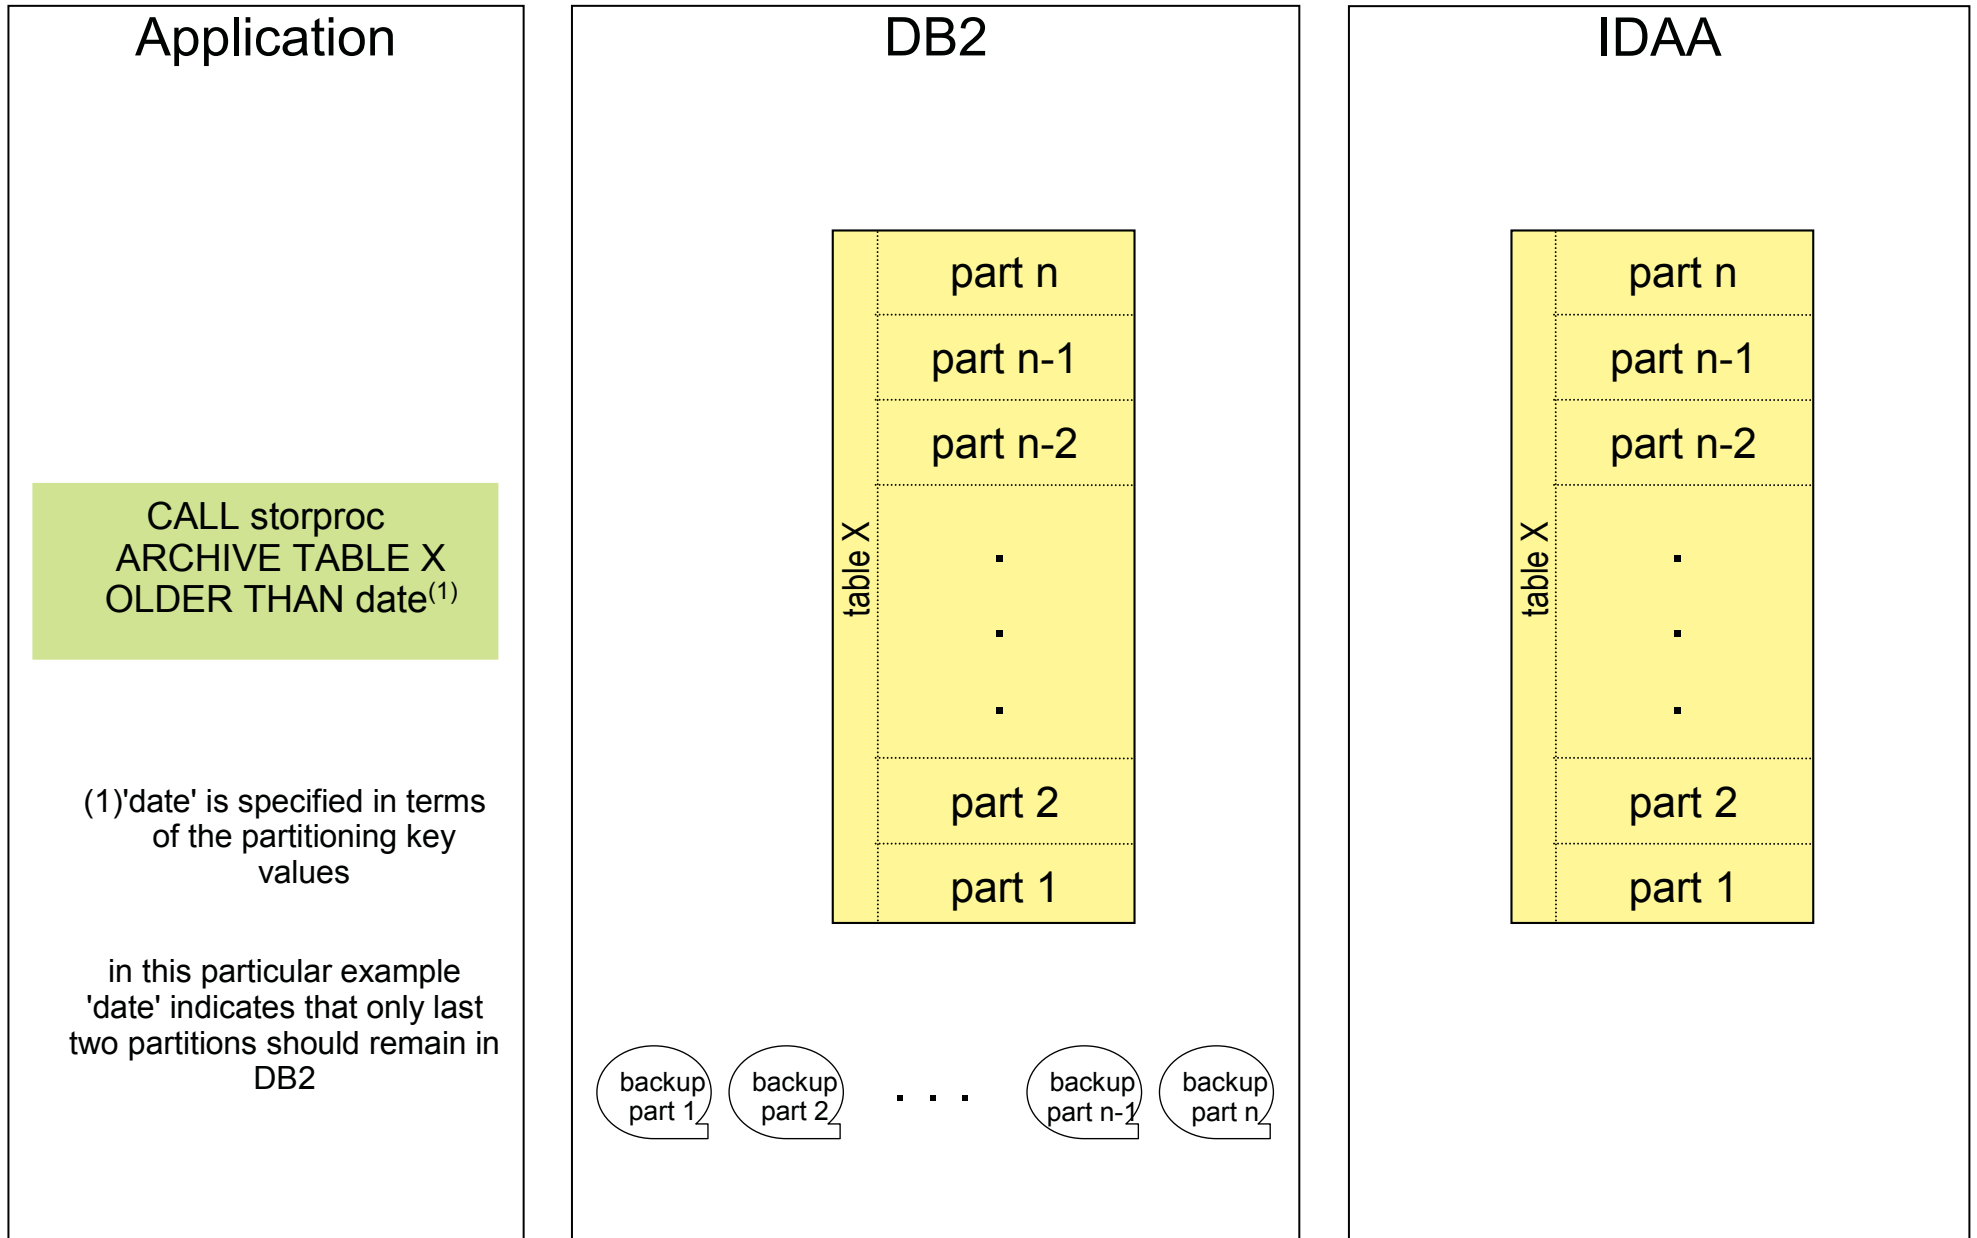


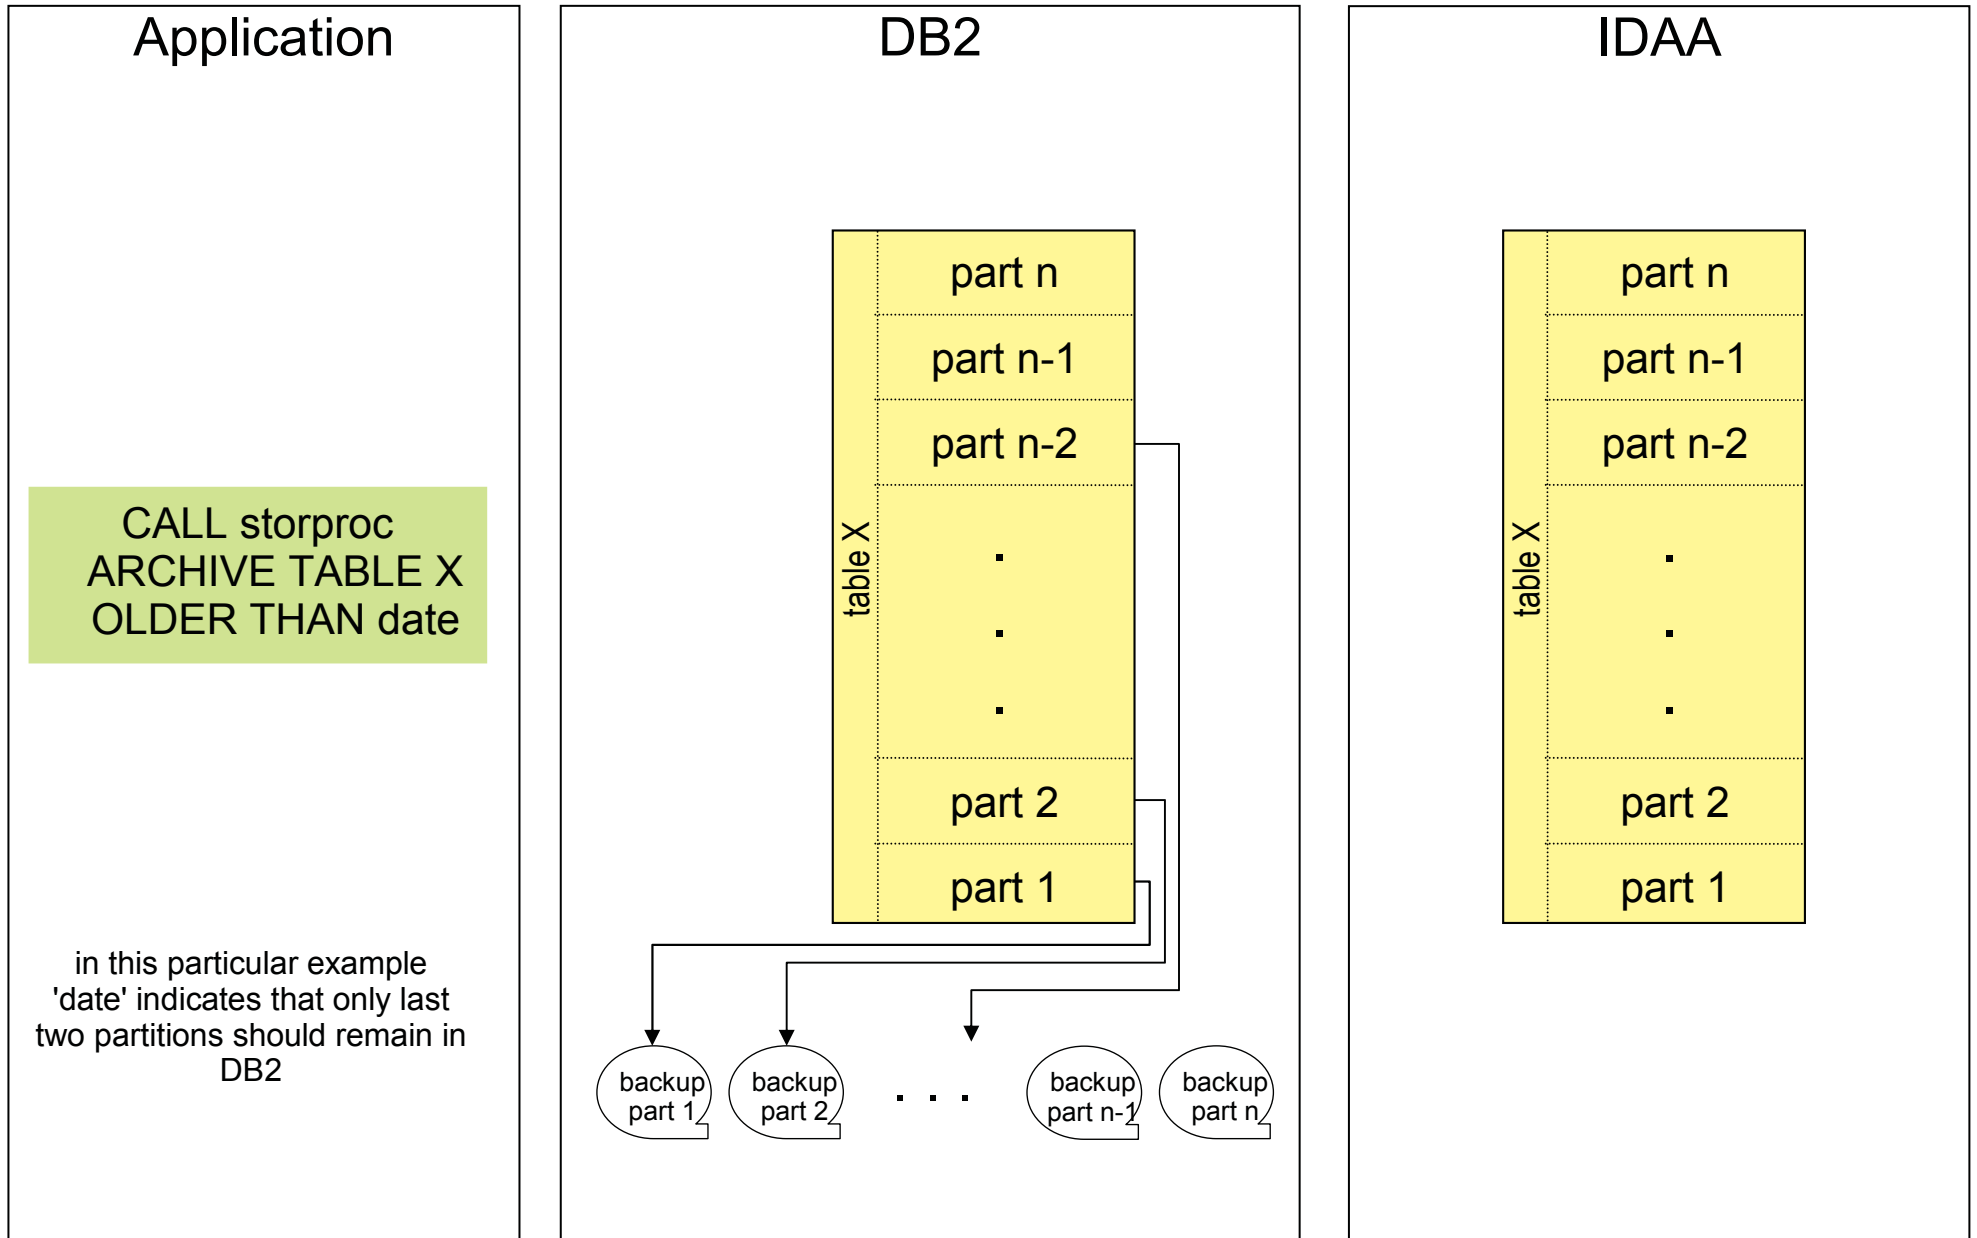
What's new in IDAA V3 ?

- Overview of V3 Enhancements
- Details on V3 “Incremental Update” Feature
- Details on V3 “Detect staleness of data” Feature
- Details on V3 WLM support
- Details on V3 “High Performance Storage Saver”

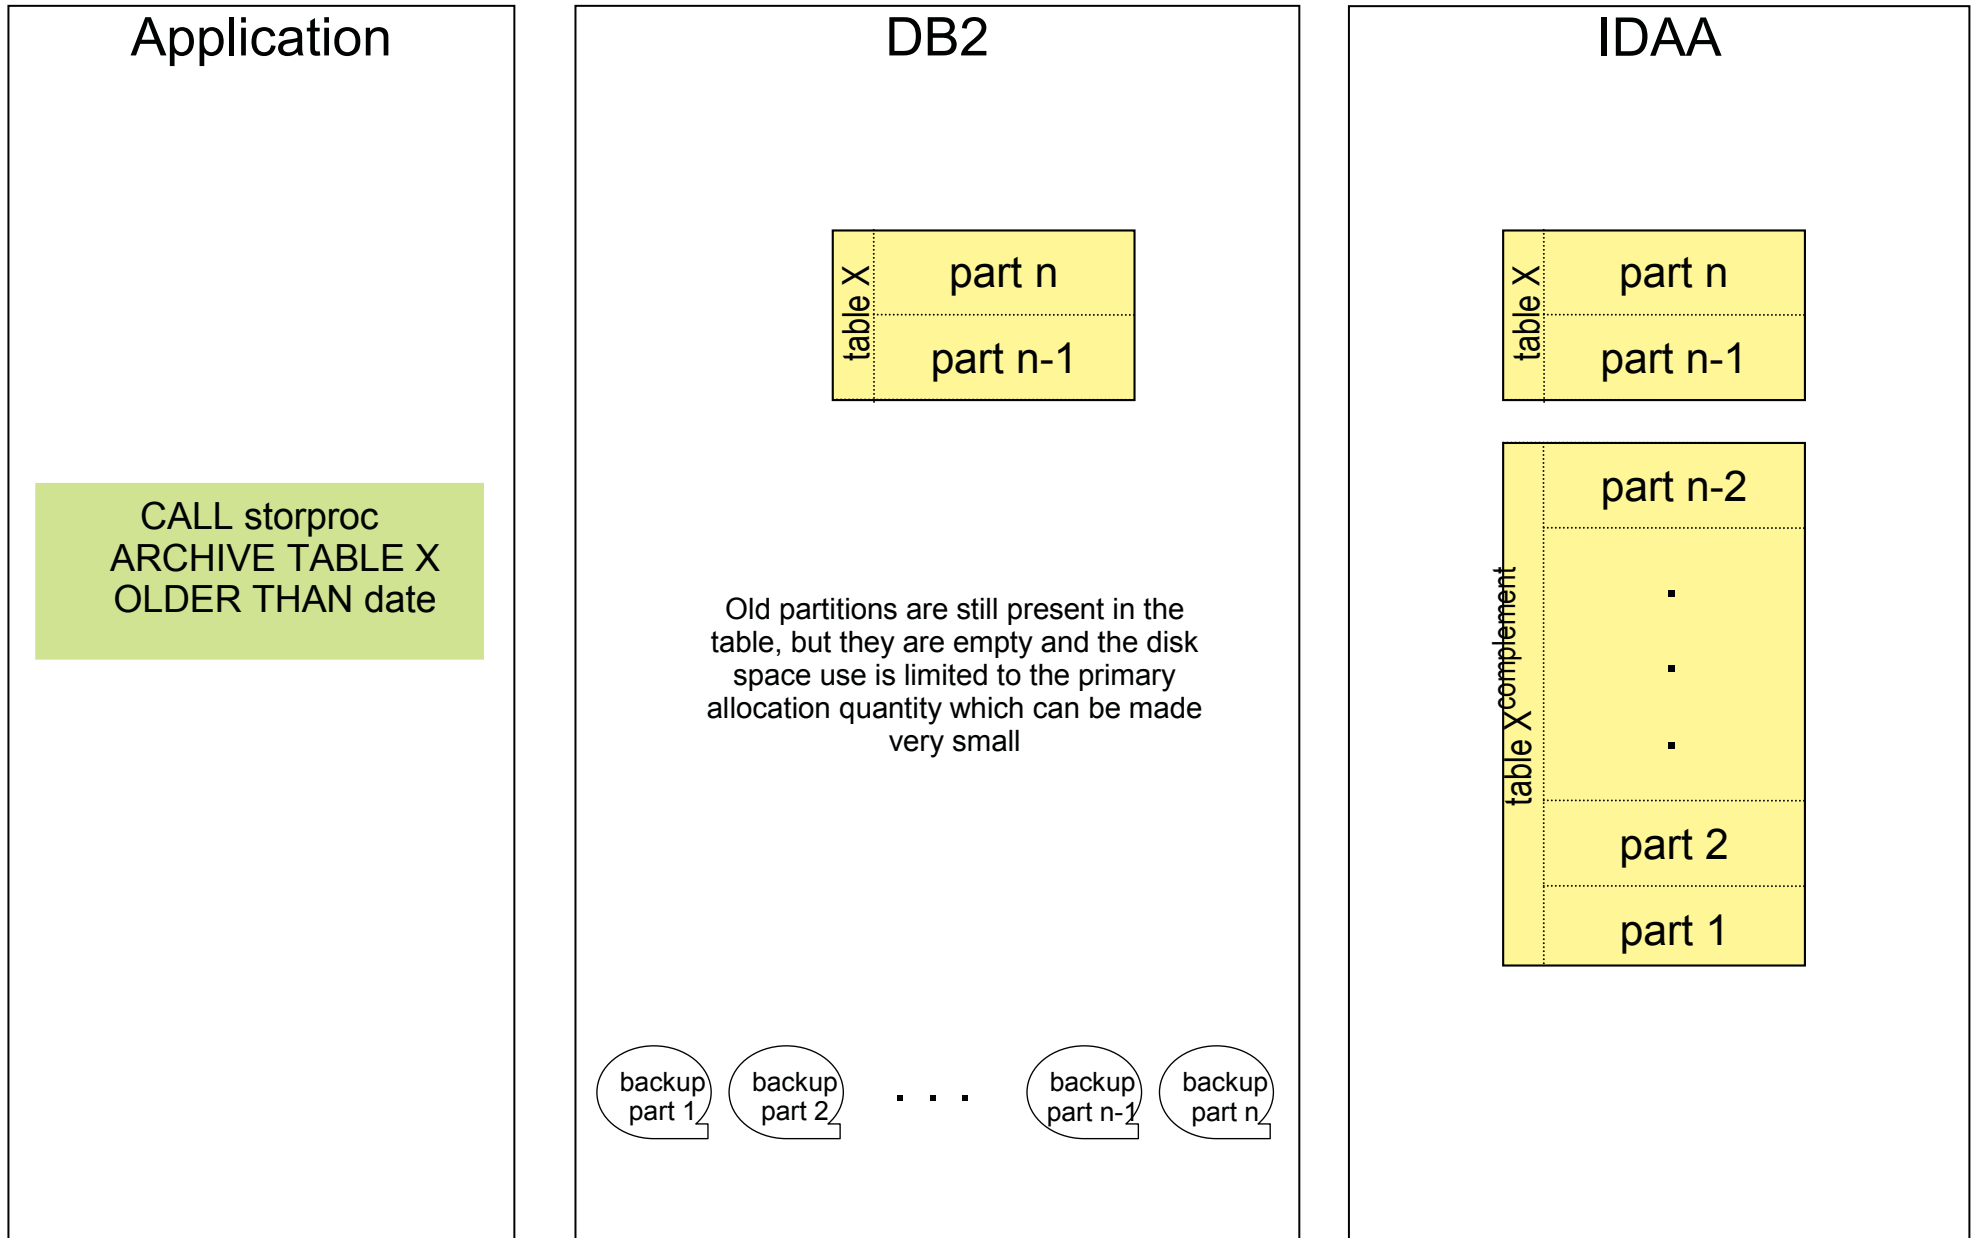
Supplied Stored Procedure encapsulates Archiving Procedure



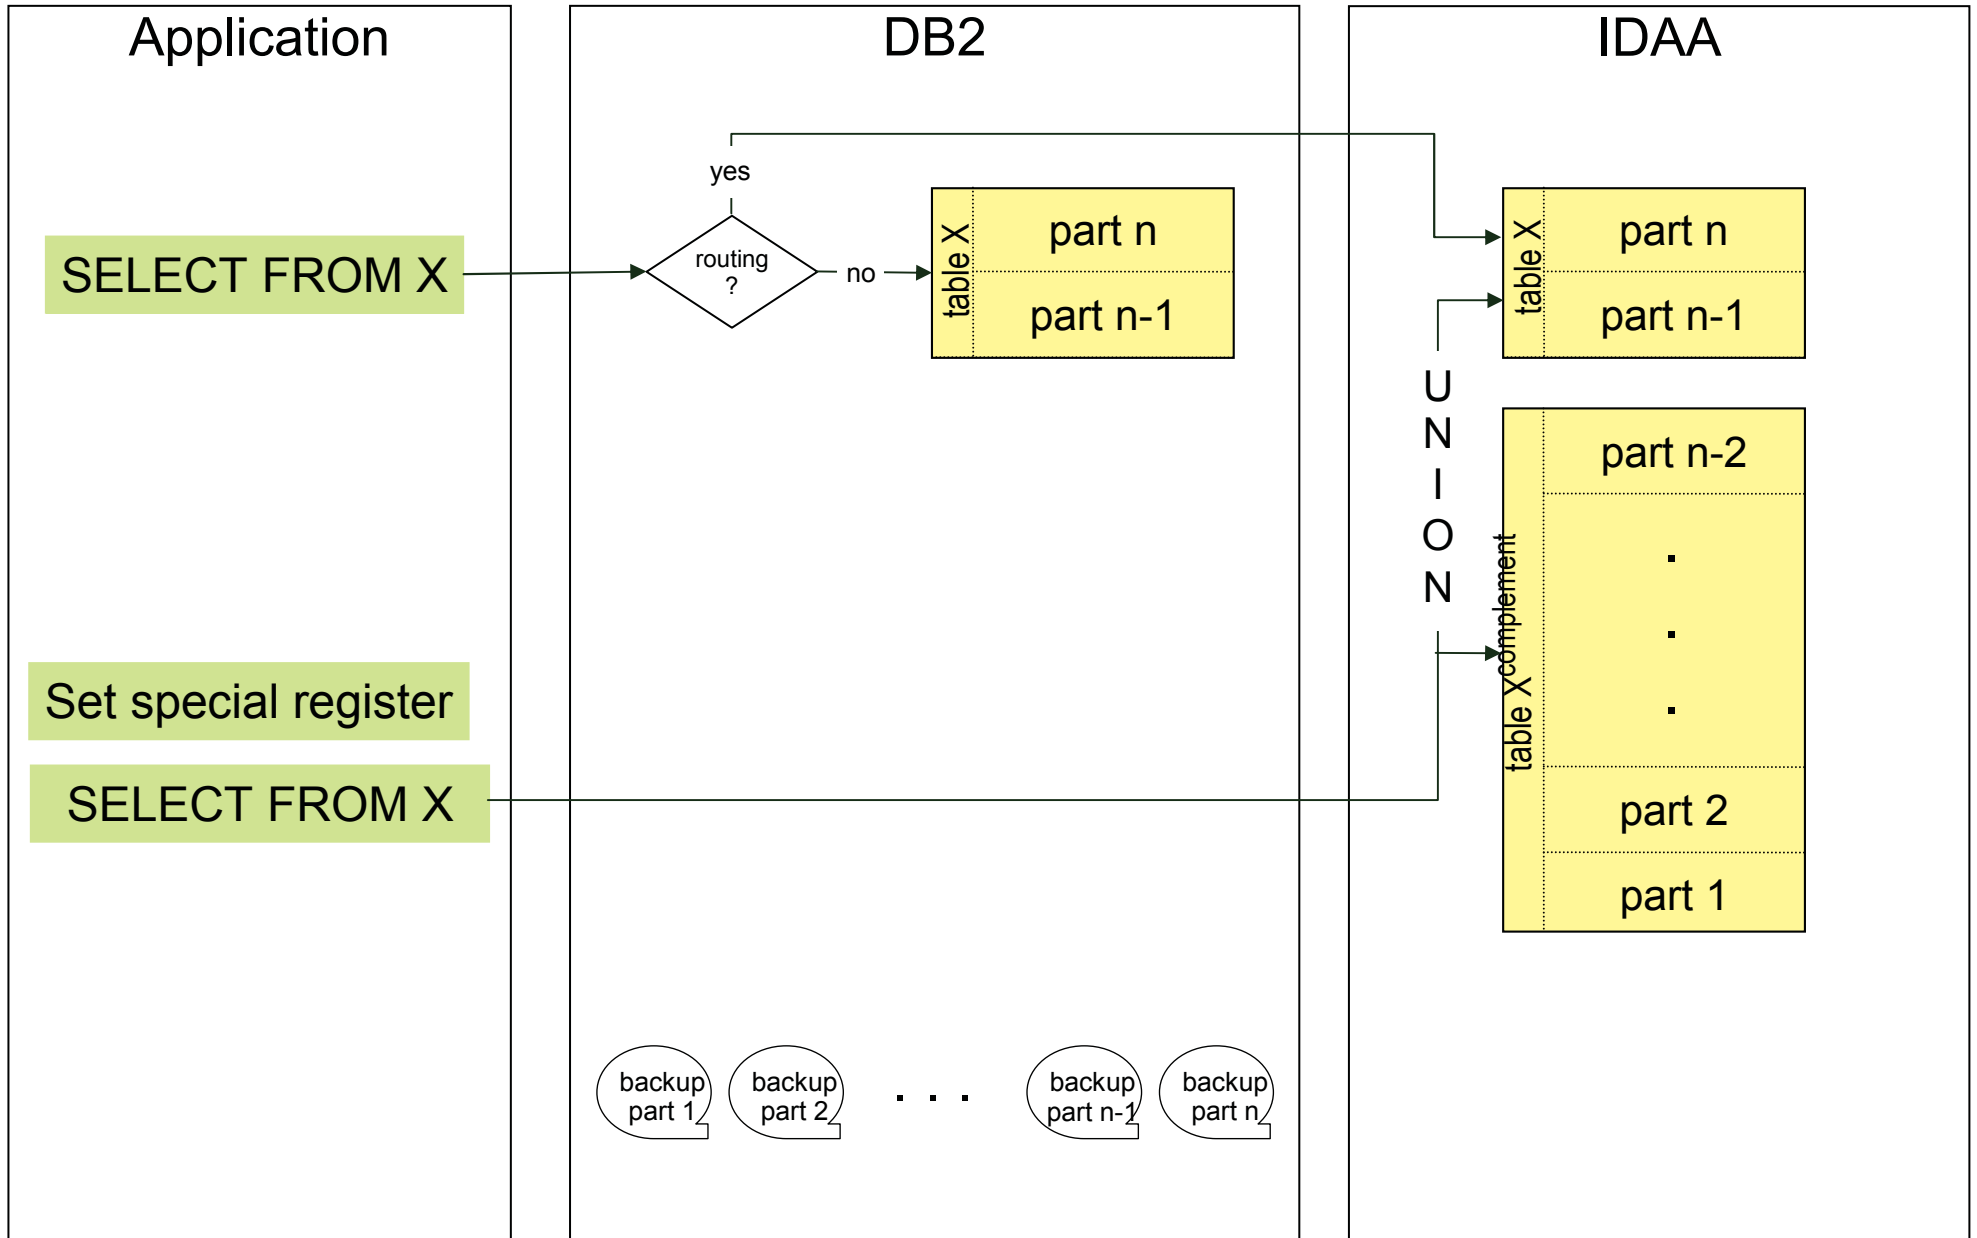
Partitions to be Archived are first backed up



Old Partitions are deleted from DB2 and Table X is split within IDAA



Applications have transparent access (no SQL statement changes needed) to the Table



Key Elements of the Implementation Approach

- Older partitions are moved to IDAA and their data no longer exists in DB2
 - DB2 is still solely responsible for the recovery and maintains all backups (copy images)
- The most recent partitions exist in both DB2 and IDAA
 - They are synchronized by existing means
 - Partition refresh or incremental update (replication-based propagation of changes)
 - Having the most recent partition in both, DB2 and IDAA, provides IDAA-driven performance acceleration for analytical queries that access most recent partitions only
- The data move process is encapsulated in a stored procedure
 - The stored procedure can be invoked directly or via IDAA Studio
- The SQL statements do not change
 - The fact that some partitions have been moved to IDAA is transparent
 - By default, queries access only the data from the most recent partitions
 - The queries can be executed in DB2 or IDAA based on the standard routing criteria
 - If all the data need to be accessed, one of the following mechanisms is used:
 - Setting a zparm which activates the 'all data' scope for the DB2 subsystem/data sharing group. This way, none of the applications need to be changed (but this setting has global impact).
 - Setting a special register "CURRENT GET_ACCEL_ARCHIVE", which allows switching between the 'all data' scope and the 'most recent data' scope at any time. This way the application can use both scopes within the same execution at choose scope at SQL statement level.

Graphical User Interface for HPSS Function

▼ Tables (8 of 8 loaded / 0 of 8 enabled for acceleration)

Name like: type filter text

Name	Size	Rows	Acceleration	Last Load	Moved Partitions	Distribution Key	Sk
TPCH	352 ...	-	0 of 8	8 of 8	1 of 8		
CUSTOMER	44 MB	4242424	Disabled	Initial load pending	-	Random	0.4
LINEITEM	44 MB	4242424	Disabled	Initial load pending	-	Random	0.4
NATION	44 MB	4242424	Disabled	Initial load pending	-	Random	0.4
ORDERS	44 MB	4242424	Disabled	9/14/10 10:15 PM	315 MB / 2000000 rows	Random	0.4
PART	44 MB	4242424	Disabled	Initial load pending	-	Random	0.4
PARTSUPP	44 MB	4242424	Disabled	Initial load pending	-	Random	0.4
REGION	44 MB	4242424	Disabled	Initial load pending	-	Random	0.4
SUPPLIER	44 MB	4242424	Disabled	Initial load pending	-	Random	0.4

Move to Accelerator...
 Restore in DB2...

Moving Partitions to IDAA via GUI

Move to Accelerator

Move Selected Partitions

For each table specify the partitions to be moved to the accelerator. Moved partitions will only be available for queries running on the accelerator and will be removed from DB2.

Table

- TPCH.LINITEM
- TPCH.NATION
- TPCH.ORDERS**

Partitions to Move for Selected Table

Table: TPC.ORDERS **Partitioned by:** O_ORDERDATE **Status:** partitions already moved

Do not move
 Move all partitions up to and including the following limit key:
 2011-12-31
 Manually select partitions

Selected Partitions:

Logical Partition	Ending At	Already Moved	DB2 Size
<input checked="" type="checkbox"/> Partition 1	2011-10-31	Yes	0.1 GB
<input checked="" type="checkbox"/> Partition 2	2011-11-30	Yes	0.1 GB
<input checked="" type="checkbox"/> Partition 3	2011-12-31	Yes	0.1 GB
<input type="checkbox"/> Partition 4	2012-01-01	No	0.1 GB
<input type="checkbox"/> Partition 5	2012-01-02	No	0.1 GB
<input type="checkbox"/> Partition 6	2012-01-03	No	0.1 GB
<input type="checkbox"/> Partition 7	2012-01-04	No	0.1 GB
<input type="checkbox"/> Partition 8	2012-01-05	No	0.1 GB
<input type="checkbox"/> Partition 9	2012-01-06	No	0.1 GB
<input type="checkbox"/> Partition 10	2012-01-07	No	0.1 GB
<input type="checkbox"/> Partition 11	2012-01-08	No	0.1 GB

Amount of data to move: 0.0 B (0 partitions)

Note: This accelerator client must stay connected until the partitions are fully moved. This can take a long time depending on the amount of data to move, for examples several minutes or hours

OK Cancel

Limitations on Moving Partitions to IDAA

- **Partitions of a given table must *always* be moved to the *same* accelerator**
 - **Currently there are no DR capabilities**, which would allow moving data of the same table to multiple, different accelerators
 - Queries accessing moved data must find all moved data on the same accelerator
 - (Plans for improvements to be delivered in 2013)
- **Initially, there is no ACCEL_RESTORE_ARCHIVE_TABLES stored procedure**
 - Manual use of DB2 RECOVER utility is required for all partitions of the table
 - RECOVER
 - MODIFY RECOVERY
 - (Plans for improvements to be delivered in 2013)
- **Only supported for tables that are *not* parent in a foreign key relationship**
- **Tables must be range-partition**
 - Table-controlled or index-controlled is both acceptable
 - Non-partitioned tables or tables partitioned by growth are not supported