THE FILLMORE GROUP
Relational Database Solutions

IBM Db2 Analytics Accelerator

The Fillmore Group – August 2019

A Premier IBM Business Partner

Premier Business Partner IBM

Agenda

- Introductions - The Fillmore Group
- A review of capabilities included in recent IDAA v7.1.x code drops
- The return of the High Performance Storage Saver (HPSS)
- Integrated Synchronization
- The closer-to-real Hybrid Transaction Analytic Processing (HTAP) capability
- Disaster Recovery and High Availability considerations

Agenda (cont)

- Form factors
  - IBM Integrated Analytics System (IIAS)
  - Accelerator on System z Docker container running on IFLs
  - Docker container running on LinuxONE
- Migration considerations
- Buyer's Guide spreadsheet

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

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*

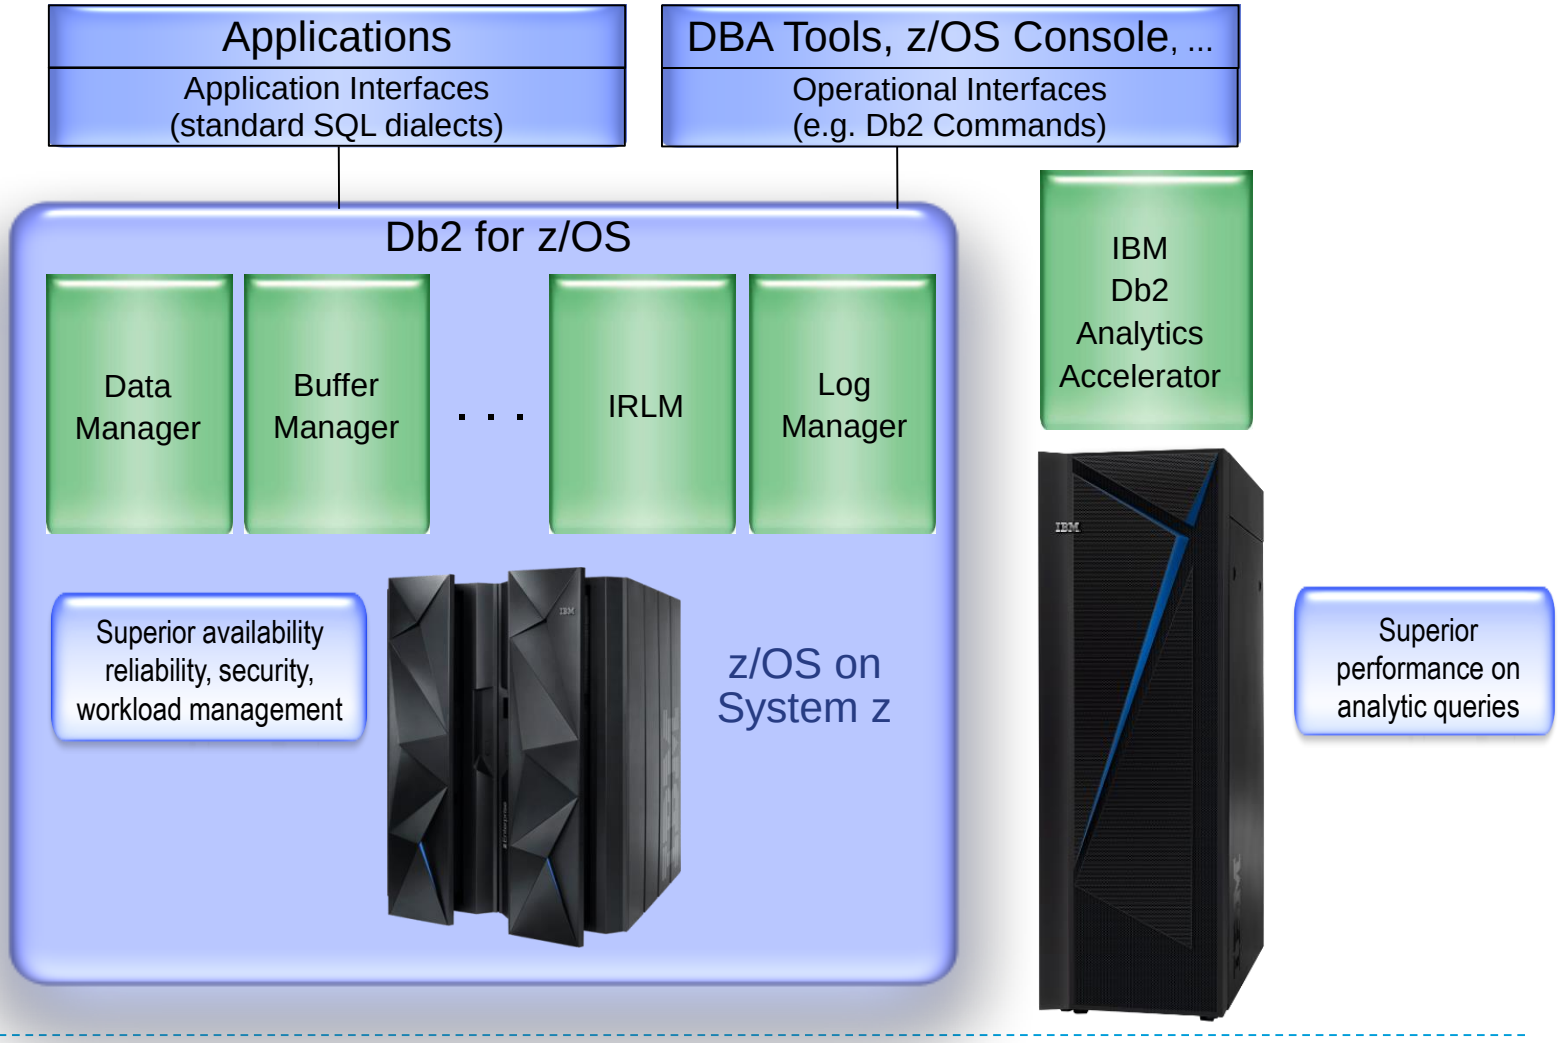
Applications
Application Interfaces
(standard SQL dialects)
DBA Tools, z/OS Console, ...
Operational Interfaces
(e.g. Db2 Commands)
Db2 for z/OS
Data Manager
Buffer Manager
. . .
IRLM
Log Manager
Superior availability reliability, security, workload management
z/OS on System z
IBM Db2 Analytics Accelerator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

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

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

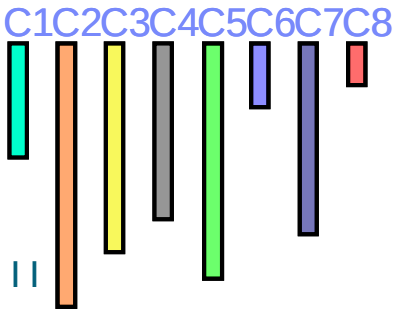

Skip Boring Data

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

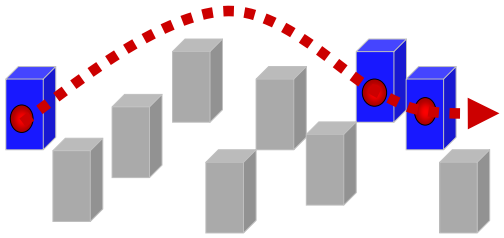

Rethink Memory

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

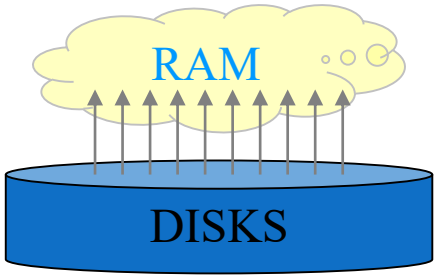

What's new in IDAA v7? – "Heart transplant"

- 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)

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.

Distribution and Clustering in the Accelerator

- 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

Distribution Keys (cont)

Identical Distribution Keys

```
CREATE TABLE ORDERS (ORDER_NO, CUST_NO, …) DISTRIBUTE BY HASH (CUST_NO)
CREATE TABLE CUSTOMERS (CUST_NO, …) DISTRIBUTE BY HASH (CUST_NO)
SELECT … FROM ORDERS O, CUSTOMERS C WHERE O.CUST_NO = C.CUST_NO
```

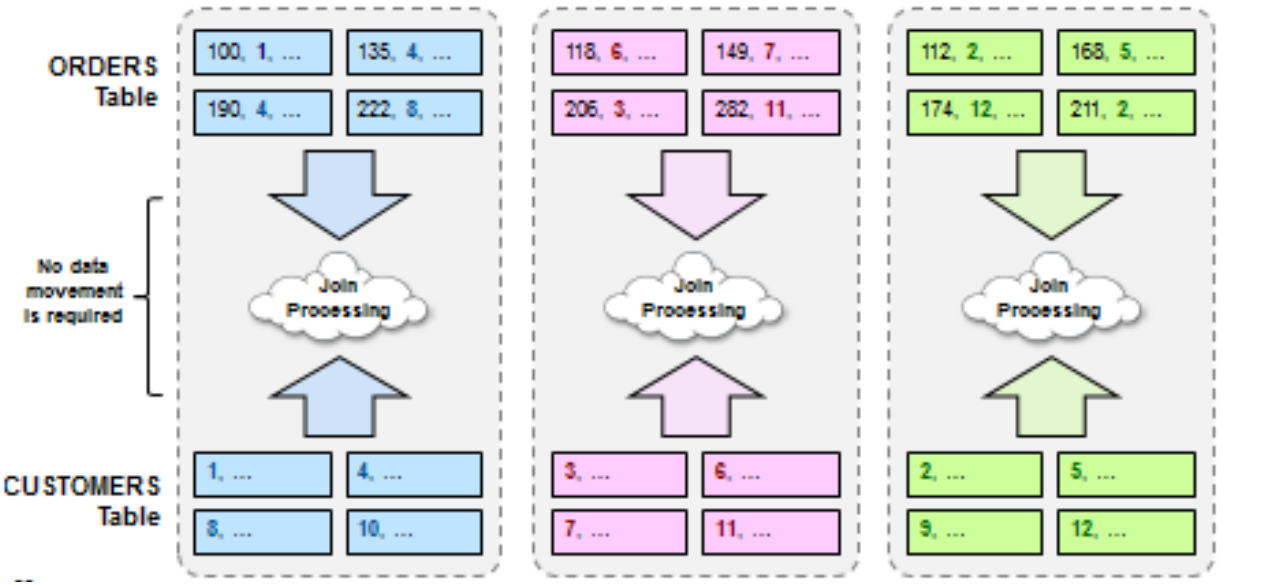

Distribution Keys (cont)

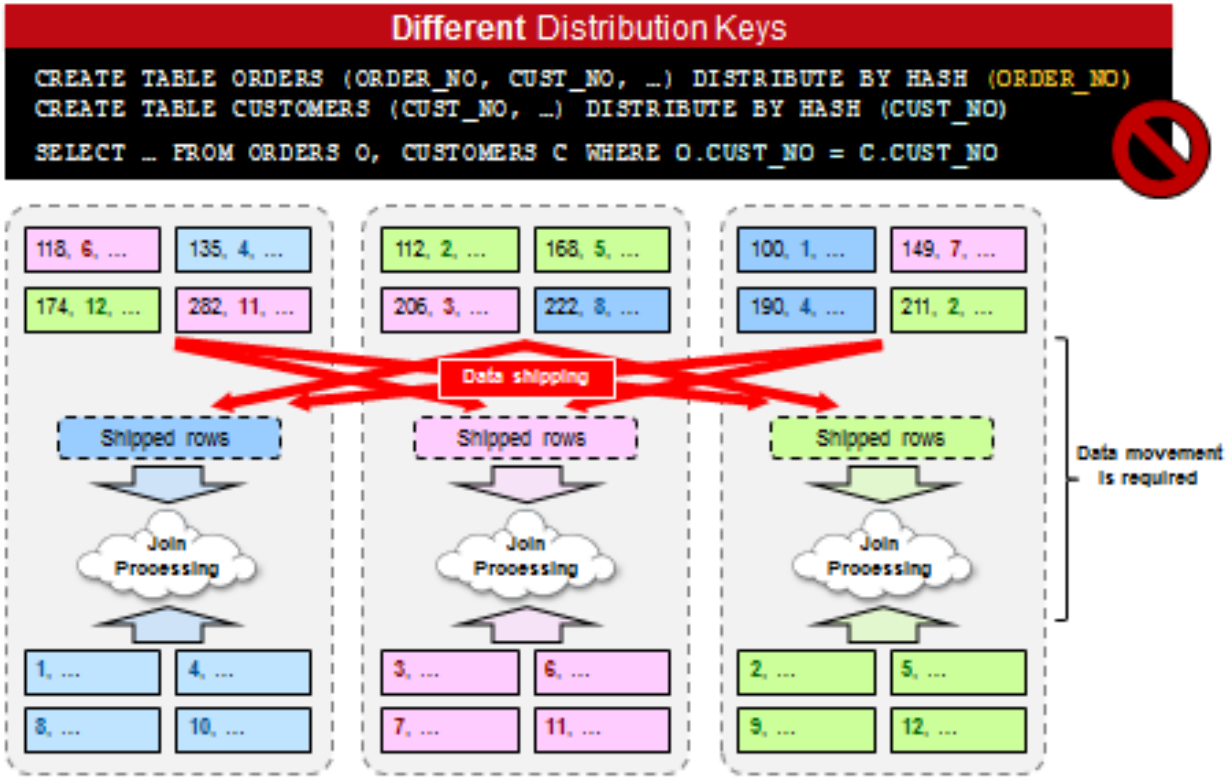

Organizing Keys

- 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.

Enhancements in 7.1.4

- Hybrid Transaction / Analytic Processing (HTAP) preview

Enhancements in 7.1.5

- Pass-through support for a set of OLAP built-in functions (Stage 1)
- Support for Local Date Format in accelerated queries
- Support for multi-row insert for accelerator-only tables
- HTAP preview: SMF counters for monitoring wait/expiration times of queries
- Accelerator on z:
  - GDPS support for fail-over scenarios
  - Script based image deployments
- Coexistence V5/V7: Option to specify default accelerator for query routing to support migration scenarios

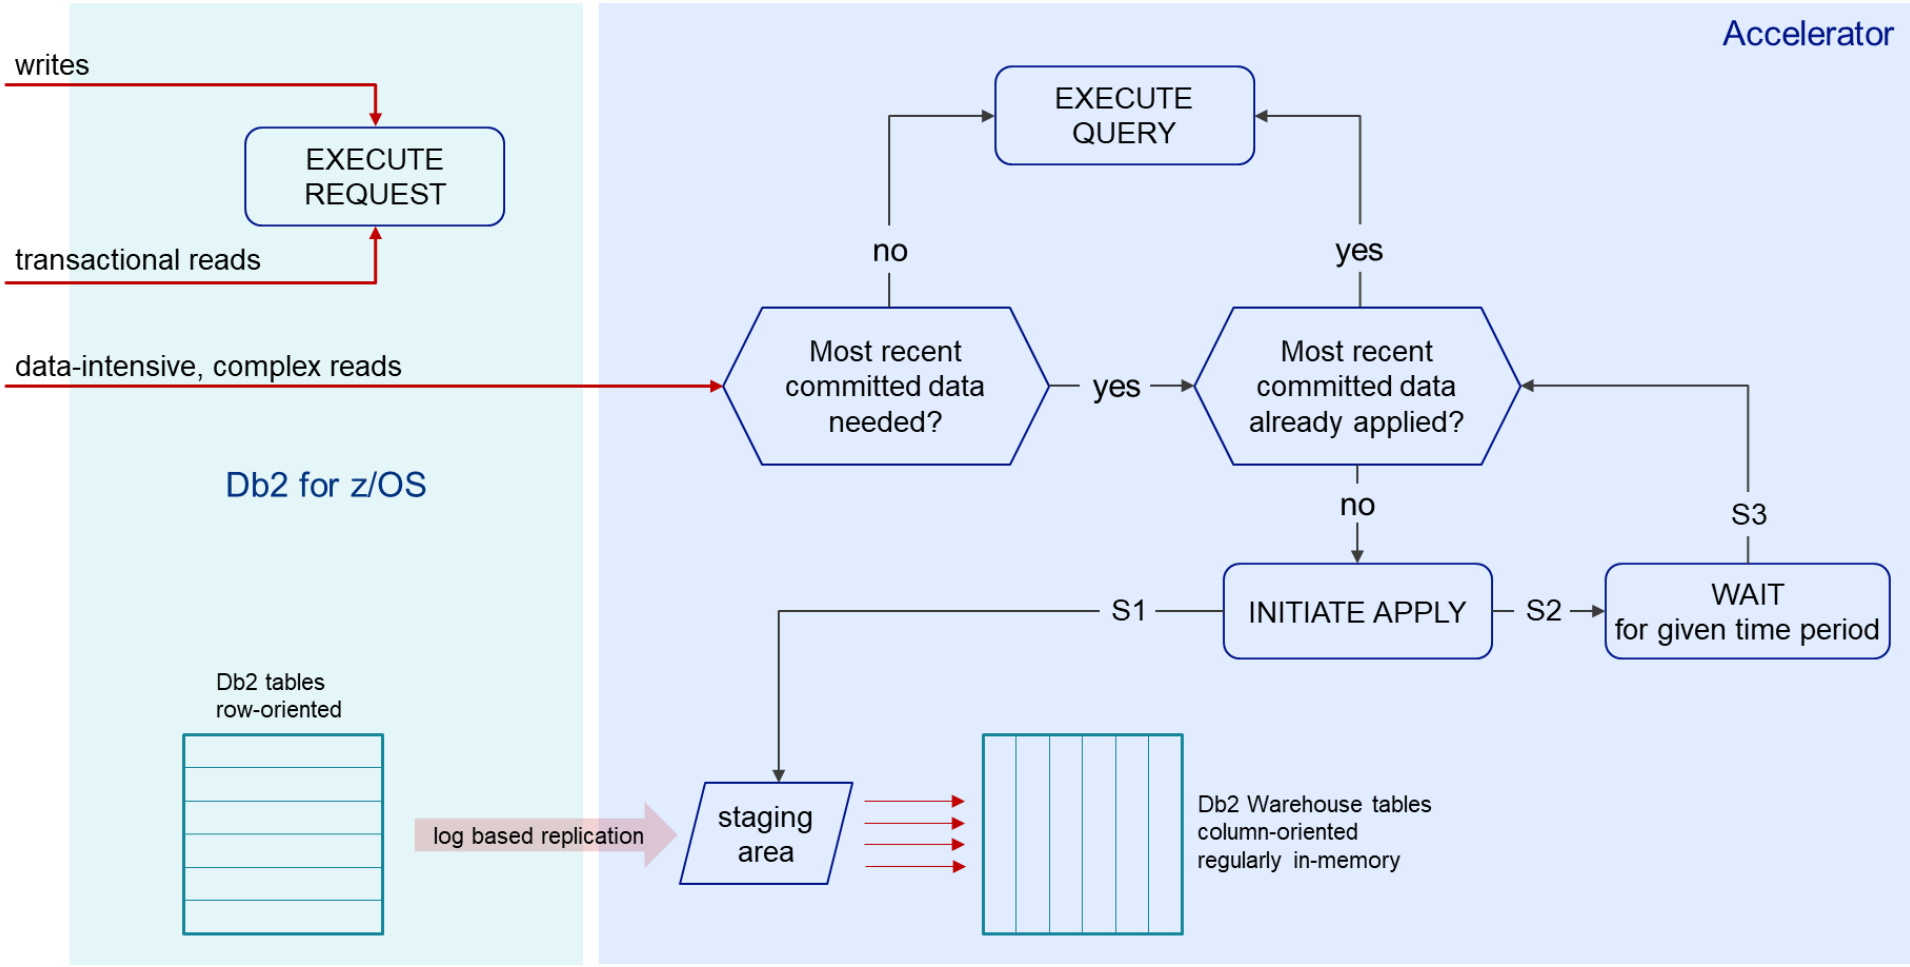
Accelerator
writes
EXECUTE REQUEST
transactional reads
data-intensive, complex reads
Db2 for z/OS
EXECUTE QUERY
no
yes
Most recent committed data needed?
yes
Most recent committed data already applied?
no
S3
S1
INITIATE APPLY
S2
WAIT for given time period
Db2 tables row-oriented
log based replication
staging area
Db2 Warehouse tables column-oriented regularly in-memory

Enhancements in 7.1.6 (Feb 2019)

- Encryption of data-in-motion: All network traffic encrypted
  - Query traffic, stored procedure traffic (e.g. data loads) and now also Incremental Update traffic
  - SSL encryption (AT-TLS on System z)
- Accelerator on IIAS:
  - Call Home via HMC
    - PMRs are opened automatically for severe hardware and software problems on the IIAS machine

Premier Business Partner IBM

Enhancements in 7.1.7 (April 2019)

- High Performance Storage Saver (HPSS) – Technical Preview
  - Supported use cases:
    - Archiving of partitions of range-partitioned tables
    - Restoring of partitions
    - High availability: Archiving to multiple accelerators
- Improved compression with full table reloads
- Support for partition reloads of replicated tables
- Support for more pass-thru built-in functions (e.g. REGEXP*)
- Support for very small deployments on IBM z

Premier Business Partner IBM

Enhancements in 7.1.8 (June 2019)

- High Performance Storage Saver (HPSS) – GA Release
- Improved compression for all use-cases (initial load, re-load and replication)
- Support for mapping Db2 for z/OS REAL column type to DOUBLE on Accelerator
- Integrated Synchronisation: Early Support Program
- Several stability and User Experience Improvements
  - Query performance: automatic column-group statistics
  - SW-Update Progress bar
  - Additional support info trace option
  - Improved table size display
  - Improved fetch performance for queries with large result sets

Premier Business Partner IBM

The core of future Accelerator Deployments

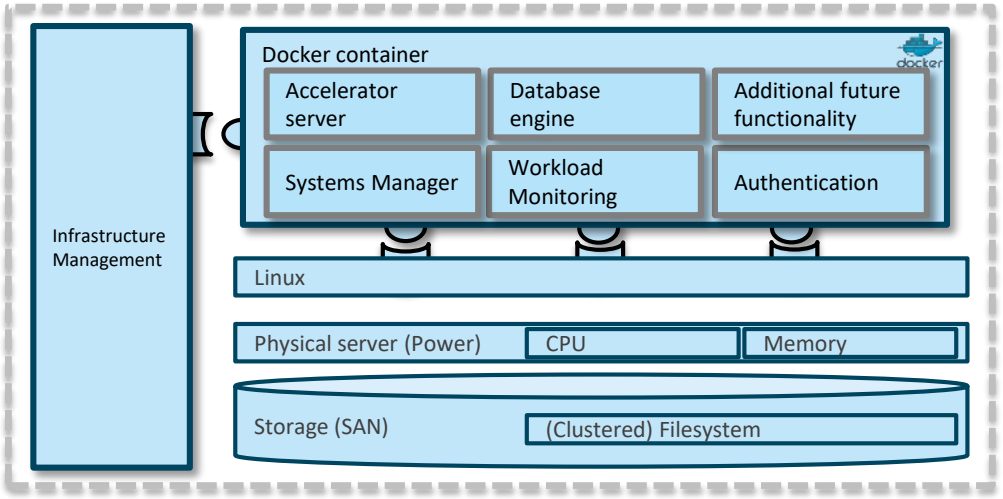

Docker is a widely adopted, open-source project that automates the deployment of applications inside software containers.

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.

Accelerator on z – Integrated Analytics System

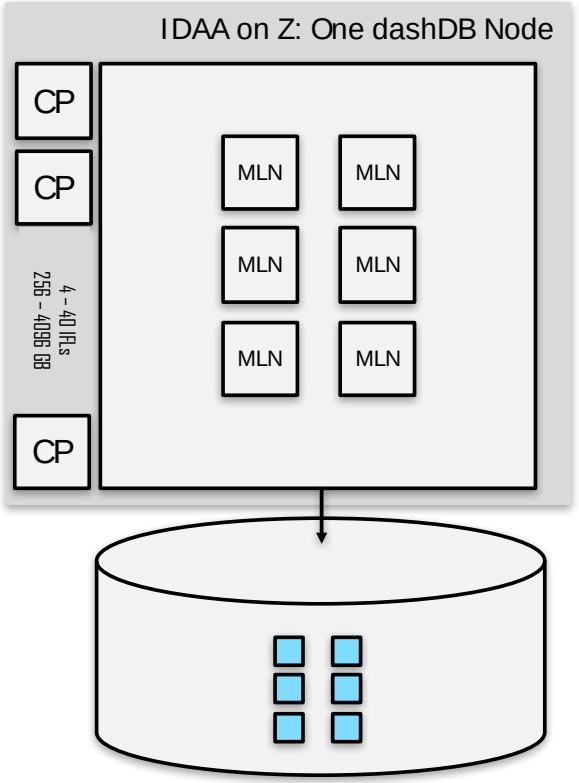

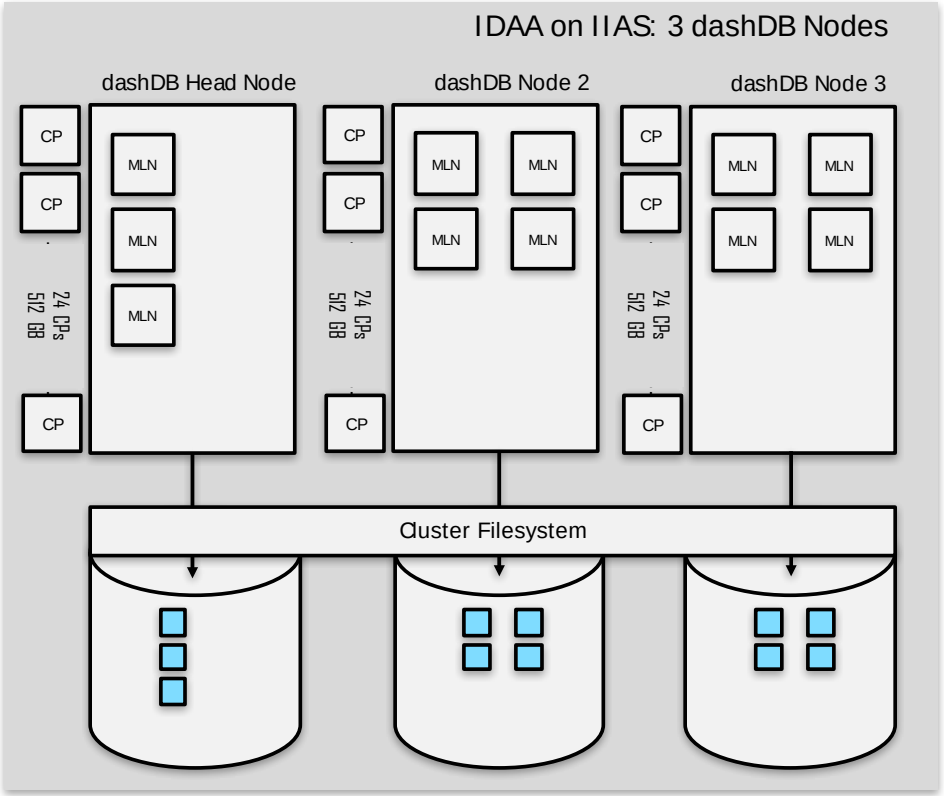

Accelerator on z – Integrated Analytics System

- Db2 Analytics Accelerator on z instances can be deployed quickly (download & go), but instances are currently limited to CPU resources in one physical z14 drawer
- Compared to IIAS based systems, higher instance granularity is possible
  - Multiple independent applications / Db2 subsystems may share a bigger IIAS system while IDAA on Z characteristics may lead to more, dedicated and smaller instances
  - Beneficial for (many) smaller test/dev instances which could be temporarily (re-) provisioned

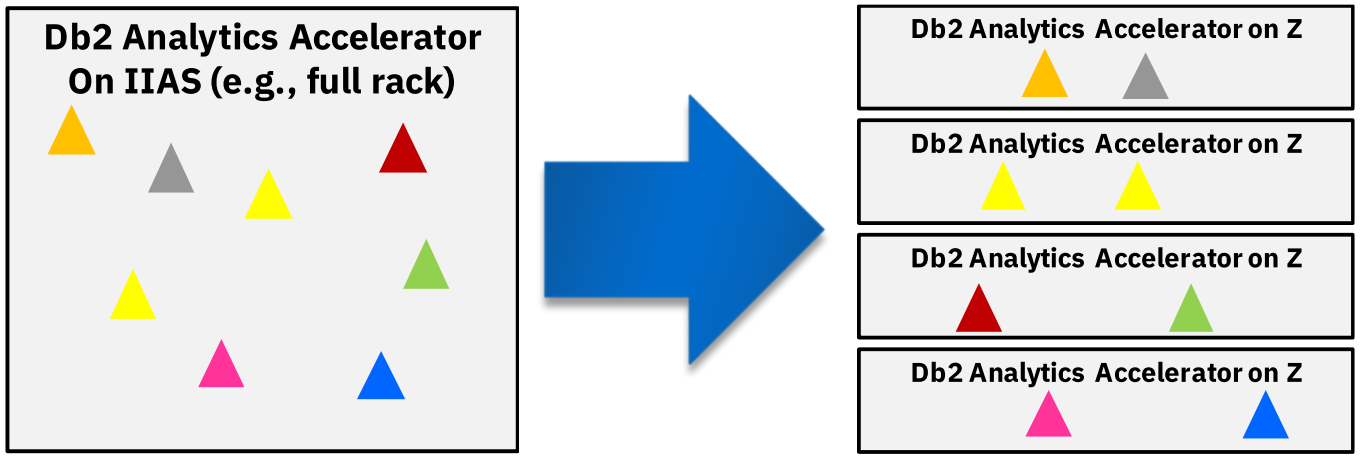

Comparison

PDA N3001 (Netezza Technology)

Integrated Analytics System M4001

Accelerator on IBM z

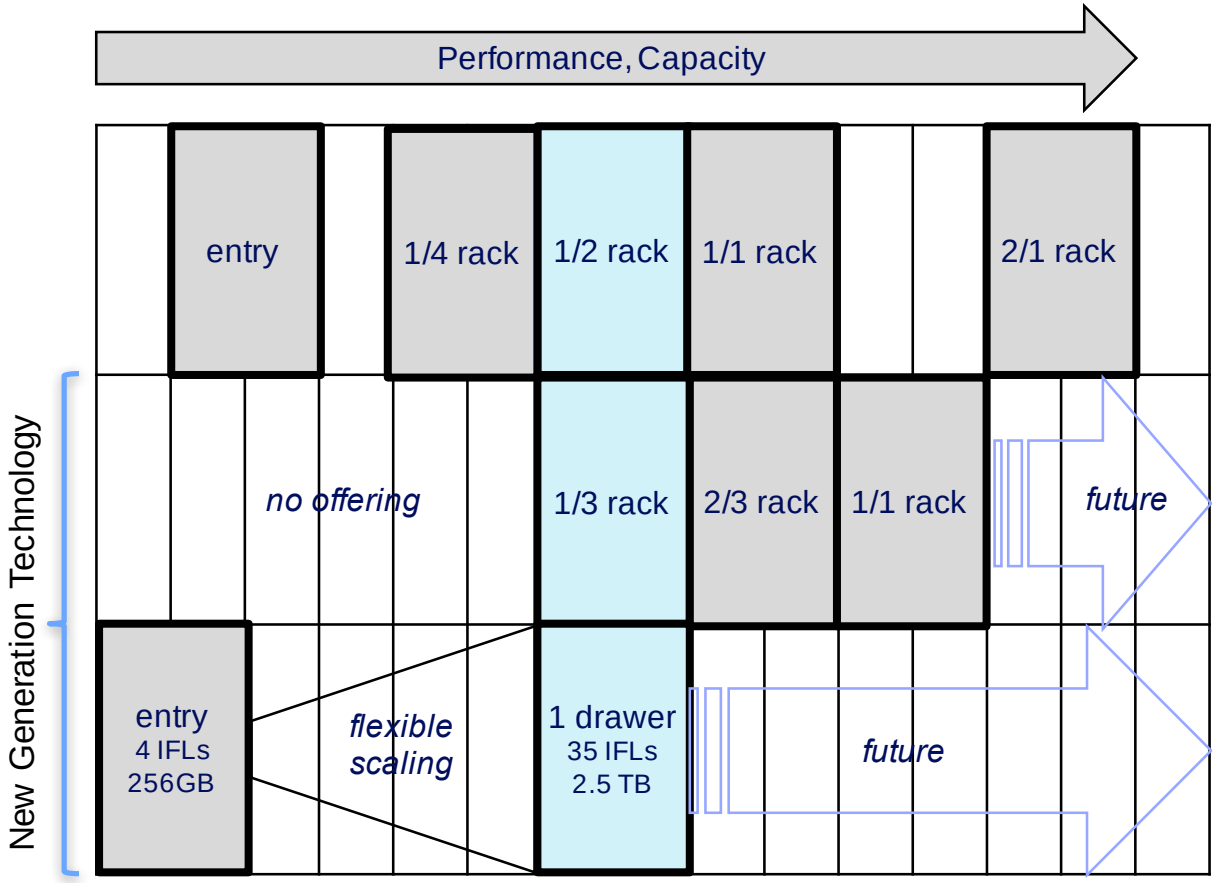

IBM Integrated Analytics System (IIAS) Configurations

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

IDAA Deployment on LinuxONE

https://www.ibm.com/it-infrastructure/linuxone

- IIAS (includes processing, memory, storage, Docker container)
- LinuxONE (includes processing, memory) *plus*
  - Storage
  - Accelerator on z Docker container

Accelerator on z on LinuxONE prerequisites:
https://www-01.ibm.com/support/docview.wss?uid=swg27050440#HWIBMZ

The Fillmore Group's IDAA "Buyer's Guide"

- Spreadsheet with three tabs – one for each form factor
  - IBM Integrated Analytics System (IIAS) (appliance)
    - All hardware/software included
  - Accelerator on z (Docker container running on IFLs)
    - Add IFLs, Memory, Storage
  - Docker container running on LinuxONE (appliance)
    - Processors and memory included; add software, storage
- Version 5 to 7 upgrade savings

THE FILLMORE GROUP
Relational Database Solutions

High Availability IIAS – 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:
  - http://www.redbooks.ibm.com/Redbooks.nsf/RedpieceAbstracts/redp5536.html?Open

HA and DR with GDPS PPRC and GDPS XRC

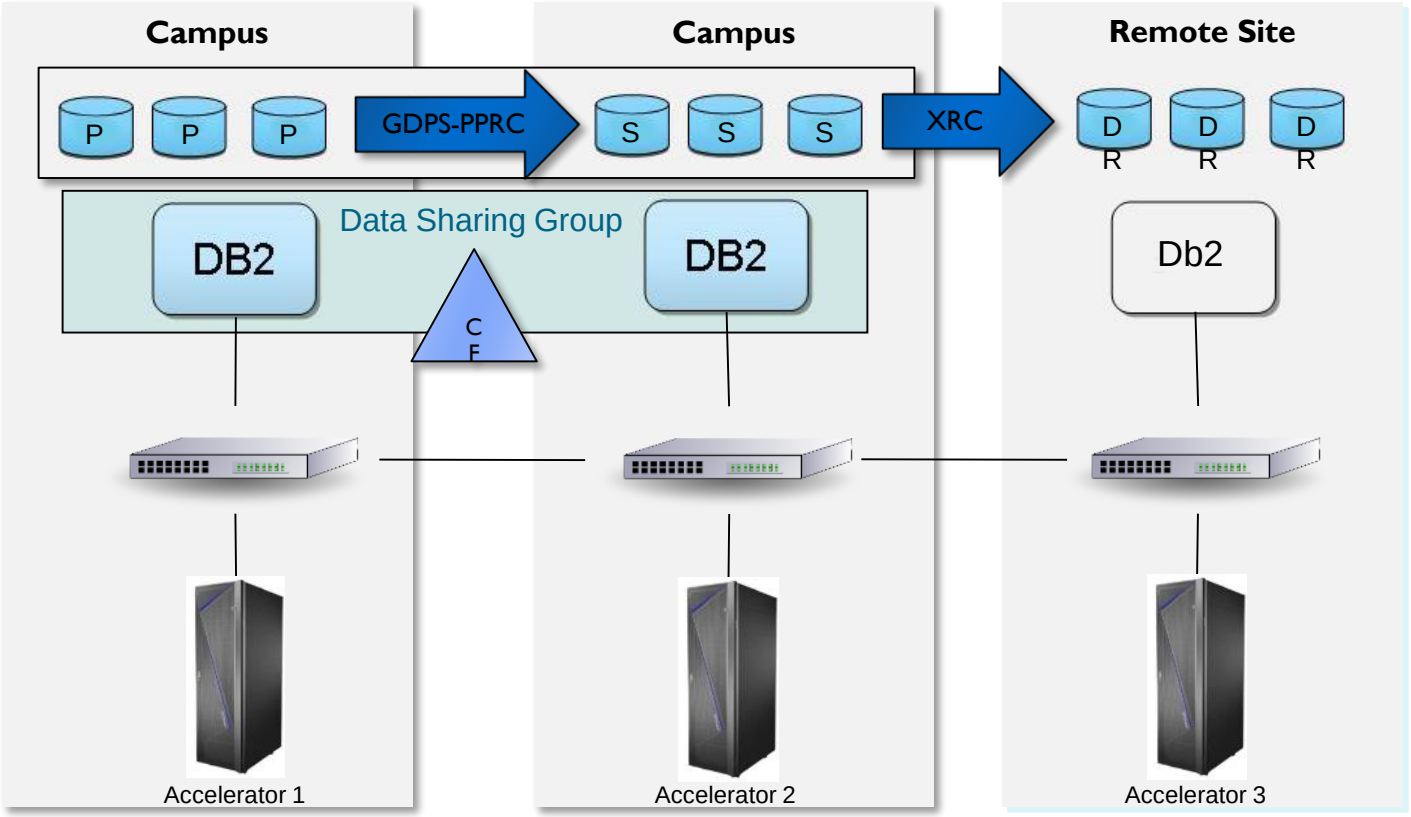

Flexible Deployment

- 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

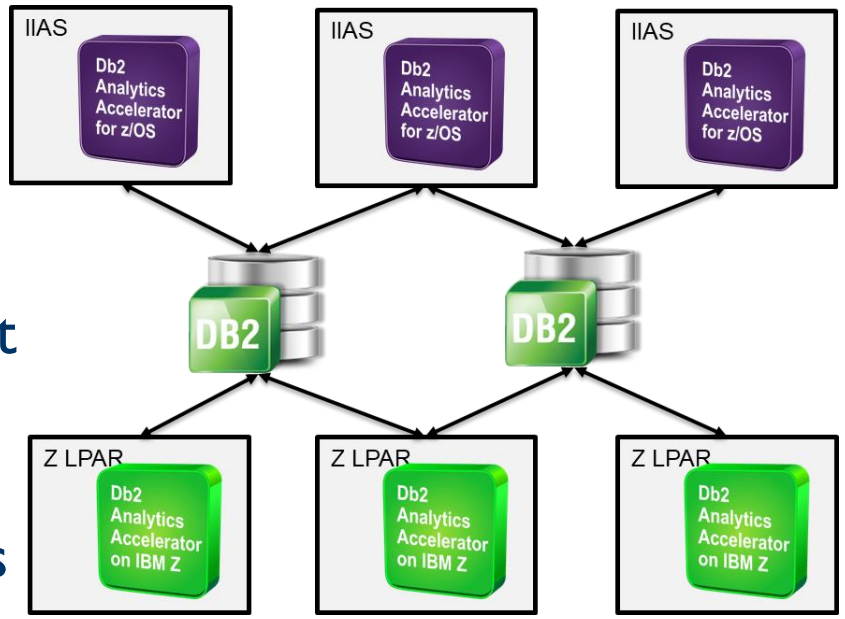

Better utilization of Accelerator resources
Scalability
High availability
Multiple options to deploy Dev/Test/QA

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

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

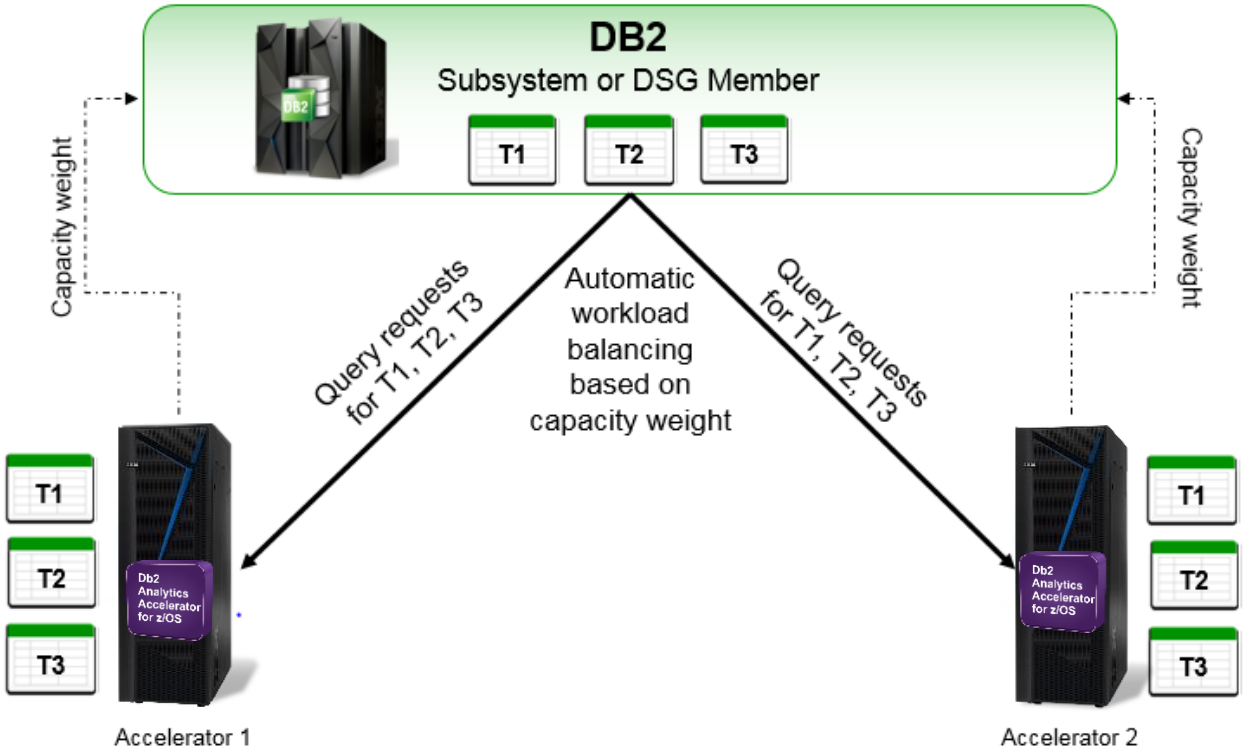

Considerations

- Coexistence and 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
  - Add IFLs for Accelerator on z

Whither IDAA v5?

- IDAA v5 was GA on **November 27, 2015**
- IBM intends to announce withdrawal from Marketing in **2019**
- There is no **End of Service** announced for v5 yet
- IDAA v5 has an Enhanced Lifecycle policy which means that this product has a minimum of five full years of standard support from the date the product was GA

Whither IDAA v5? (cont)

- Pure Data System for Analytics (PDA) N3001-XXX were GA on **October 17, 2014**
- PDA were used as appliance base for Db2 Analytics Accelerator v5 and cannot be used for v7.
- PDA were **End of Marketing** on **April 10, 2018** They cannot be ordered any more, also nothing left in stock
- There is no **End of Service** announced for PDA N3001-XXX models yet

IDAA v5 to v7 Data Migration

- Non-accelerated Db2 table – N/A
- Accelerator "shadow" table – Reload*
- Archived table or partitions (HPSS) – Restore from Db2 Image Copy
- Accelerator-only Table (AoT) – Recreate, Reload*

*Db2 Analytics Accelerator Loader

Db2 Analytics Accelerator Loader

High Availability Load

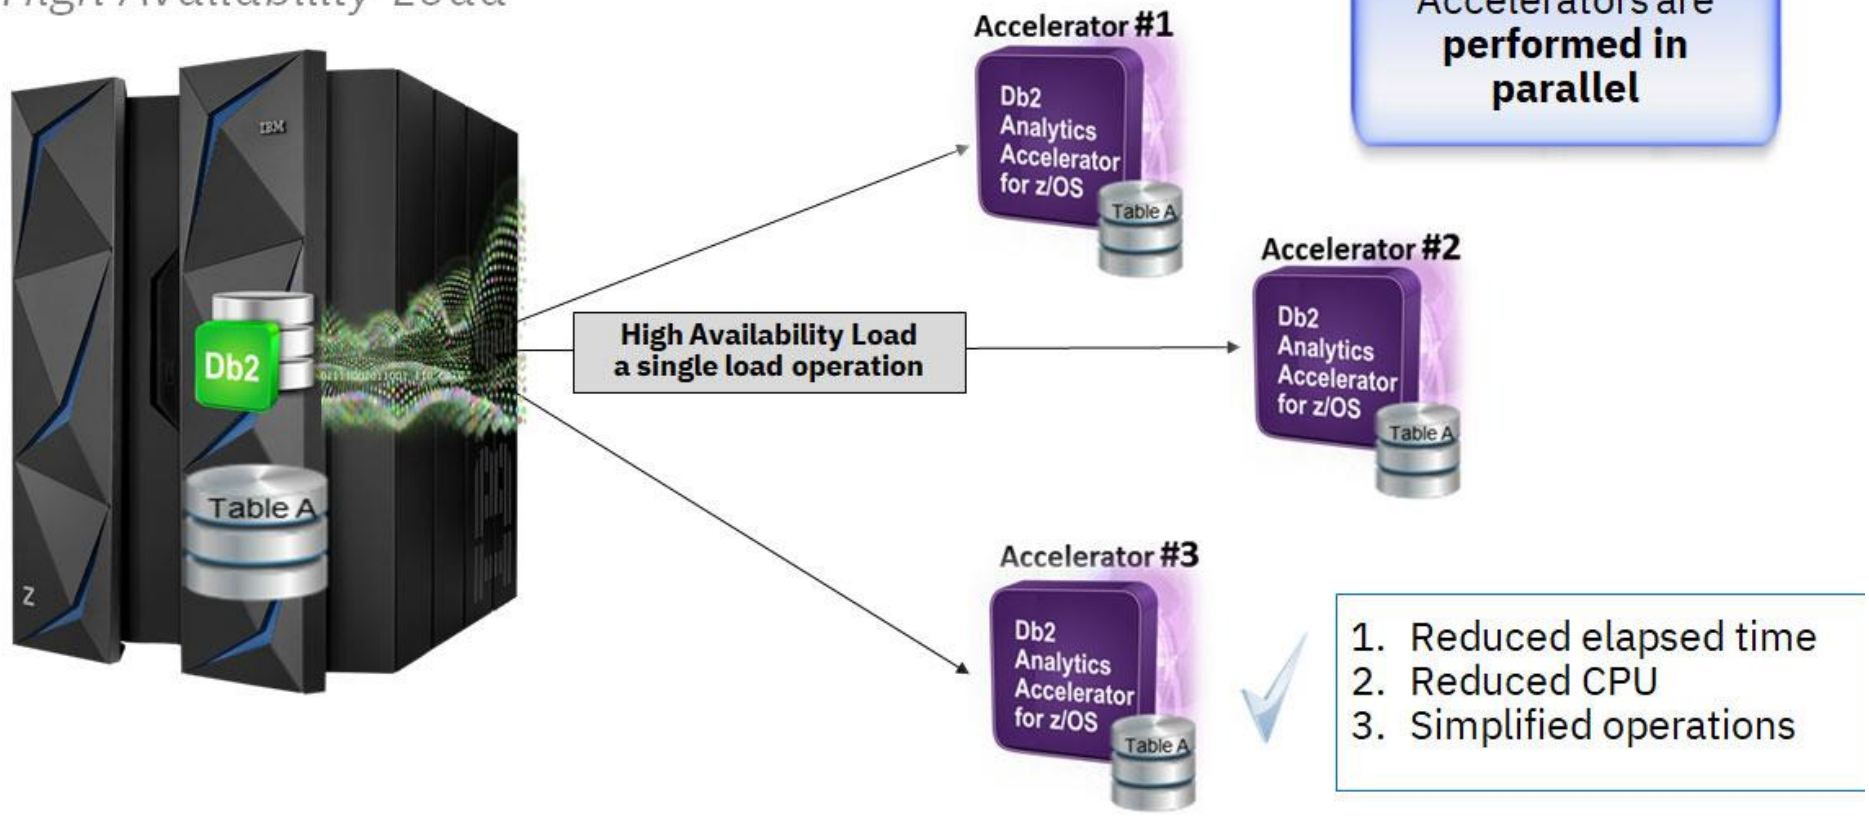

High Performance Storage Saver (HPSS)

- 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

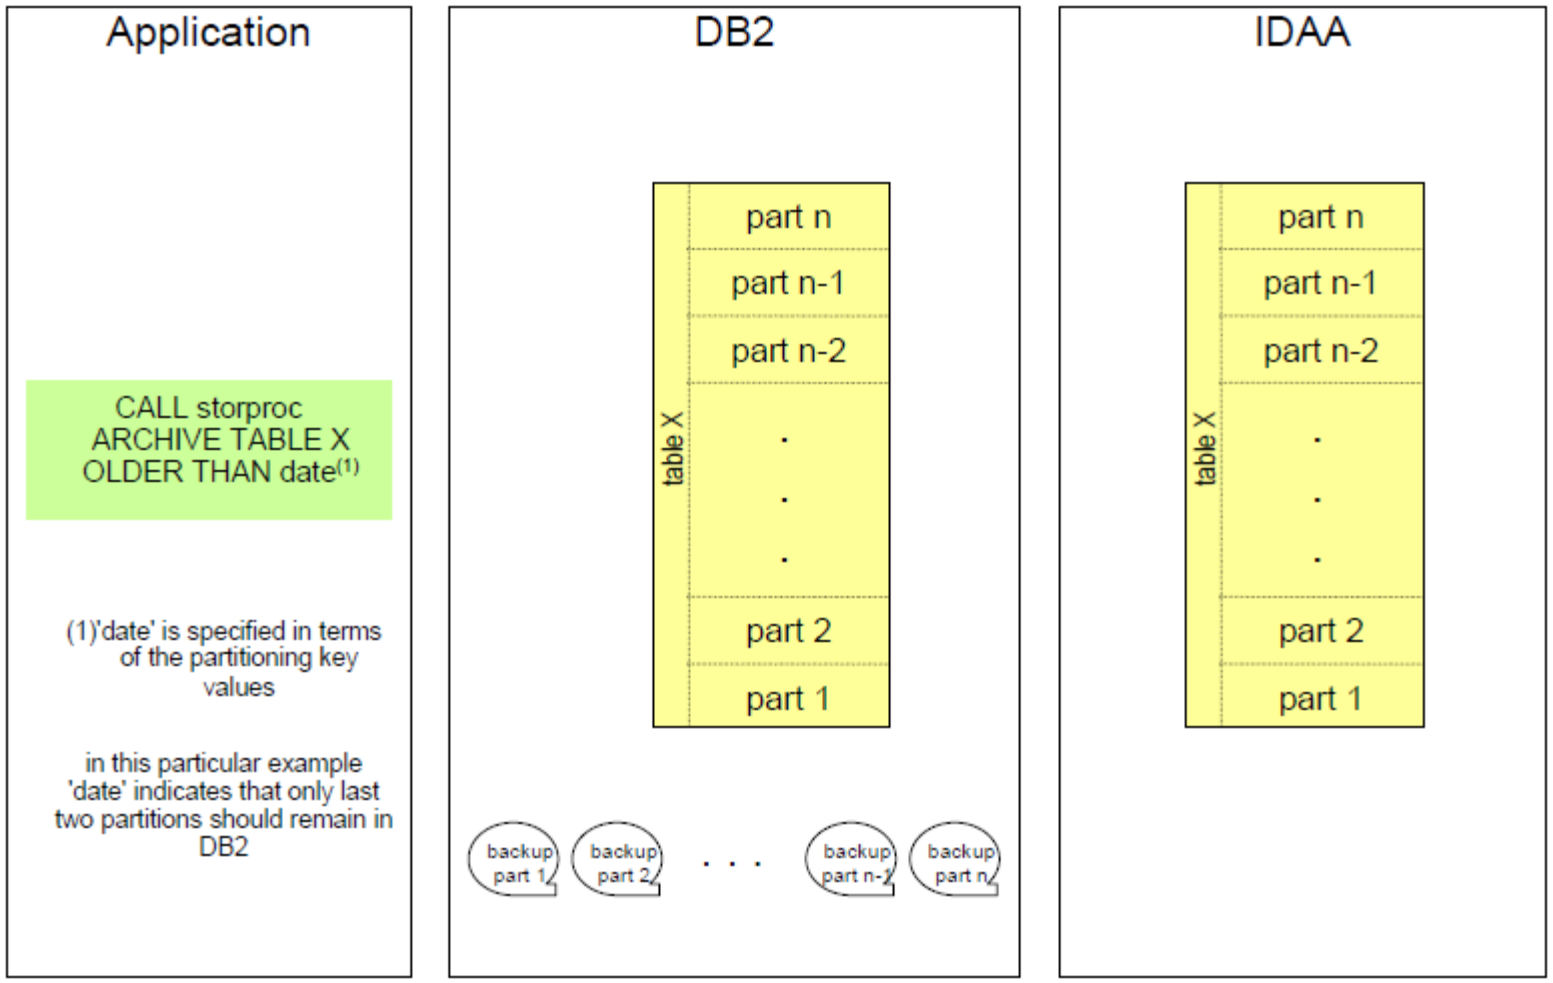
Application
CALL storproc ARCHIVE TABLE X OLDER THAN date(1)
(1)'date' is specified in terms of the partitioning key values
in this particular example 'date' indicates that only last two partitions should remain in DB2
DB2
table X
part n
part n-1
part n-2
part 2
part 1
backup part 1
backup part 2
backup part n-1
backup part n
IDAA
table X
part n
part n-1
part n-2
part 2
part 1

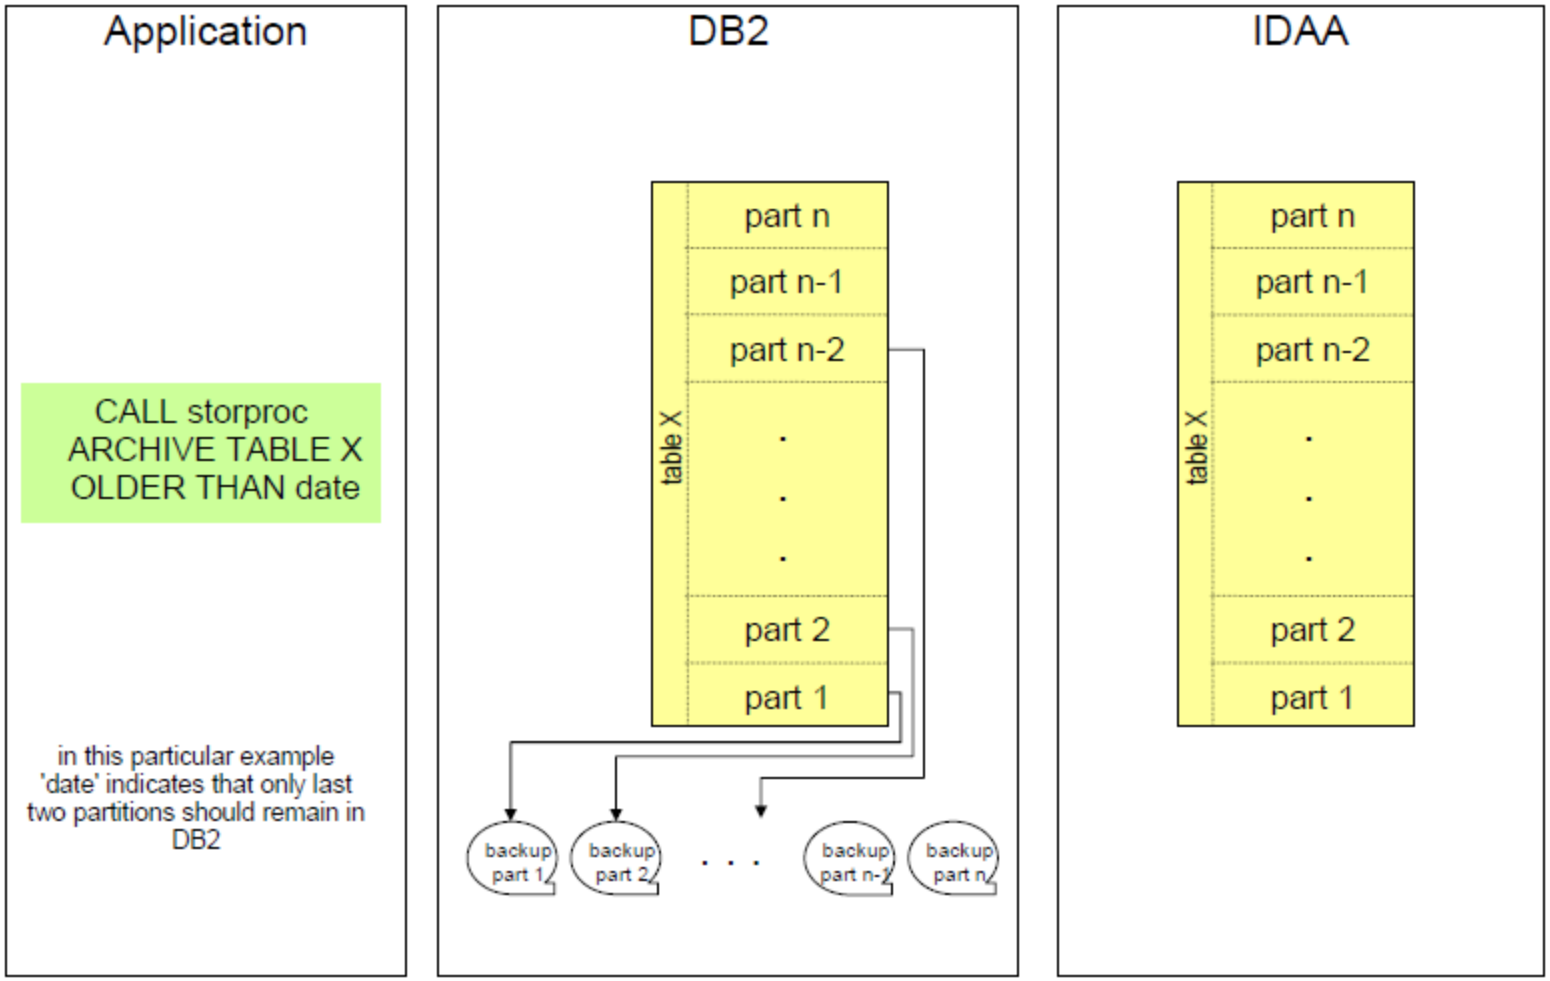
Application
CALL storproc
ARCHIVE TABLE X
OLDER THAN date
in this particular example
'date' indicates that only last
two partitions should remain in
DB2
DB2
table X
part n
part n-1
part n-2
part 2
part 1
backup part 1
backup part 2
backup part n-2
backup part n
IDAA
table X
part n
part n-1
part n-2
part 2
part 1

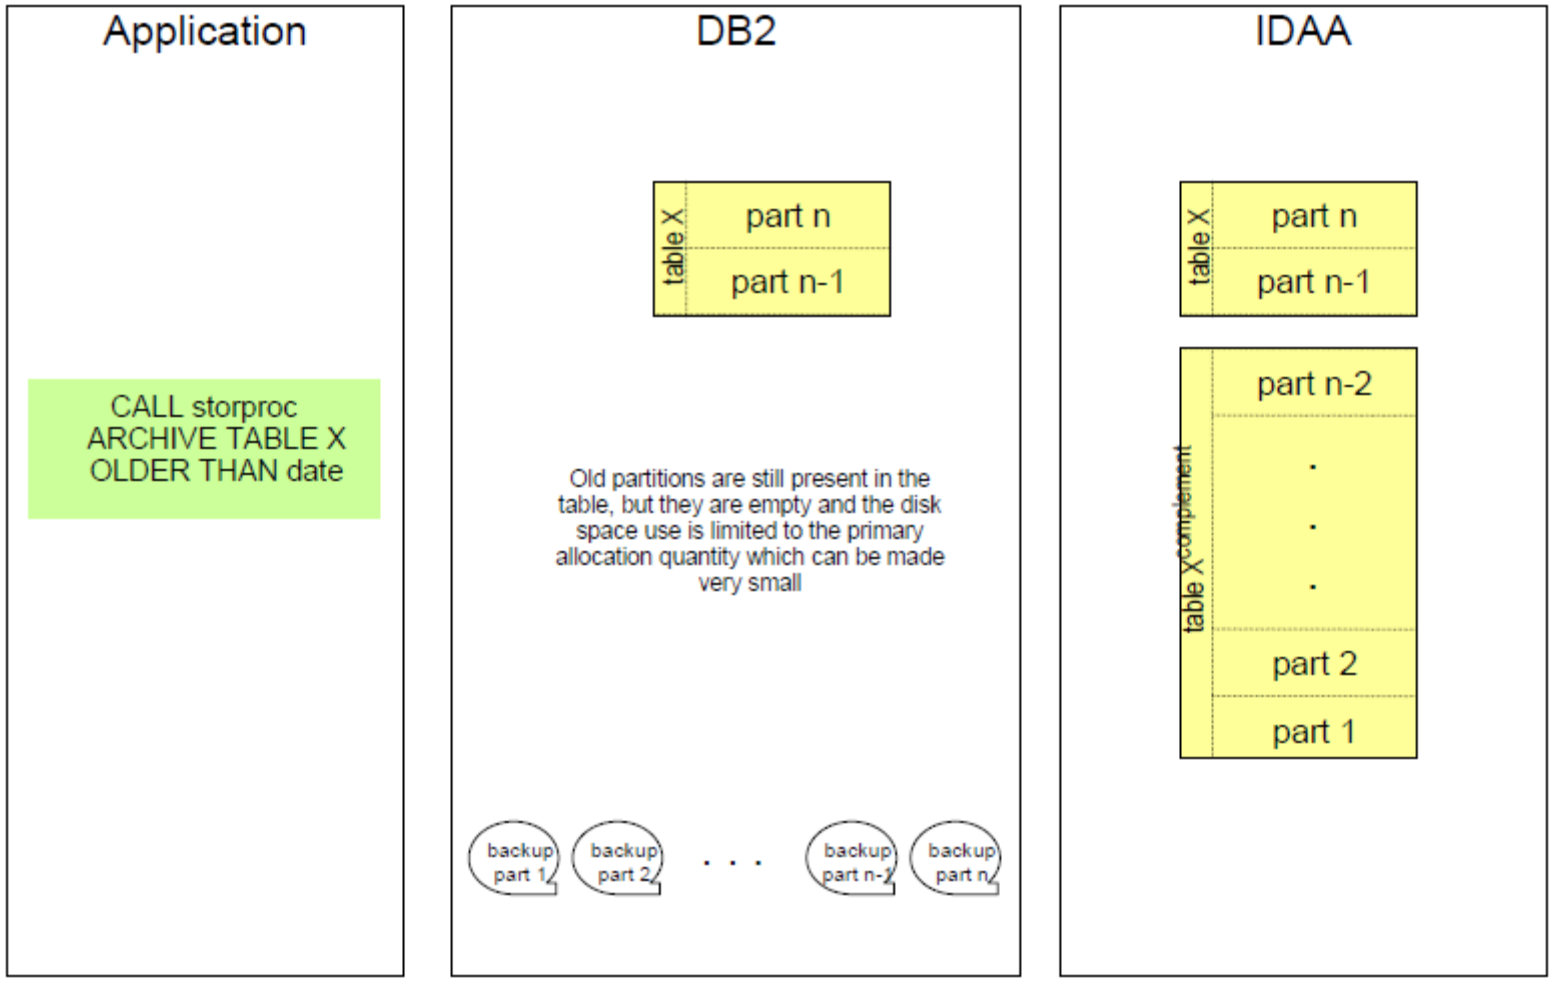
Application
CALL storproc
ARCHIVE TABLE X
OLDER THAN date
DB2
table X
part n
part n-1
Old partitions are still present in the table, but they are empty and the disk space use is limited to the primary allocation quantity which can be made very small
backup part 1
backup part 2
backup part n-1
backup part n
IDAA
table X
part n
part n-1
table X complement
part n-2
part 2
part 1

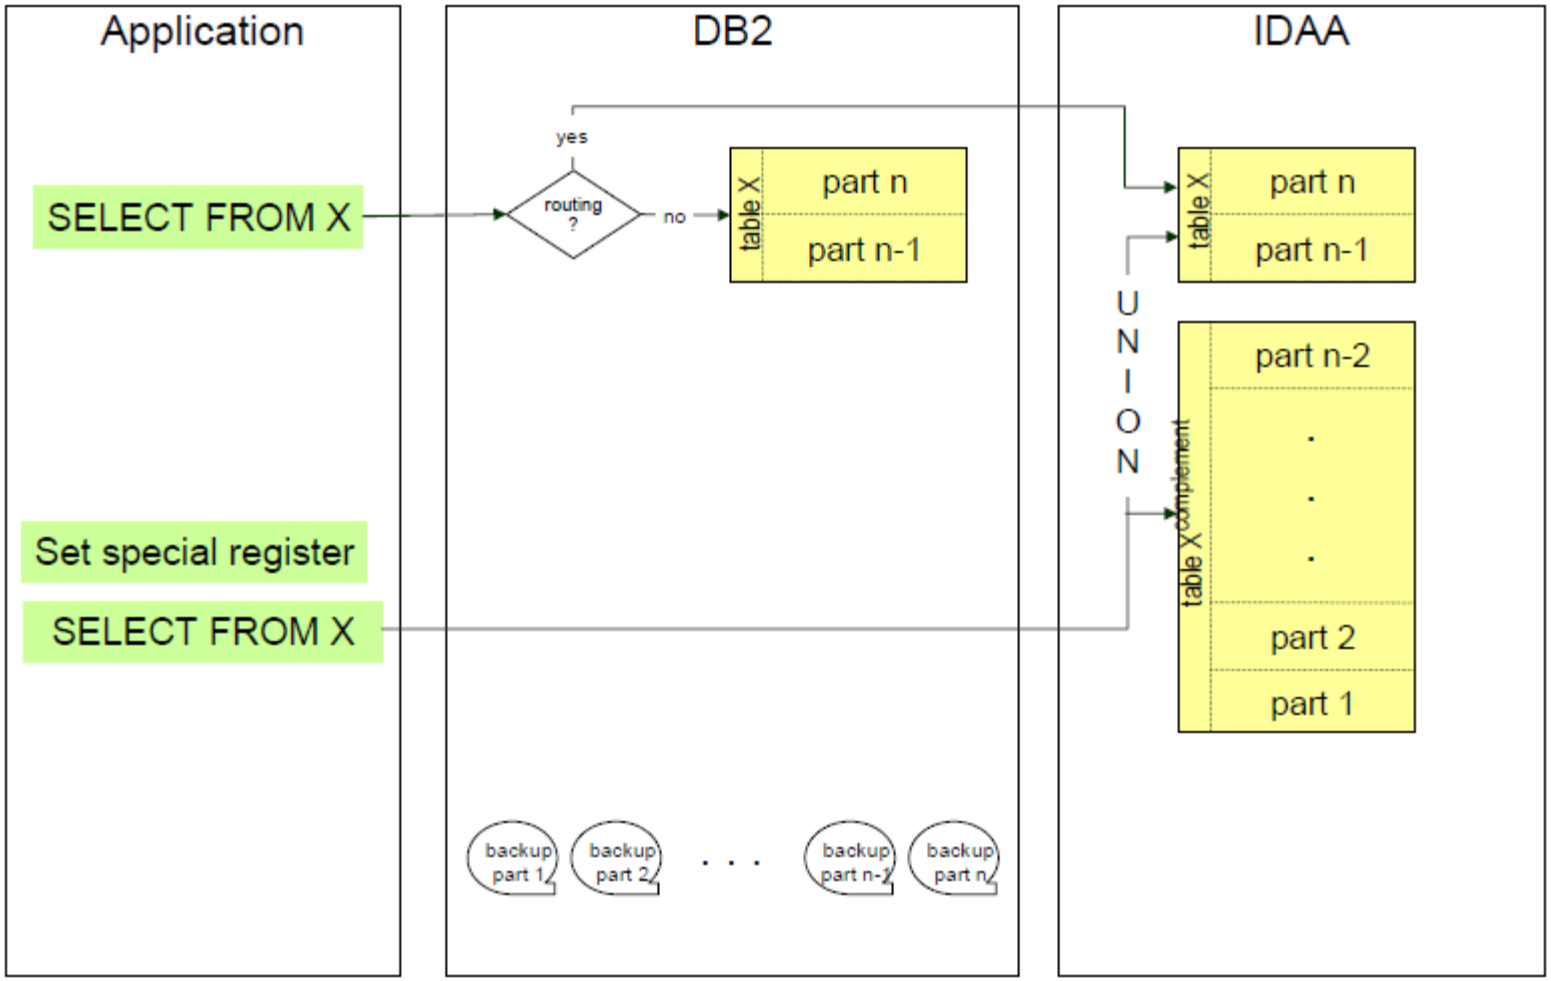
Application
DB2
IDAA
SELECT FROM X
yes
routing ?
no
table X
part n
part n-1
table X
part n
part n-1
U
N
I
O
N
part n-2
.
.
.
part 2
part 1
table X complement
Set special register
SELECT FROM X
backup part 1
backup part 2
. . .
backup part n-2
backup part n

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

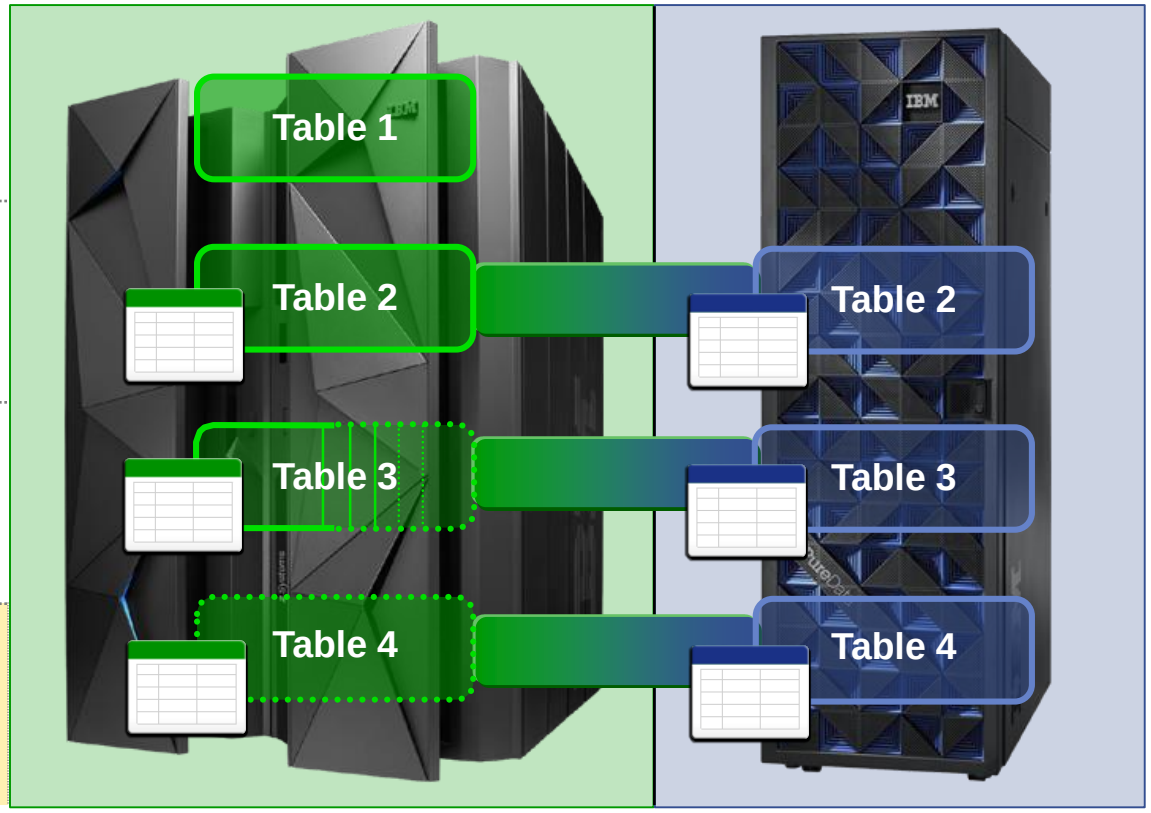

Transaction Processing Systems (OLTP)

Analytics

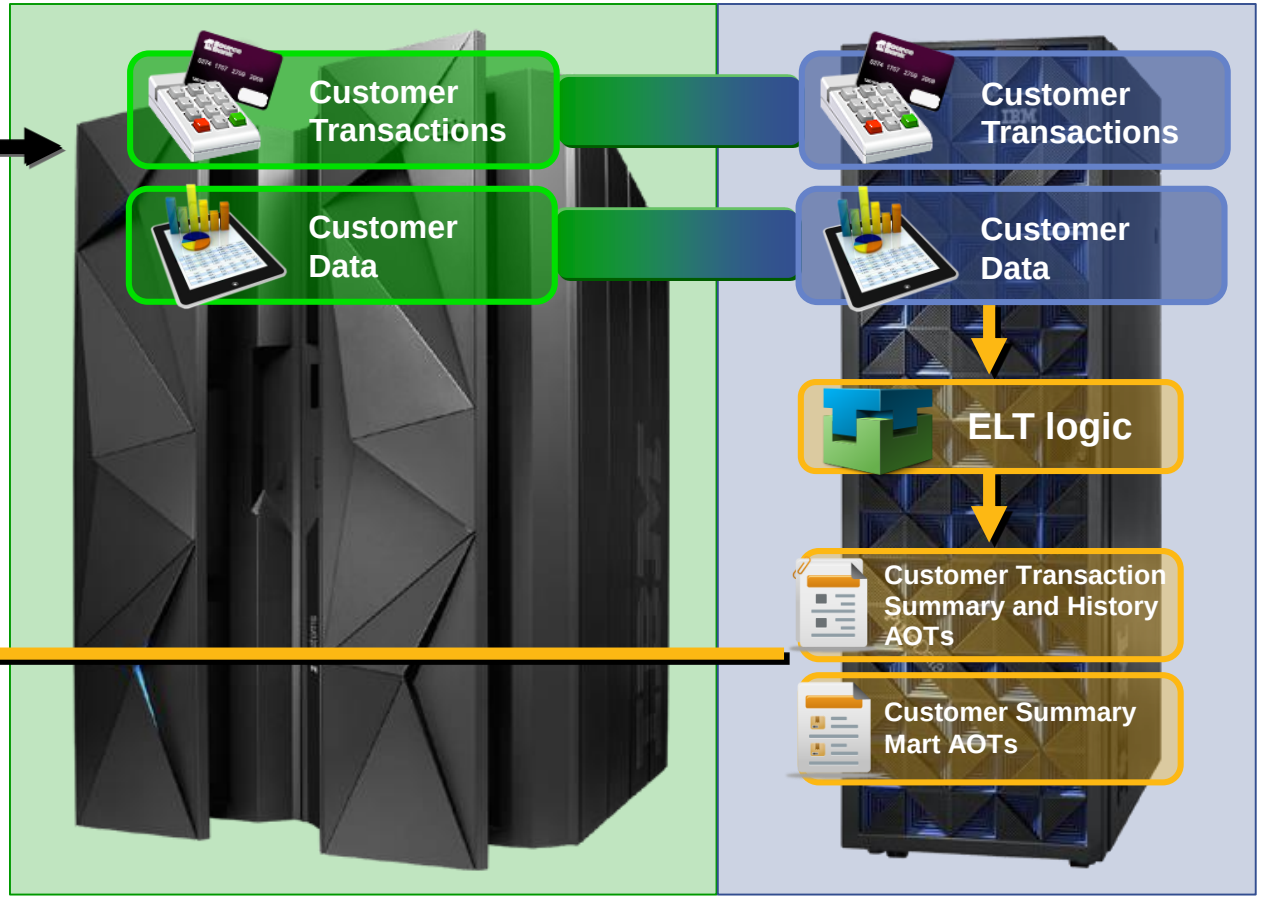

Advantages:

- Simpler to manage
- Better performance and reduced latency

Data for transactional and analytical processing

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**

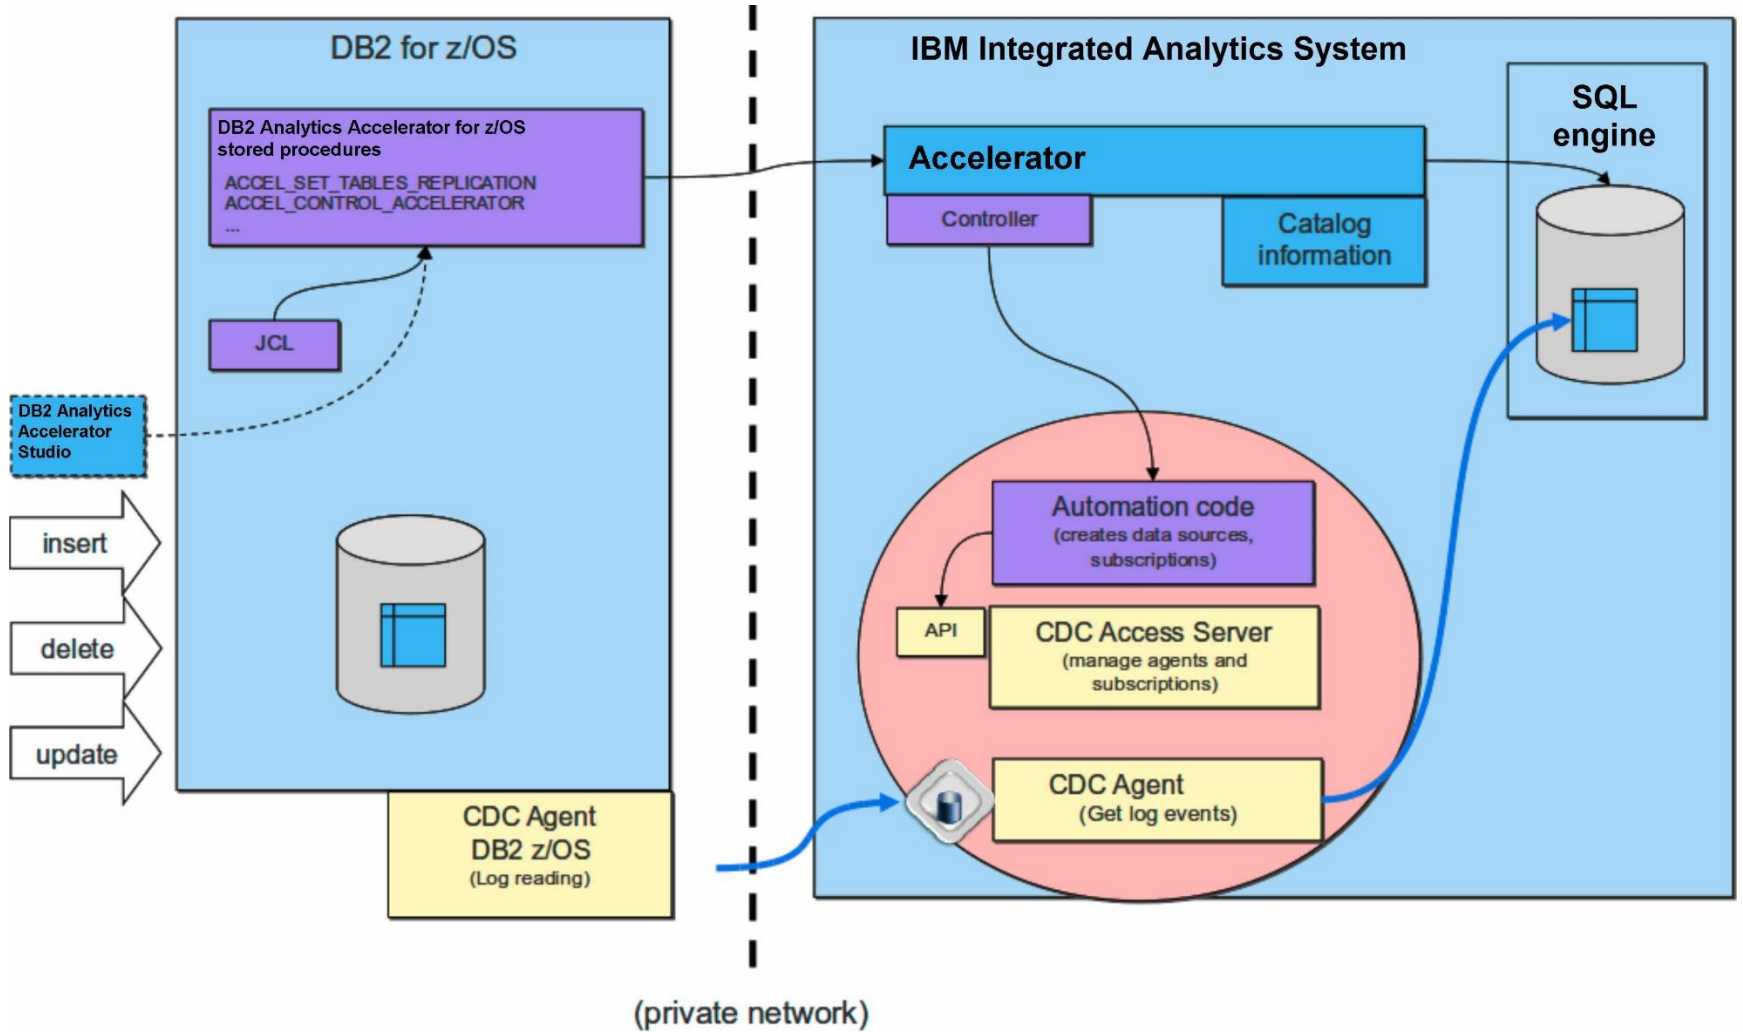
DB2 for z/OS
DB2 Analytics Accelerator for z/OS stored procedures
ACCEL_SET_TABLES_REPLICATION
ACCEL_CONTROL_ACCELERATOR
...
JCL
DB2 Analytics Accelerator Studio
insert
delete
update
CDC Agent DB2 z/OS
(Log reading)
IBM Integrated Analytics System
Accelerator
Controller
Catalog information
SQL engine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

Integrated Synchronization

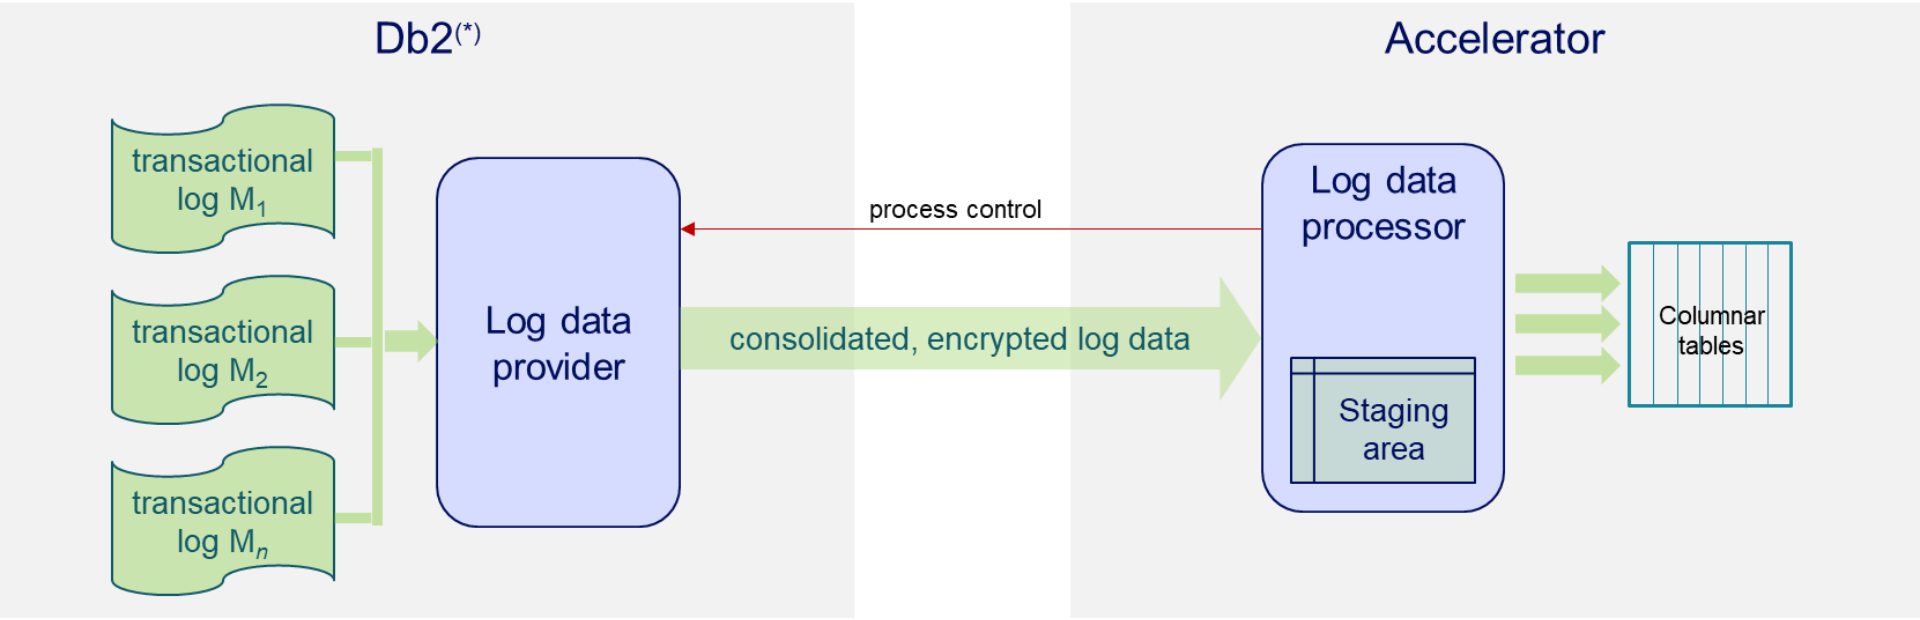

Integrated Synchronization (cont)

- **Log data provider** is a newly developed, internal Db2 for z/OS component
- Adheres to Db2 life-cycle management resulting in simplified installation, packaging, administration, upgrade, support, … as compared to external data capture tools
- Fully zIIP enabled - MSU savings potential
- Streamlined design resulting in reduced CPU usage and higher throughput (see next chart)
- Supports transactional consistency protocol that guarantees queries executed by IDAA return most recently committed data: the cornerstone of application transparency and HTAP

Integrated Synchronization (cont)

- **Log data processor** is a newly developed, internal accelerator component
- Adheres to the accelerator life-cycle management resulting in simplified installation, packaging, administration, upgrade, support, … as compared to external data capture tools
- Custom-built and optimized resulting in higher throughput and lower latency
- Numerous performance optimizations such as 'bulk vs. trickle' mode and optimistic apply
- Significant enhancements in DB2 Warehouse insert/update/delete performance
- Supports transactional consistency protocol that guarantees queries executed by IDAA return most recently committed data: the cornerstone of application transparency and HTAP

Hybrid Transaction / Analytic Processing (HTAP) Performance

- HTAP scalability with number of concurrent queries
  - 4M rows updated, commit then starts 30 concurrent queries
    All waited for ~6s and finished almost at same time
  - Single query wait time is also ~6s
  - No more impact of HTAP queries on replication throughput, huge HTAP scalability improvement

- HTAP with Change Data Capture (CDC) vs Integrated Synchronization
  - Average wait time dropped from 51 seconds to 6 seconds across a total of 89 queries
  - Average latency of the workload dropped from 50 seconds to 3

Resources

- Redbooks
  - "Enabling Real-time Analytics on IBM z Systems Platform" SG24-8272
  - "Managing IBM Db2 Analytics Accelerator by using IBM Data Server Manager" (Web Doc)
- www.thefillmoregroup.com/blog

- Release Notes v7.1.4
  - https://www-01.ibm.com/support/docview.wss?uid=ibm10732149
- Release Notes v7.1.5
  - https://www-01.ibm.com/support/docview.wss?uid=ibm10743339
- Release Notes v7.1.6
  - https://www-01.ibm.com/support/docview.wss?uid=ibm10795428
- Release Notes v7.1.7
  - http://www-01.ibm.com/support/docview.wss?uid=ibm10876158
- Prerequisites and maintenance
  - http://www-01.ibm.com/support/docview.wss?uid=swg27050440
- Knowledge Center
  - https://www.ibm.com/support/knowledgecenter/SS4LQ8_7.1.0/com.ibm.datatools.aqt.doc/idaa_kc_welcome.html

THE FILLMORE GROUP
Relational Database Solutions

Thank you

Kim May, Vice President Business Development

Frank Fillmore, President

www.thefillmoregroup.com

Premier Business Partner IBM