIZE_THRESHOLD (in thousands, default: -1)
    - ACCEL_TOTALCOST_THRESHOLD (default: 5,000)

Hands-on Lab 1

- Data Studio GUI
- Managing the IDAA
  - Adding, Loading, Enabling Tables
  - Running queries without/with acceleration
  - EXPLAIN Access Paths

Premier Business Partner IBM

Adding a new Accelerator – Pairing Code

Use TSO TELNET to access the IIAS to get the pairing information

```
OPTION ===> TSO TELNET 10.101.8.100 1600
(1) - Generate a pairing code and display IP-address and
  port.
Accelerator pairing information:
Pairing code   : 7189
IP address     : 10.101.8.100
Port           : 1400
Valid for      : 30 minutes
 Completing the authentication using the Add New
  Accelerator
```

Distribution and Clustering in the Accelerator

- In general, use defaults
  - No Distribution Keys - random/balanced distribution
  - No Organizing Keys – no sequencing of data in a node
  - Incremental Update is an exception
- For larger tables, joined frequently, orders of magnitude improvement can be seen by specifying Organizing and or Distribution keys
  - Distribution key:
    - Determines how data is partitioned across worker nodes
  - Organizing key:
    - Determines the row clustering on disk within each node

Premier Business Partner IBM

Distribution Keys

- Db2 table data added to the accelerator is distributed among multiple worker nodes within the accelerator
  - By default, data is distributed among the worker nodes randomly
- Sometimes, large amounts of data needs to be redistributed or broadcasted within the cluster of worker nodes to join tables
- The IBM Db2 Analytics Accelerator therefore allows specifying columns for the distribution
  - Using the join columns of two tables as Distribution key, causes the rows with the same join values to be stored on the same worker node
- This enables co-located joins without data broadcasts

Premier Business Partner IBM

Organizing Keys

- Without an Organizing key, the data range between MIN and MAX could be the full spectrum of values
- Db2 BLU maintains a synopsis table that stores MIN and MAX value of Organizing key columns
- By specifying Organizing keys, the data is "ordered" for the specified columns, allowing smaller value ranges for these columns
- The synopsis table is created and maintained automatically for column-organized tables and doesn't require any configuration or maintenance to use.

Hands-on Lab 1

- What did we learn?
- Questions?
- Problems?

Premier Business Partner IBM

Hands-on Lab 2

- Manage IDAA through XML and JCL
- Incremental Update
- Static SQL

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**	▪ **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**

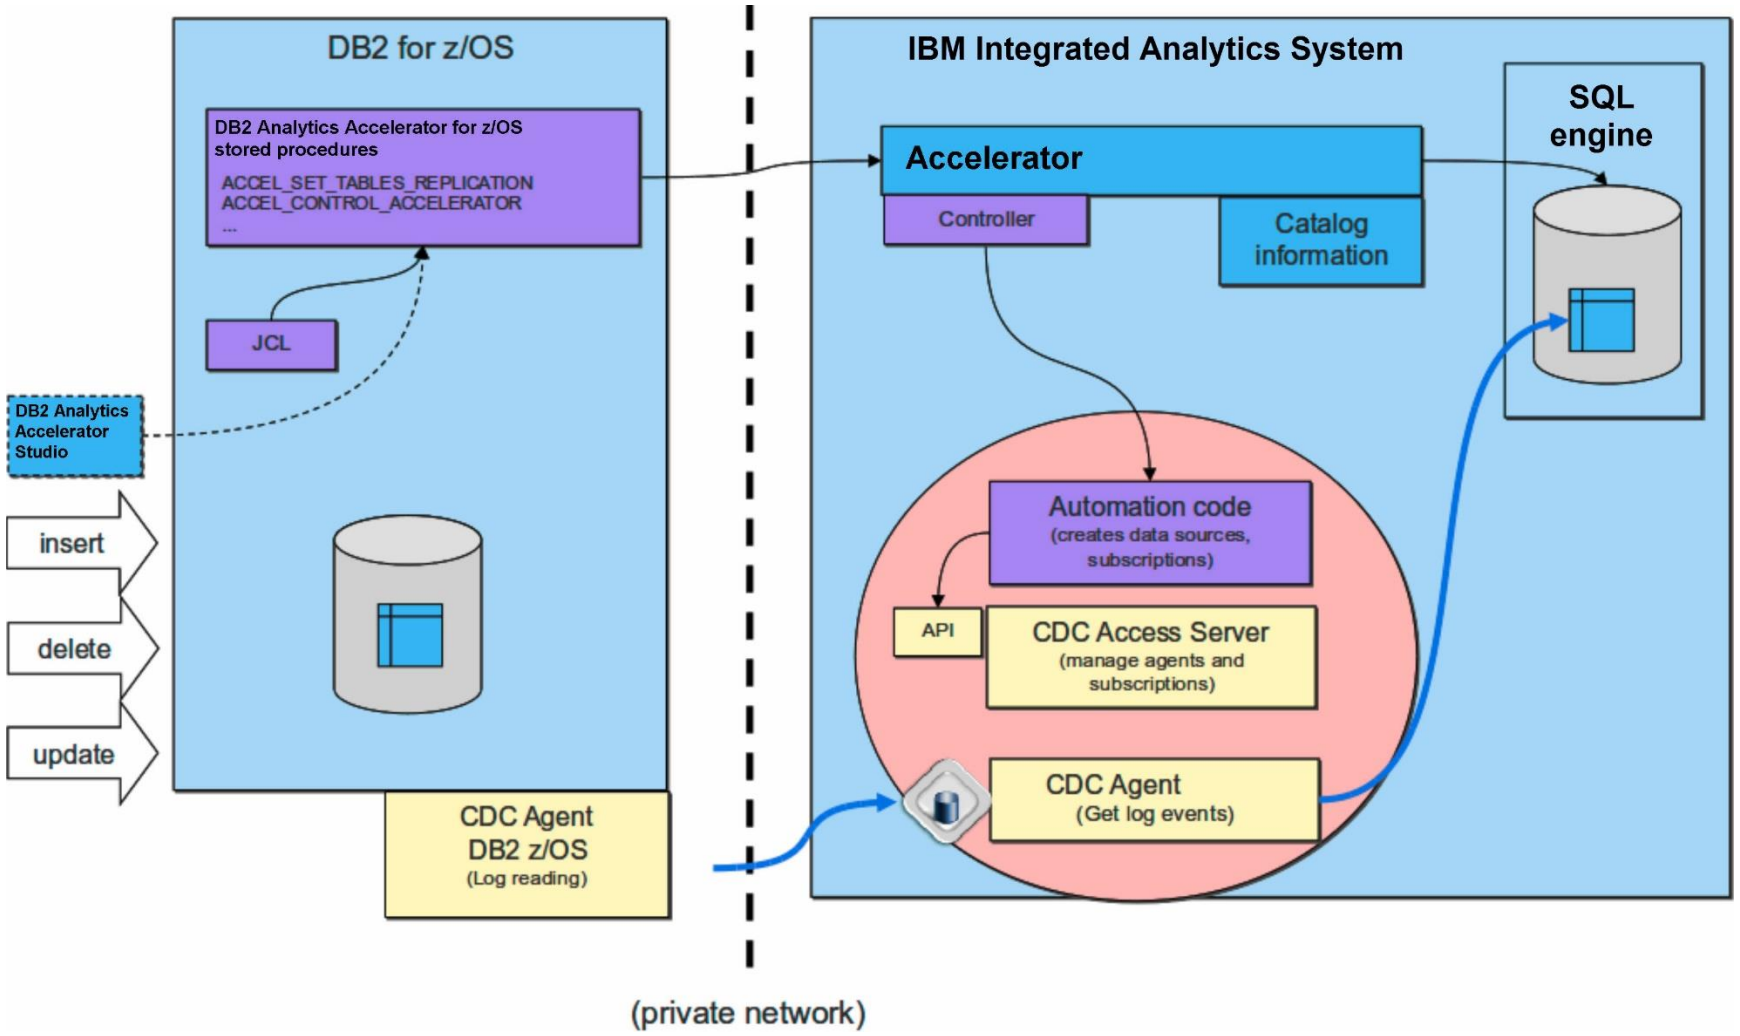
DB2 for z/OS
DB2 Analytics Accelerator for z/OS
stored procedures
ACCEL_SET_TABLES_REPLICATION
ACCEL_CONTROL_ACCELERATOR
...
JCL
DB2 Analytics Accelerator Studio
insert
delete
update
CDC Agent
DB2 z/OS
(Log reading)
IBM Integrated Analytics System
Accelerator
Controller
Catalog information
SQL engine
Automation code
(creates data sources, subscriptions)
API
CDC Access Server
(manage agents and subscriptions)
CDC Agent
(Get log events)
(private network)

Hybrid Transaction/Analytical Processing (HTAP)

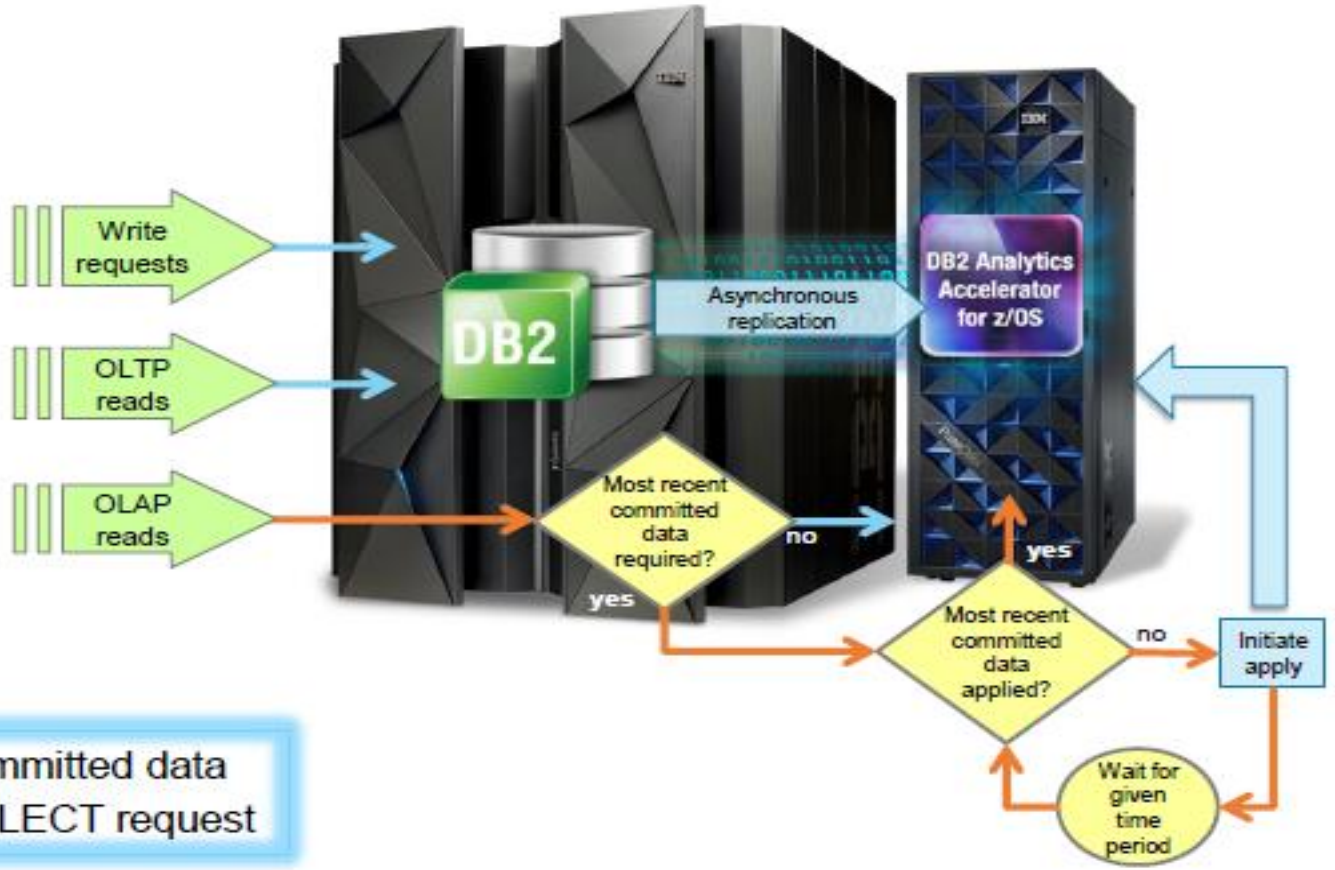

Static SQL

- New BIND options
  - QUERYACCELERATION
    - NONE
    - ENABLE
    - ENABLEWITHFAILBACK
    - ELIGIBLE
    - ALL
  - *GETACCELARCHIVE*
    - *NO*
    - *YES*

Premier Business Partner IBM

Static SQL (cont)

- The SYSIBM.SYSPACKSTMT column STATUS will have a new value, 'O' , for a static query that has been bound for acceleration during bind package
- To accelerate queries in a static package
  - First add and enable the referenced Db2 table(s) to the accelerator
  - Then BIND/REBIND the query's package specifying the desired QUERYACCELERATION behavior to bind queries in that package for acceleration

Static SQL (cont)

- *To use data that is **archived** to the accelerator*
  - *First archive at least one partition of the referenced Db2 table(s)*
  - *Then BIND/REBIND the package specifying GETACCELARCHIVE(YES) and the desired QUERYACCELERATION behavior to bind queries in that package for acceleration and retrieve archived data*

EXPLAIN

- For each query (irrespective of number of query blocks) row inserted into following tables:
  - PLAN_TABLE's ACCESSTYPE column is set to a value of 'A'
  - DSN_QUERYINFO_TABLE's QI_DATA column shows the converted query text
  - If the query is not re-routed REASON_CODE and QI_DATA columns provide details

Column Name	Column Contents
QUERYNO	The statement identification, the same value as in PLAN_TABLE. Use it with EXPLAIN_TIME to correlate DSN_QUERYINFO_TABLE and PLAN_TABLE
QBLOCKNO	
QINAME1	If REASON_CODE = 0, the name of the accelerator
QINAME2	If REASON_CODE = 0, the location of the accelerator
APPLNAME	The name of the application plan for the row. Applies only to embedded EXPLAIN statements that are executed from a plan or to statements that are explained when binding a plan. A blank indicates that the column is not applicable.
PROGNAME	The name of the program or package containing the statement being explained. Applies only to embedded EXPLAIN statements and to statements explained as the result of binding a plan or package. A blank indicates that the column is not applicable.
VERSION	The version identifier for the package. Applies only to an embedded EXPLAIN statement executed from a package or to a statement that is explained when binding a package. A blank indicates that the column is not applicable.
COLLID	The collection ID for the package. Applies only to an embedded EXPLAIN statement that is executed from a package or to a statement that is explained when binding a package. A blank indicates that the column is not applicable.
GROUP_MEMBER	The member name of the Db2 that executed EXPLAIN. The column is blank for non-data sharing.
SECTNOI	The section number of the statement.
SEQNO	
EXPLAIN_TIME	The time at which the statement is processed. This time is the same as the BIND_TIME column in PLAN_TABLE.
TYPE	'A' identifies a query that is considered for acceleration. REASON_CODE identifies if the query qualifies for acceleration or not.
REASON_CODE	If 0, the query qualifies for acceleration. Otherwise, the query cannot be accelerated. More details on the next chart.
QI_DATA	If REASON_CODE = 0, the text of the converted SQL statement (sent to IDAA). Otherwise, the description of the reason for not qualifying for acceleration
SERVICE_INFO	IBM internal use only
QB_INFO_ROWID	IBM internal use only

Value	Description
0	Query qualifies for acceleration
1	No active accelerator was found when EXPLAIN was executed.
2	The special register CURRENT QUERY ACCELERATION is set to NONE.
3	The query is a Db2 short running query or re-routing to the accelerator is not considered advantageous.
4	The query is not read-only
5	The query is running under the private protocol.
6	The cursor is defined as scrollable or rowset cursor.
7	The query refers to multiple encoding schemes.
8	The query FROM clause specifies a data-change-table-reference.
9	The query contains a correlated table expression.
10	The query contains a common table expression reference.
11	The query contains an unsupported expression. QI_DATA contains the expression text.
12	The query references table *table-name* that is either not defined in accelerator, or the table is defined, but is not enabled for query re-routing.
13	The accelerator *accelerator-name* containing the tables of the query is not started.
14	The column *column-name* referenced in the query is altered in Db2 after the data is loaded in the accelerator.
900 through 999	IBM internal use

Hands-on Lab 2

- What did we learn?
- Questions?
- Problems?

Premier Business Partner IBM

Hands-on Lab 3

- Accelerator-only Tables (AoTs)

Premier Business Partner IBM

Accelerator-only Tables (AoTs)

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

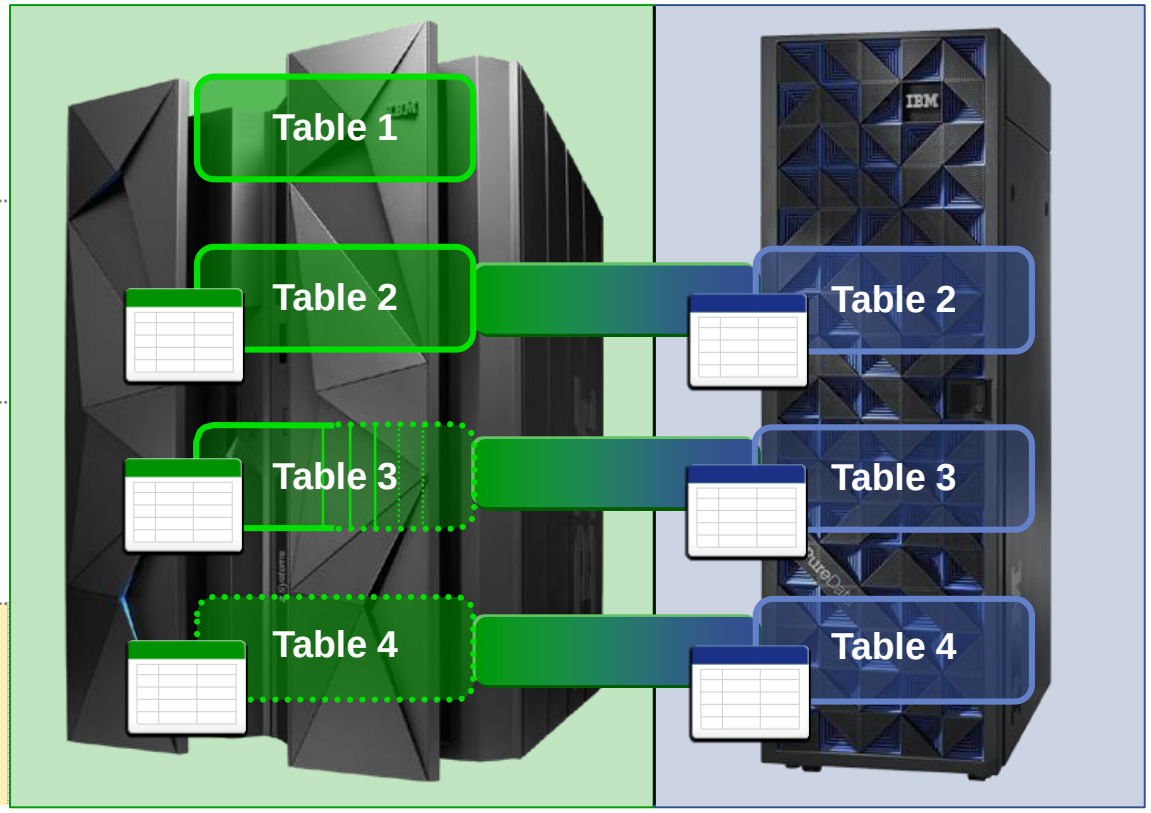

Transaction Processing Systems (OLTP)

Analytics

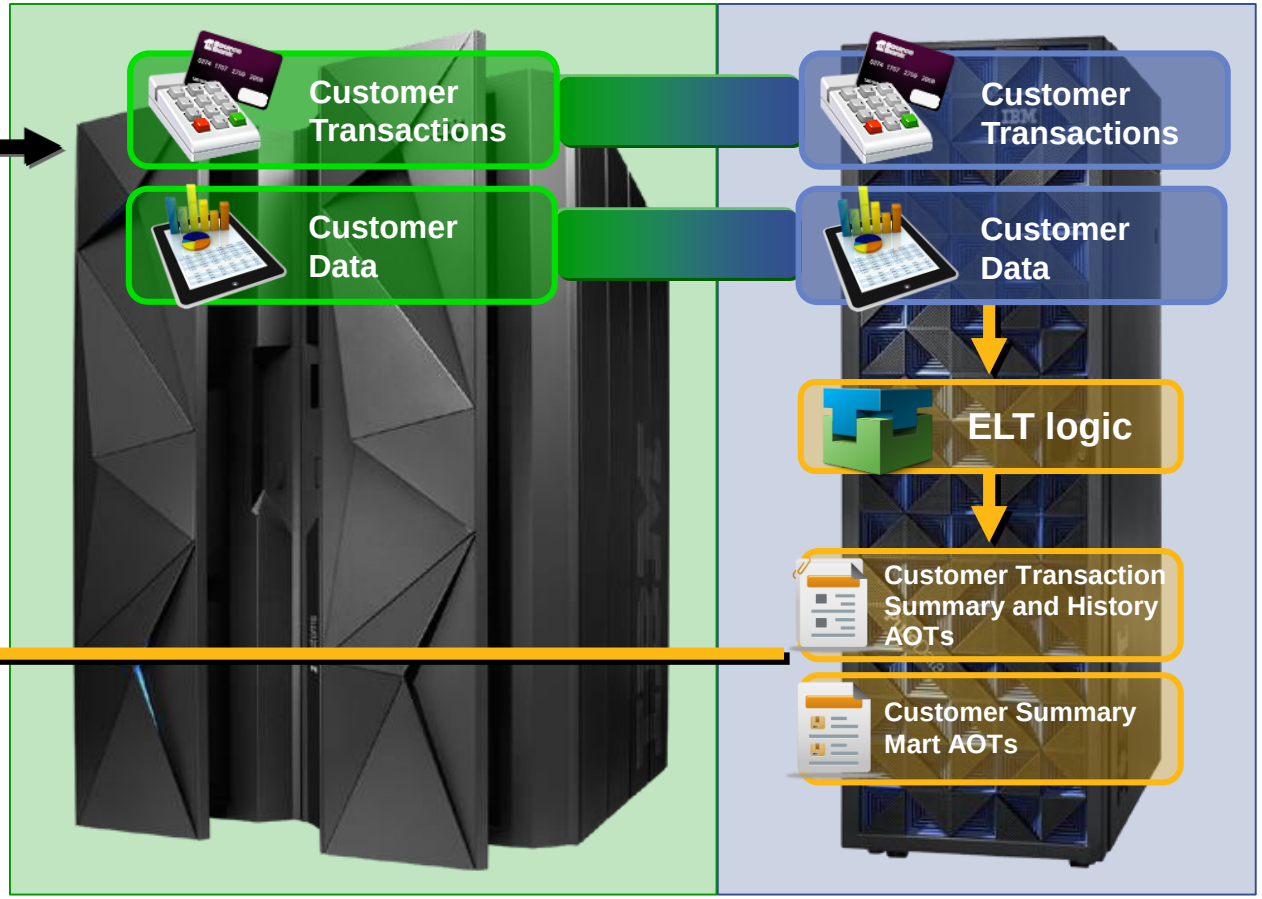

Advantages:

- Simpler to manage
- Better performance and reduced latency

Data for transactional and analytical processing

Hands-on Lab 3

- What did we learn?
- Questions?
- Problems?

Premier Business Partner IBM

Next Steps

- Whiteboarding session
- Workload Assessment
  - Workload on Db2
  - Competitor workload (e.g. Teradata, MS SQL Server)
- Cost Benefit Analysis (CBA)
- Proof-of-concept (POC)

Business Use Case White Boarding Session

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

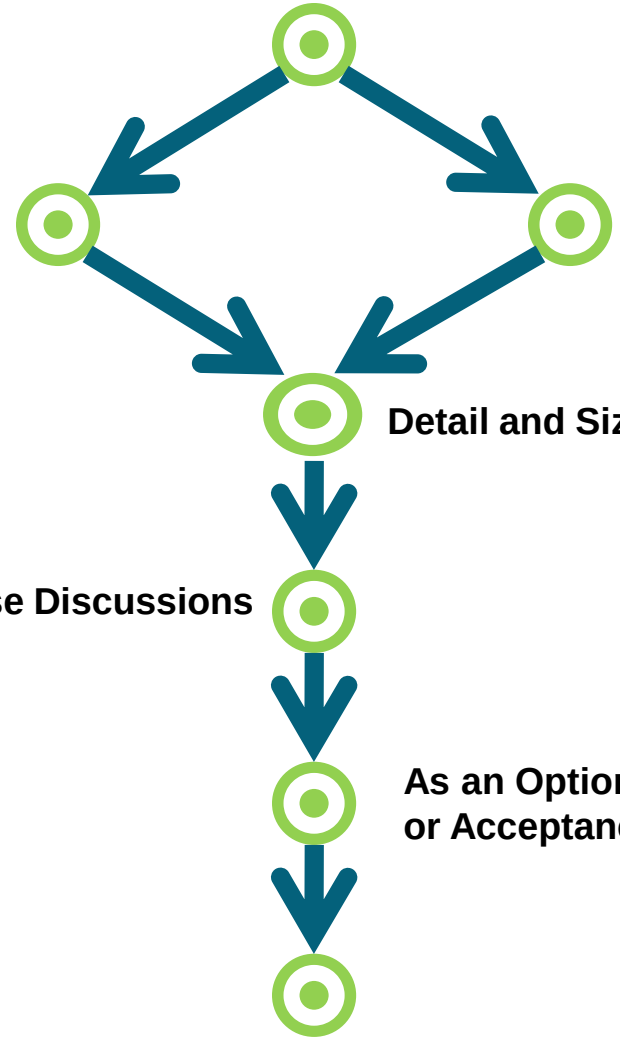

Hands on Workshop

- 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

Workload Assessment

- Activate Dynamic Statement Cache
  - CACHEDYN subsystem parameter
- Activate relevant IFCID, start trace classes
  - `-START TRACE(MON) CLASS(30) IFCID(316,317,318) DEST(SMF)`
  - `-START TRACE(ACCTG) CLASS(1,2) DEST(xxx)`
- Create objects for collecting workload information
  - DDL in <hlq>.SDSNSAMP(DSNTESC)

Workload Assessment (cont)

- Collect workload information from the Dynamic Statement Cache
  - `EXPLAIN STMTCACHE ALL`
  - Review DSN_STATEMENT_CACHE_TABLE
- Verify the Dynamic Statement Cache contains Analytic Queries (SQL provided)
- Optional: Mask sensitive information in SQL statements

Premier Business Partner IBM

Workload Assessment (cont)

- EXPLAIN extracted SQL
  - `%REXXEXPL <subs> <schema>`
- Unload workload, EXPLAIN and catalog information (UNLOAD job provided)
- Deliver datasets for evaluation
- Review the results

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	Stater
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)		(Prime 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%

Workload Assessment – Static SQL

- Set zParm ACCELMODEL to YES
- BIND or REBIND PACKAGE for all Db2 packages to be evaluated
- Use option APREUSE(ERROR) with the REBIND PACKAGE command to make sure that the current access path remains unchanged

Premier Business Partner IBM

Workload Assessment – Alternatives

- OMEGAMON XE for Db2 Performance Expert
- Optim Query Workload Tuner
- Manual Assessment

Premier Business Partner IBM

Cost Benefit Analysis (CBA)

- Financial justification for IDAA
- Examples:
  - Reduce expenses for System z DASD by deploying HPSS – IDAA storage costs < $1/GB/month
  - Defer a System z hardware upgrade or refresh by offloading processing to IDAA

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

Proof-of-concept – Key Players

- Project Sponsor – Person responsible for the business success of the solution
- Project Manager – Responsible for keeping the deployment tasks staffed and on track
- Db2 Systems Programmer – Installation and integration of Db2 and Accelerator software
- z/OS Systems Programmer – Configuration and integration of System z and z/OS related resources, such as the z/OS-based Workload Manager (WLM)
- Operations / Infrastructure – Preparation of the physical environment for the accelerator
- Networking – Preparation of the communications environment for the accelerator
- Database administrators and application developers – Integration of the product into applications and evaluation of the product

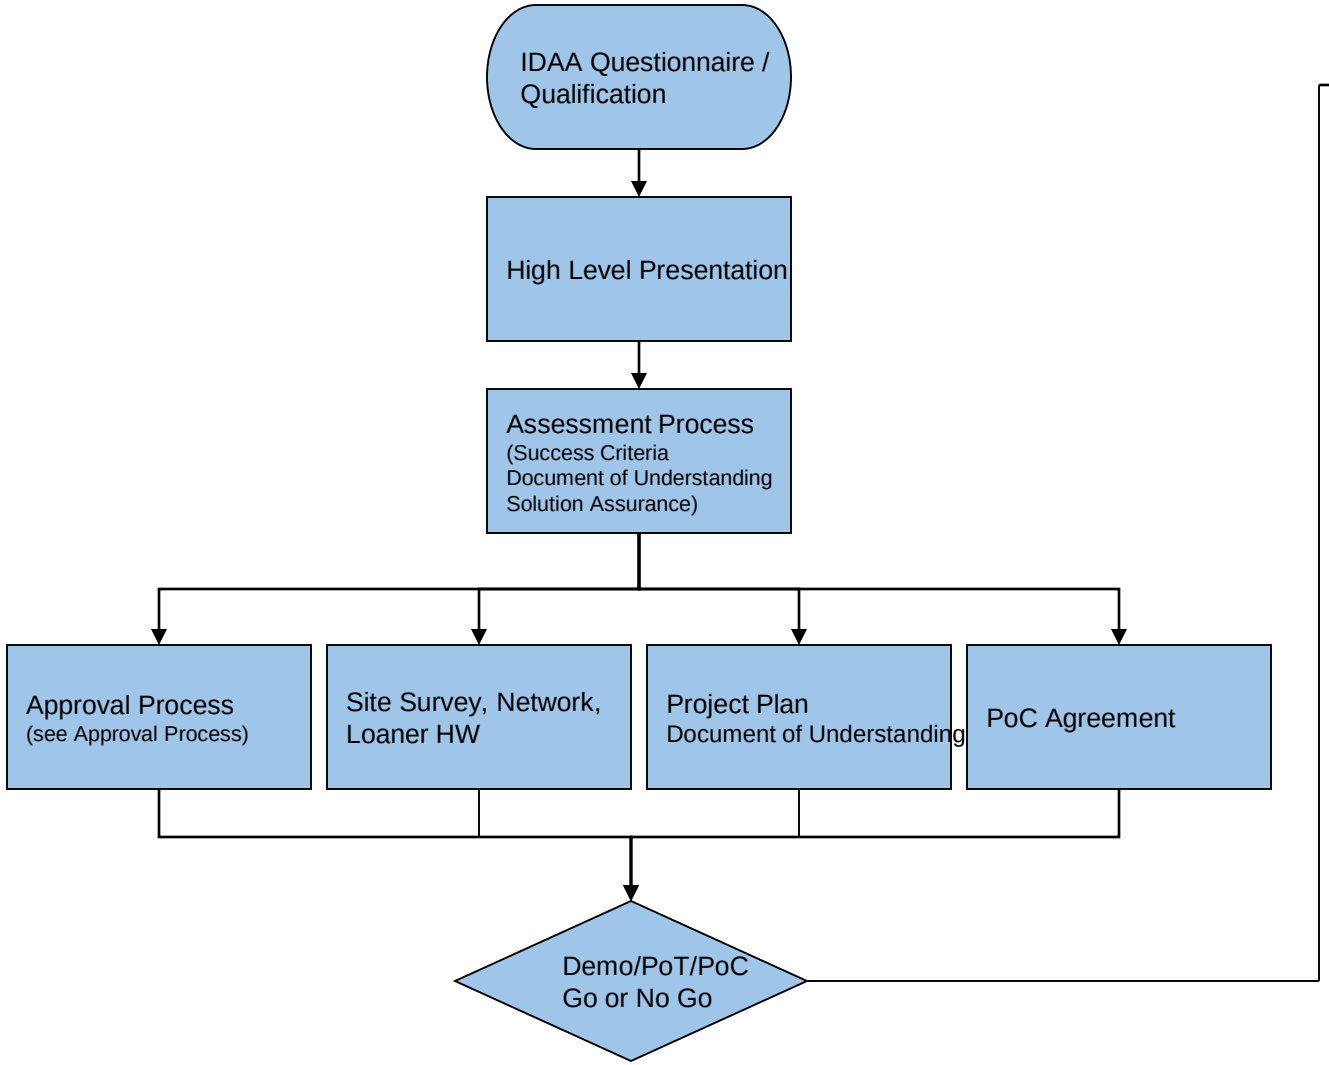
IDAA Questionnaire / Qualification
High Level Presentation
Assessment Process
(Success Criteria
Document of Understanding
Solution Assurance)
Approval Process
(see Approval Process)
Site Survey, Network, Loaner HW
Project Plan
Document of Understanding
PoC Agreement
Demo/PoT/PoC Go or No Go
Demo/PoT/PoC
(see PoC Process)
Technical Buy-in

Attributions

- Dwaine Snow, IBM
- Jeff Feinsmith, IBM
- Mike Hood, IBM
- Patric Becker, IBM Boeblingen Lab
- Knut Stolze, IBM Boeblingen Lab
- Poughkeepsie Benchmark Center Team
  - John Theos
  - John Borrero

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
  - "Enabling Real-time Analytics on IBM z Systems Platform" SG24-8272
- www.thefillmoregroup.com/blog

THE FILLMORE GROUP
Relational Database Solutions

Thank you

Kim May, Vice President Business Development

Frank Fillmore, President

www.thefillmoregroup.com

Premier Business Partner IBM