

Tips and tricks to get the most out of DB2 V11 with IBM DB2 Utilities and Tools



Neale Armstrong

System Z Information Management

Technical Sales for IBM Australia

Neale.armstrong@au1.ibm.com

Agenda

- **IBM's DB2 Tools Portfolio Overview**
- **Using IBM DB2 Tools to Streamline V2V Migration**
- **DB2 V11 Utilities Enhancements**
- **Positioning of DB2 Archiving Options**
- **IDAA Operational Efficiencies**

DB2 Tools Portfolio

DB2 Utilities Solution Pack

DB2 Automation Tool
 DB2 HPU
 DB2 Sort
 DB2 Utilities Enhancement Tool

Super-charge the IBM utilities

DB2 Performance Solution Pack

DB2 SQL Performance Analyzer
 Tivoli Omegamon XE for DB2 PE
 Optim Query Workload Tuner
 DB2 Query Monitor

Master the performance lifecycle

DB2 Fast Copy Solution Pack

DB2 Cloning Tool
 DB2 Recovery Expert

Superior avail & cost for copy,
 backup & recovery operations

DB2 Administration Solution Pack

DB2 Administration Tool
 DB2 Object Comparison Tool
 DB2 Table Editor
 Optim Configuration Manager

Manage objects & schema

QMF

Infosphere Data
 Replication

IDAA V4.1
 IDAA Loader

DB2 Application
 Management
 Tools

Governance Solutions

- Audit
- Test Data Management
- Data Privacy
- Data Archiving

Agenda

- **IBM's DB2 Tools Portfolio Overview**
- **Using IBM DB2 Tools to Streamline V2V Migration**
- **DB2 V11 Utilities Enhancements**
- **Positioning of DB2 Archiving Options**
- **IDAA Operational Efficiencies**

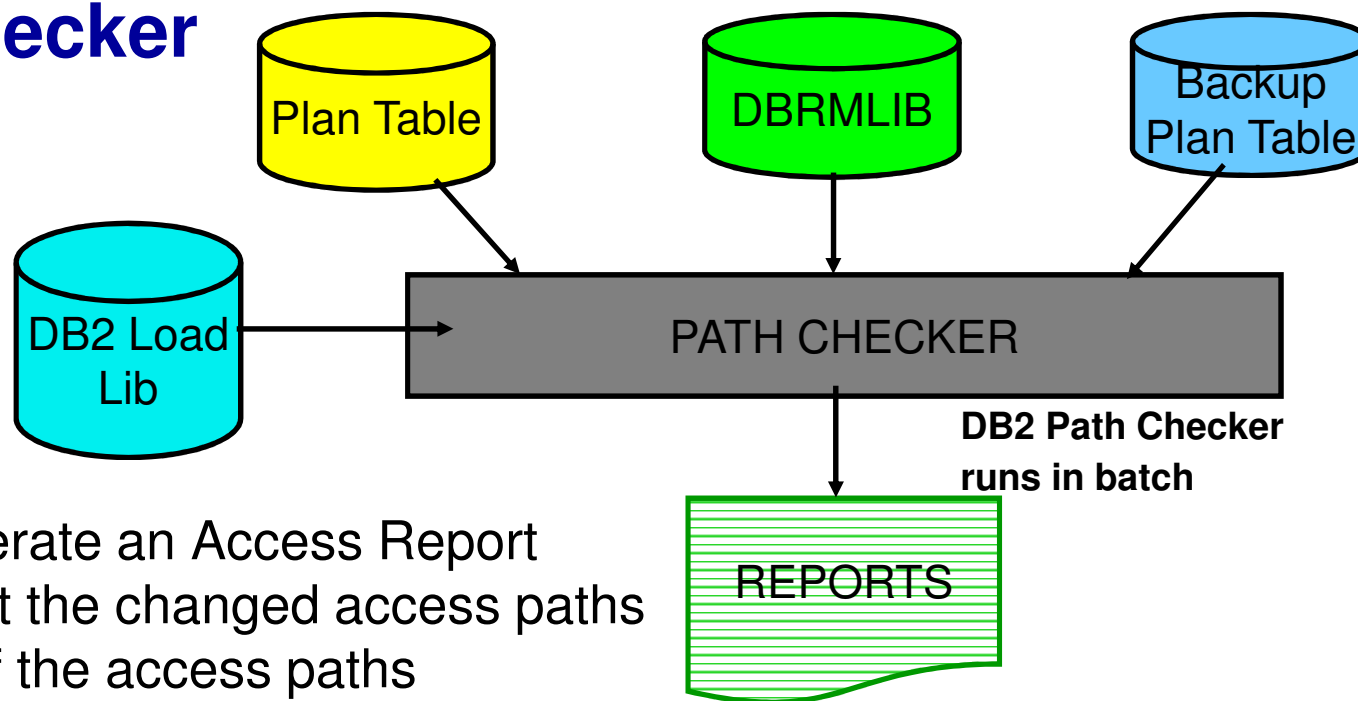
DB2 V2V Migration Best Practices

- 1. Establish a project team and project plan**
 - Review the Installation Guide checklists
 - **Identify DB2 Tools and ISV DB2 11 support requirements**
 - **Upgrade tools to required levels before starting DB2 11 migration**
- 2. Develop conversion and coexistence goals**
 - How did your V8 / V9 / V10 test plans work?
 - Reuse and improve upon your experiences
- 3. Establish performance baselines**
 - **Identify Key Performance Indicators (KPI) with OMEGAMON for DB2**
 - **Proactive tuning with Optim Query Workload Tuner (OQWT)**
- 4. Create Pre-production test environment**
 - **Use Optim Query Workload Replay to capture production workloads**
 - **Use DB2 Cloning Tool to iteratively refresh production catalog & data**
 - **Monitor KPIs to identify performance problems**
- 5. REBIND while in CMx**
 - **Path Checker identifies access path changes**
 - Use Plan Management features in DB2
 - **Invoke OQWT to resolve access path regressions**
- 6. Proactive tuning with OQWT**

DB2 Rebind Strategy

- **Rebind whilst in DB2 V11 CM**
- **Use Plan Management (Package/Bind stability)**
 - Consider FREEing original packages to establish a new DB2 V10 backup
- **Consider**
 - REBIND ... EXPLAIN (YES) APREUSE (WARN or ERROR)
 - REBIND ... EXPLAIN (YES) APCOMPARE (WARN or ERROR)
- **Rebind again in DB2 V11 NFM**
 - Again using package stability

DB2 Path Checker



- **REPORT** – generate an Access Report
 - Can view just the changed access paths
 - Or view all of the access paths
- **TEST** – validate access path **before** committing to a bind
- **COMPARE** – compare the access paths **after** an explain has been done (ea. in their own plan table)
- Unacceptable access paths can be avoided by favoring the old access path via optimizer hint (MAKE command) generated by DB2 Path Checker

**Create a clone of DB2 subsystem for
pre-migration testing ?**

Use DB2 Cloning Tool

Clone is great idea ! How to create one ?

Traditional methods

There are two traditional methods:

The first method is resource intensive

- Create a new DB2 subsystem (the clone)
- Create the objects from the source DB2 on the new DB2 via
DDL
- UNLOAD/LOAD

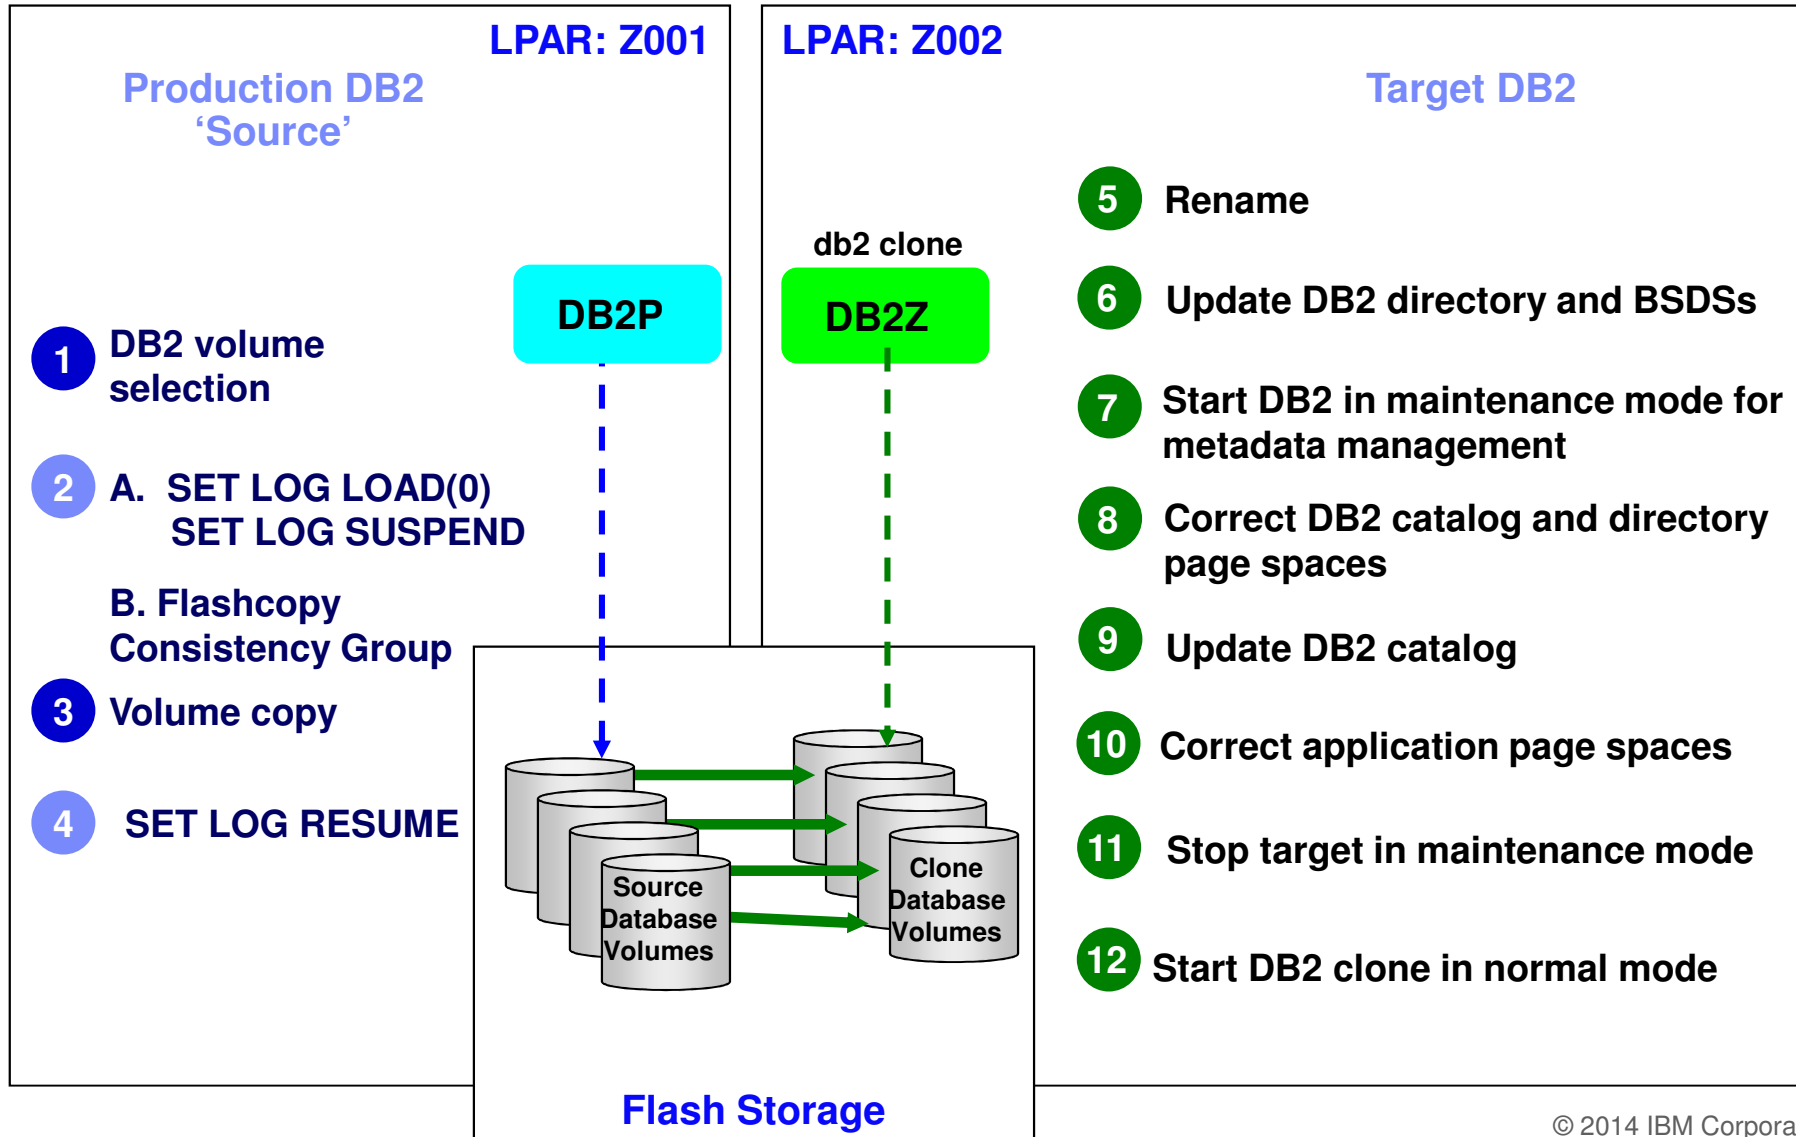
The second method requires an isolated LPAR and involves the following steps:

- Dump the DB2 data sets to removable media
- On an isolated LPAR, restore the data sets of the source DB2
- Start the clone DB2

Storage-aware cloning method using **CLONING Tool**

- Fast replication to replicate a DB2 subsystem in minutes
- To same LPAR or another LPAR
- CPU and I/O costs are born by the storage system

DB2 Cloning Tool : Dress Rehearsal



Pre-Migration Testing using the CLONE

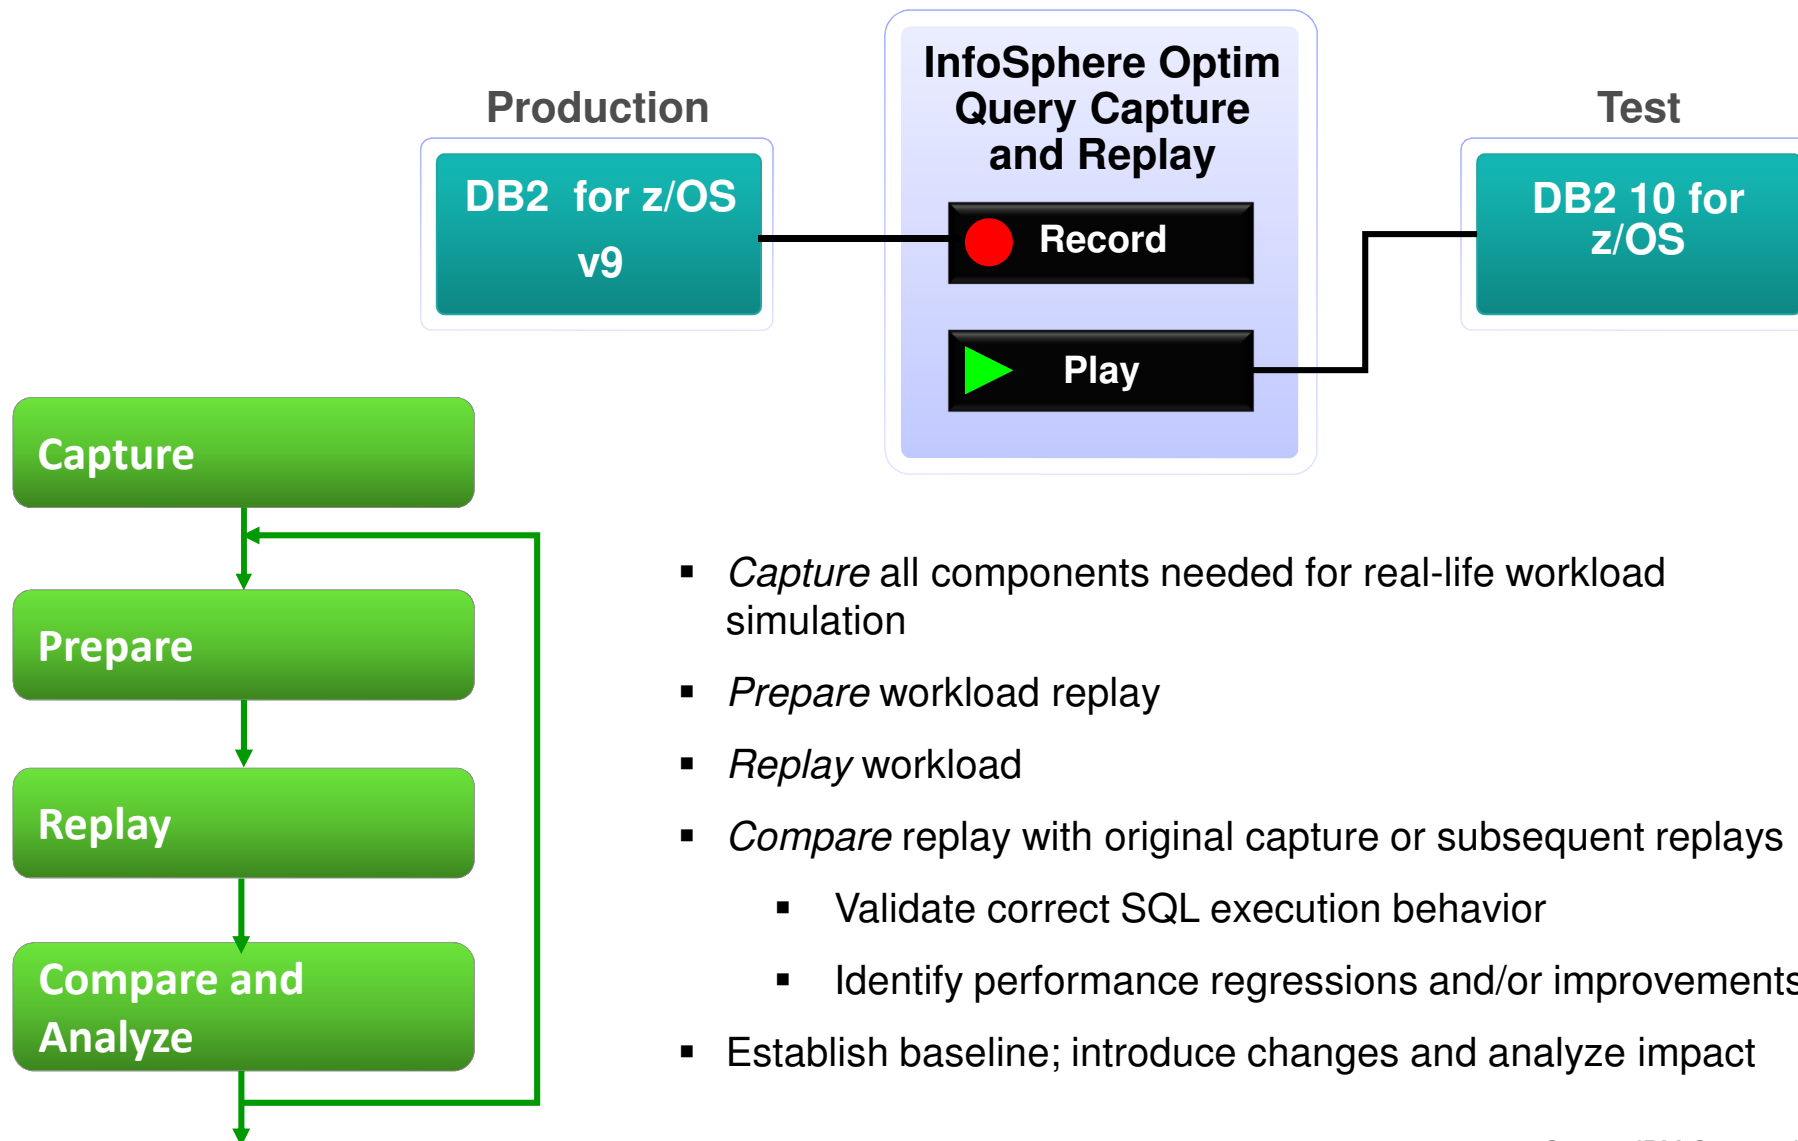
1. Create a DB2 subsystem clone using DB2 Cloning Tool.
2. Migrate the clone to the new version of DB2, using the documented IBM migration steps.
3. TEST your applications – Batch jobs, CICS applications, and Web transactions on the clone.
4. When issues are uncovered, correct them at your leisure during normal business hours, with no fear or anxiety of executing an emergency back-off procedure in the early morning hours.

White Paper – Smoothly migrating to a new version of DB2 using Cloning Technology

<http://public.dhe.ibm.com/common/ssi/ecm/en/imw14709usen/IMW14709USEN.PDF>

How about cloning SQL workload ?

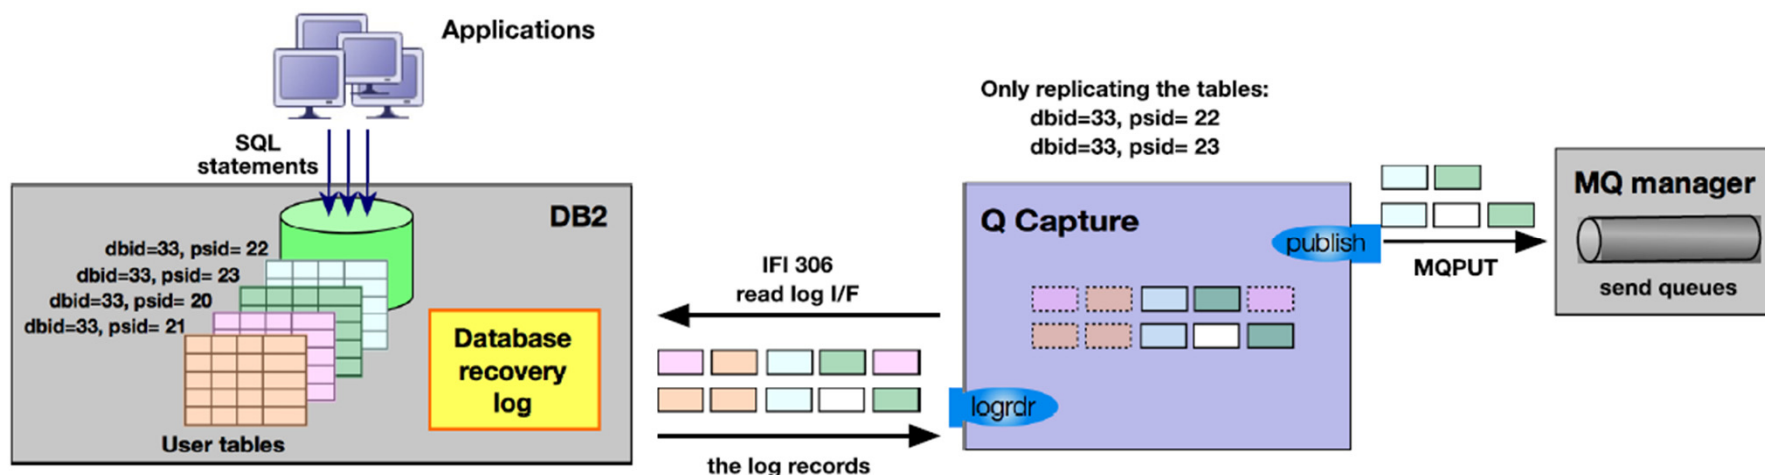
Solution overview



Infosphere Data replication (QRep) with DB2 11

- **Q Capture 10.2.1 is required for DB2 V11**
- Q Apply '1001' arch_level supports DB2 V11
- Q Capture and Q Apply 10.2.1 support DB2 V9, V10 as well

Infosphere Data Replication – IFI Log Filter



CPU Savings potential.

Previously: DB2 IFI returns all DML log records for all tables with DCC

**Now: Capture passes list of objects actually used in Pubs / Subs
IFI filters out the DML log records not used**

Available in DB2 V10 APAR PM90568 and DB2 V11

Infosphere Data Replication

Impact of DB2 11 changes on compression dictionary

▪ DB2 10 and prior

- Compression dictionary is stored in the Table space (or partition)
- If the table was REORG without KEEPDICTIONARY (which creates a new dictionary in the tablespace) and IFI cannot find the old compression dictionary to decode the log record. Error - [00C90064](#)
- Capture deactivates the subscription, forcing a potential full refresh.

▪ DB2 11

- Compression dictionary is still stored in the Table space (or partition)
- But when a new dictionary is built or if the table is altered to compress no, LOAD REPLACE or REORG will store the prior decompression dictionary on the log and write a SYSCOPY record identifying the position in the log.
- Capture will no longer see 00C90064 from IFI.

Agenda

- **IBM's DB2 Tools Portfolio Overview**
- **Using IBM DB2 Tools to Streamline V2V Migration**
- **DB2 V11 Utilities Enhancements**
- **Positioning of DB2 Archiving Options**
- **IDAA Operational Efficiencies**

DB2 Utilities Suite for z/OS v11.1

- **DB2 11 Utilities Suite provide support for ALL DB2 11 core function.**
- **In addition, new features in the DB2 11 Utilities:**
 - Reduce CPU usage, elapsed time and resource consumption
 - Maximize availability
 - Remove constraints and limitations
 - Simplify data management

In DB2 11, IBM provides an impressive list of improvements to its utilities. Whether they are overall improvements or new user choices, the **DB2 11 utilities** bring us **more availability, better performance, more control, more cost reduction** and better day to day usability. It's a win-win-win as far as I'm concerned.
Kurt Struyf, SuadaSoft



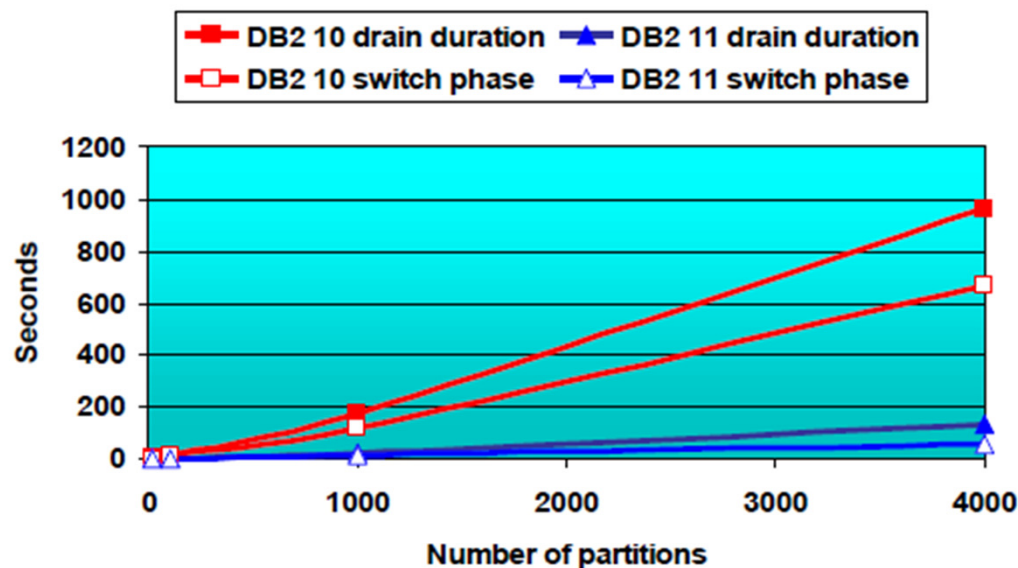
DB2 11 Utilities by the Numbers



- **More zIIP offload**
 - Up to 81% zIIP-eligible CPU with RUNSTATS COLGROUP
 - Up to 40% zIIP-eligible CPU in REORG & LOAD with inline distribution stats
- **REORG improvements**
 - REORG switch phase improvements can reduce the application outage window by up to 90%
 - Up to 71% elapsed time reduction for REORG of subset of Partitions
 - Increased parallelism elapsed time in REORG of 21%
- **CPU improvement for utilities with EXCLUDE NULL KEYS:**
 - LOAD 12%, REORG 25%, REORG INDEX 88%, REBUILD INDEX 72%, CHECK INDEX 79%, RUNSTATS 90%
- **RECOVER from part-level image copies reduced CPU by up to 50%, elapsed time by up to 40%**
- **LOAD from single input dataset elapsed time reduced by up to 70%**
- **Inline Stats vs Separate RUNSTATS**
 - 40% ET reduction for inline histogram stats
 - 28% ET & 19% CPU reduction with inline distribution stats

REORG improvements

- **Part-level REORG with NPSIs –**
 - Option to defer shadow index build until all keys passed through sort
 - Retrofit to DB2 9 & 10 in PM55051
 - REORG of 40% of partitions yields 55% ET reduction for 22% CPU increase
- **REORG – easier drain acquisition**
 - Prevents new claims on all target partitions whilst waiting for drains
 - Switch phase restructure : 91% ET reduction when REORGing 20 parts
 - SWITCHTIME NEWMAXRO parameter



REORG Improvements

- **Automated REORG mapping table management**
- **REORG without sorting data**
- **REORG with LOGRANGES NO and SHRLEVEL CHANGE**
 - Useful if problems with SYSLGRNX, and IBM Support requests REORG LOGRANGES NO

STATISTICS Enhancements

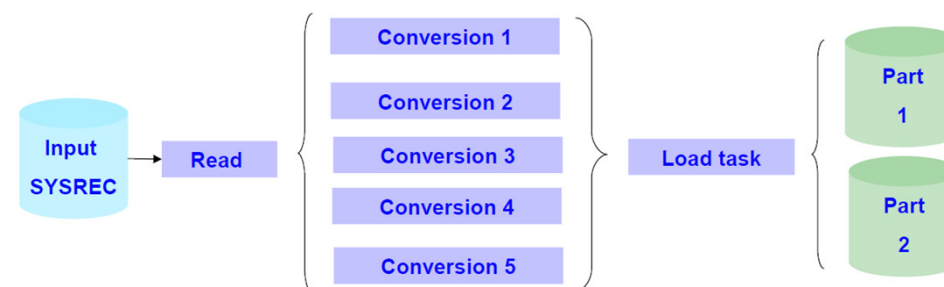
- **More zIIP offload for RUNSTATS distribution statistics**
 - Up to 80% zIIP-eligible
- **zIIP offload for inline statistics**
 - Additional 30% offload to zIIP
- **Enhance inline statistics for RUNSTATS avoidance**
 - Inline statistics collection on NPSIs during REORG with SORTNPSI
 - Inline histogram statistics

LOAD improvements

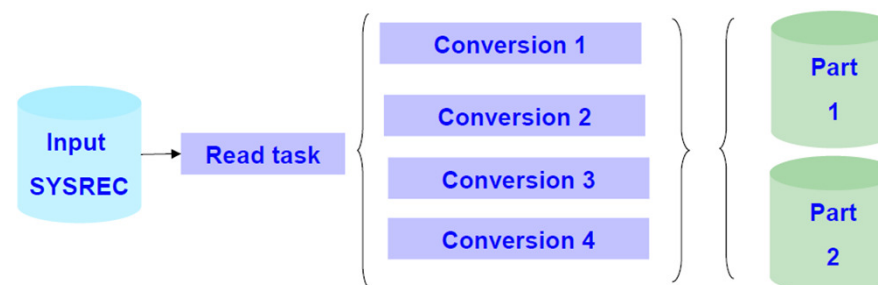
- **Crossloader support for XML data**
- **Exploit FETCH CONTINUE for processing large LOBs & XML data in Crossloader**
 - Reduce vstor requirement
 - Avoid DSNU1178i errors
 - 28% CPU reduction
 - Load of 1Mb LOBs
- **ZIIP offload for LOAD REPLACE PART clearing of NPSIs**
 - 100% offload to zIIP for LOAD REPLACE with dummy input

LOAD improvements

- LOAD SHRLEVEL NONE PARALLEL with single input dataset
 - Parallel data conversion
 - Not supported for PBGs
 - 50% ET reduction possible on single SYSREC load



- LOAD SHRLEVEL CHANGE PARALLEL
 - Supports non-partitioned as well as partitioned
 - Single input dataset
 - Not supported for PBGs
 - >80% ET reduction




Agenda

- **IBM's DB2 Tools Portfolio Overview**
- **Using IBM DB2 Tools to Streamline V2V Migration**
- **DB2 V11 Utilities Enhancements**
- **Positioning of DB2 Archiving Options**
- **IDAA Operational Efficiencies**


Organizations are Challenged with Data Growth

Increasing Costs



Buying more storage is not a “cheap” fix when you add the operational burden

Poor Application Performance



Business users & customers wait for application response; DBA's spend majority of time fixing performance issues

Manage Risk & Compliance



The “keep everything” strategy can impact disaster recovery and data retention & disposal compliance

DB2 z/OS V11 Transparent Archiving

```
SET SYSIBMADM.GET_ARCHIVE = 'N';
```

```
SELECT DISTINCT ACCOUNT  
FROM TRANSACTION_AET ;
```

ACCOUNT
A
D

```
SET SYSIBMADM.GET_ARCHIVE = 'Y';
```

```
SELECT DISTINCT ACCOUNT  
FROM TRANSACTION_AET ;
```

ACCOUNT
A
B
C
D

TABLE: TRANSACTION_AET

ACCOUNT	TXN_ID	TXN_AMT	TXN_DATE
A	1234	55	1/3/2014
D	1235	123	1/3/2014
A	1236	-66	2/3/2014

```
SET SYSIBMADM.MOVE_TO_ARCHIVE = 'Y' or 'E'
```

Archive Data for Accounts B & C

- by SQL Delete
- By Utilities (REORG, LOAD etc...)





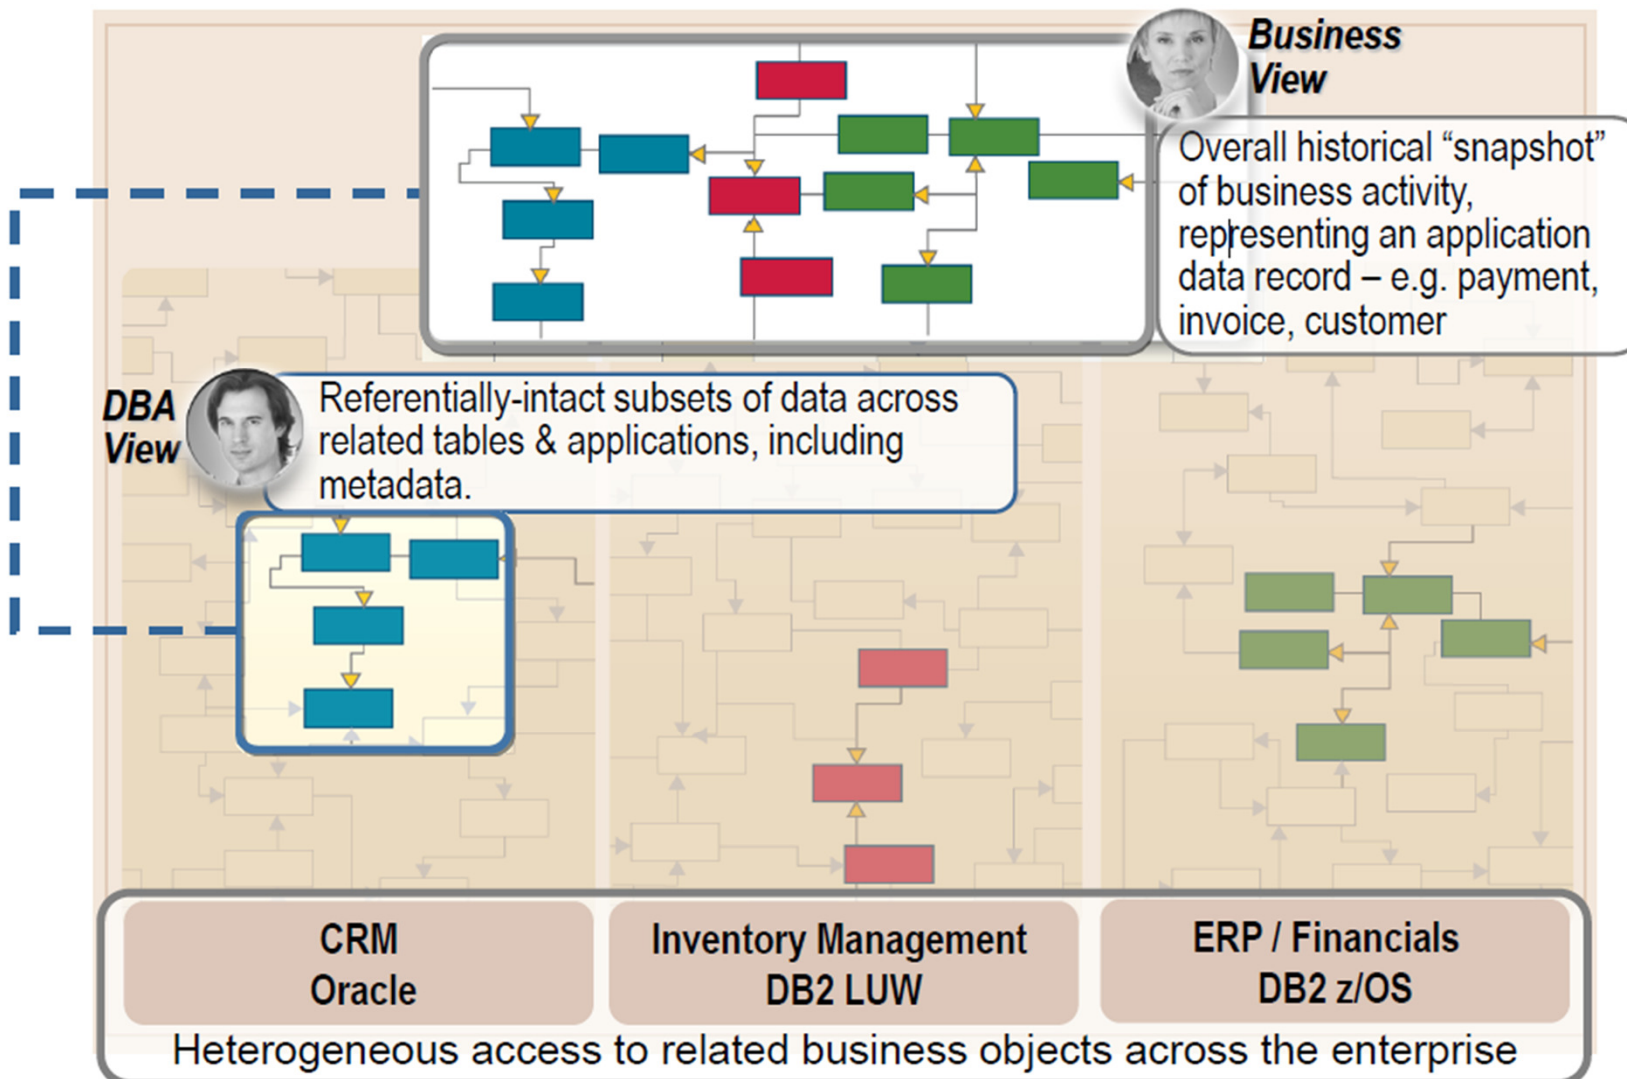
TABLE: TRANSACTION_ARCHIVE

ACCOUNT	TXN_ID	TXN_AMT	TXN_DATE
B	991	28	1/3/2014
B	992	-33	1/3/2014
C	993	182	2/3/2014

IDAA High Performance Space Saver



Optim Data Growth Solutions



Comparisons

	DB2 / HPSS	Optim
Comply with Regulations		
Policy-Based Archival & Disposal	No	Yes
Business-Object Based Archiving	No	Yes
Immutable Archived Data with Full Audit	Manual	Yes
Full Restore of Archived Data	Manual	Yes
Lower Cost of Dormant Data		
Archived Data can reside outside DB2	No / Yes	Yes
Archived Data Compressed	Yes	Yes
Access Dormant Data	Yes	Yes
Tooling for High Productivity	Yes	Yes
Maintain High Performance Data Access		
Archive Data can reside in DB2	Yes	No
Transparent Access to Active and Archive Data	Yes	No
SQL Access to Archived Data	Yes	Manual

New Australian Law Privacy Principles

- The Privacy Amendment (Enhancing Privacy Protection) Act 2012 (Privacy Amendment Act) was introduced to Parliament on 23 May 2012 and was passed with amendments on 29 November 2012. The Privacy Amendment Act introduces many significant changes to the Privacy Act and becomes effective *12 March 2014*
- Includes a set of new, harmonised, privacy principles that will regulate the handling of personal information by both Australian government agencies and businesses.
 - Known as the Australian Privacy Principles (APPs).
 - Replacement for existing Information Privacy Principles (IPPs) that currently apply to Australian government agencies and the National Privacy Principles (NPPs) that currently apply to businesses.
- And this time, the penalties can hurt!
 - The Office of the Australian Information Commissioner can
 - Accept enforceable undertakings
 - Seek civil penalties
 - Conduct assessments of privacy performance
 - Impose fines
 - ❑ Up to \$340K for individuals, \$1.7M for entities

Optim Solution Portfolio

included in all Optim Enterprise Editions



Identify Relevant & Sensitive Data

Find what data must be retained, protected or removed

IBM InfoSphere Optim Archive

Dispose of Unnecessary Data

Remove unnecessary data from critical transactional or analytics applications



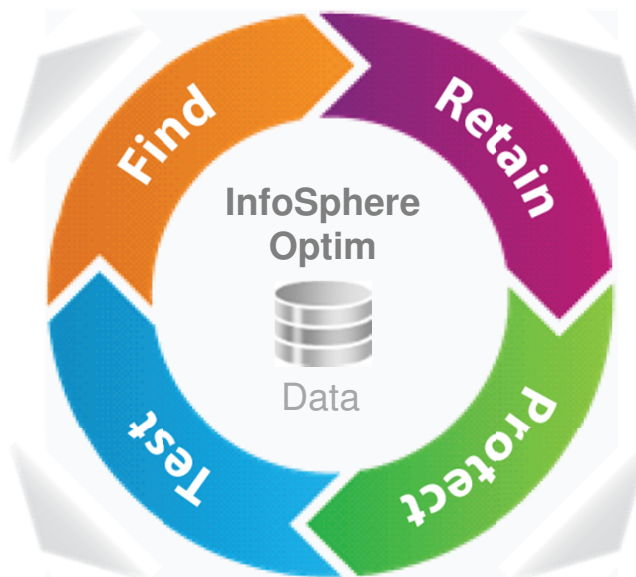
Retain Essential Data

Historical inactive data is safely retained while easily accessible for reports and compliance



Protect Sensitive Data

Privatize Test Data: Customer IDs, credit cards and financial data are masked or redacted



Optimize Test Data

Automate and optimize the application test processes that rely on data to enable continuous testing & DevOps

IBM InfoSphere Optim Test Data Management

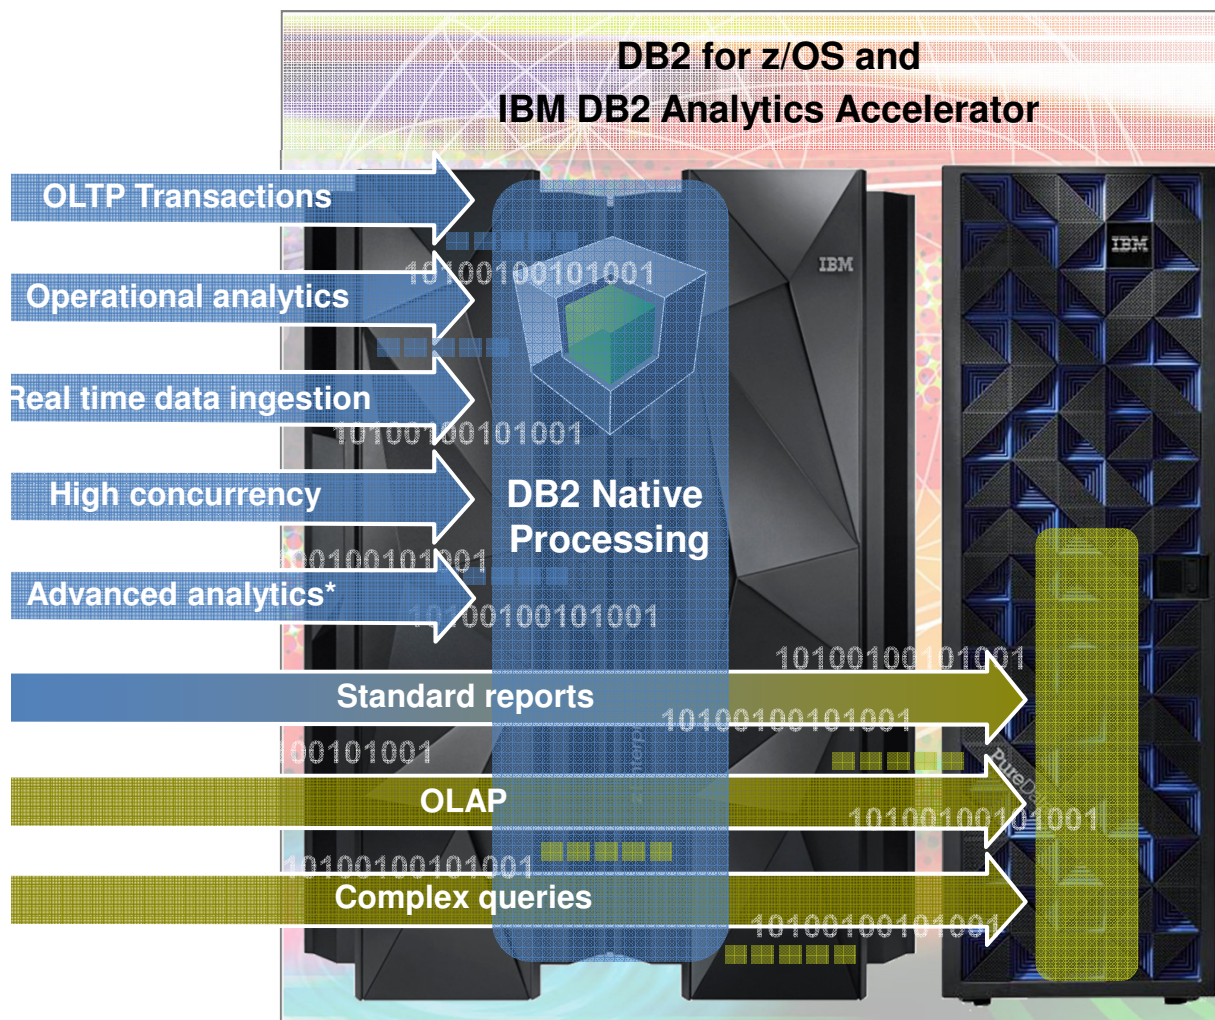
IBM InfoSphere Optim Workload Replay

IBM InfoSphere Optim Data Privacy

Agenda

- **IBM's DB2 Tools Portfolio Overview**
- **Using IBM DB2 Tools to Streamline V2V Migration**
- **DB2 V11 Utilities Enhancements**
- **Positioning of DB2 Archiving Options**
- **IDAA Operational Efficiencies**

IDAA Recap



- Hybrid DB Server
- Best of Both Worlds
 - OLTP & Batch
 - Analytics
- Dynamic Routing of SQL
- Speed (up to 2000 times)
- No tuning “load ‘n go” Appliance
- Cost Effective (an extreme offload engine)
- Merge Operational & Analytics Processing
- Analytics without waiting for populating the Data Warehouse
- Real Time Scoring on Real Time data

Performance & Savings

Query	Total Rows Reviewed	Total Rows Returned	DB2 Only		DB2 with IDAA		Times Faster
			Hours	Sec(s)	Hours	Sec(s)	
Query 1	2,813,571	853,320	2:39	9,540	0.0	5	1,908
Query 2	2,813,571	585,780	2:16	8,220	0.0	5	1,644
Query 3	8,260,214	274	1:16	4,560	0.0	6	760
Query 4	2,813,571	601,197	1:08	4,080	0.0	5	816
Query 5	3,422,765	508	0:57	4,080	0.0	70	58
Query 6	4,290,648	165	0:53	3,180	0.0	6	530
Query 7	361,521	58,236	0:51	3,120	0.0	4	780
Query 8	3,425,29	724	0:44	2,640	0.0	2	1,320
Query 9	4,130,107	137	0:42	2,520	0.1	193	13

Queries run faster

- Save CPU resources
- People time
- Business opportunities

Actual customer results, October 2011


DB2 Analytics Accelerator: “we had this up and running in days with queries that ran over 1000 times faster”

DB2 Analytics Accelerator: “we expect ROI in less than 4 months”


Accelerating decisions to the speed of business

CPU off-load example (BIDAY – IBM internal benchmark)

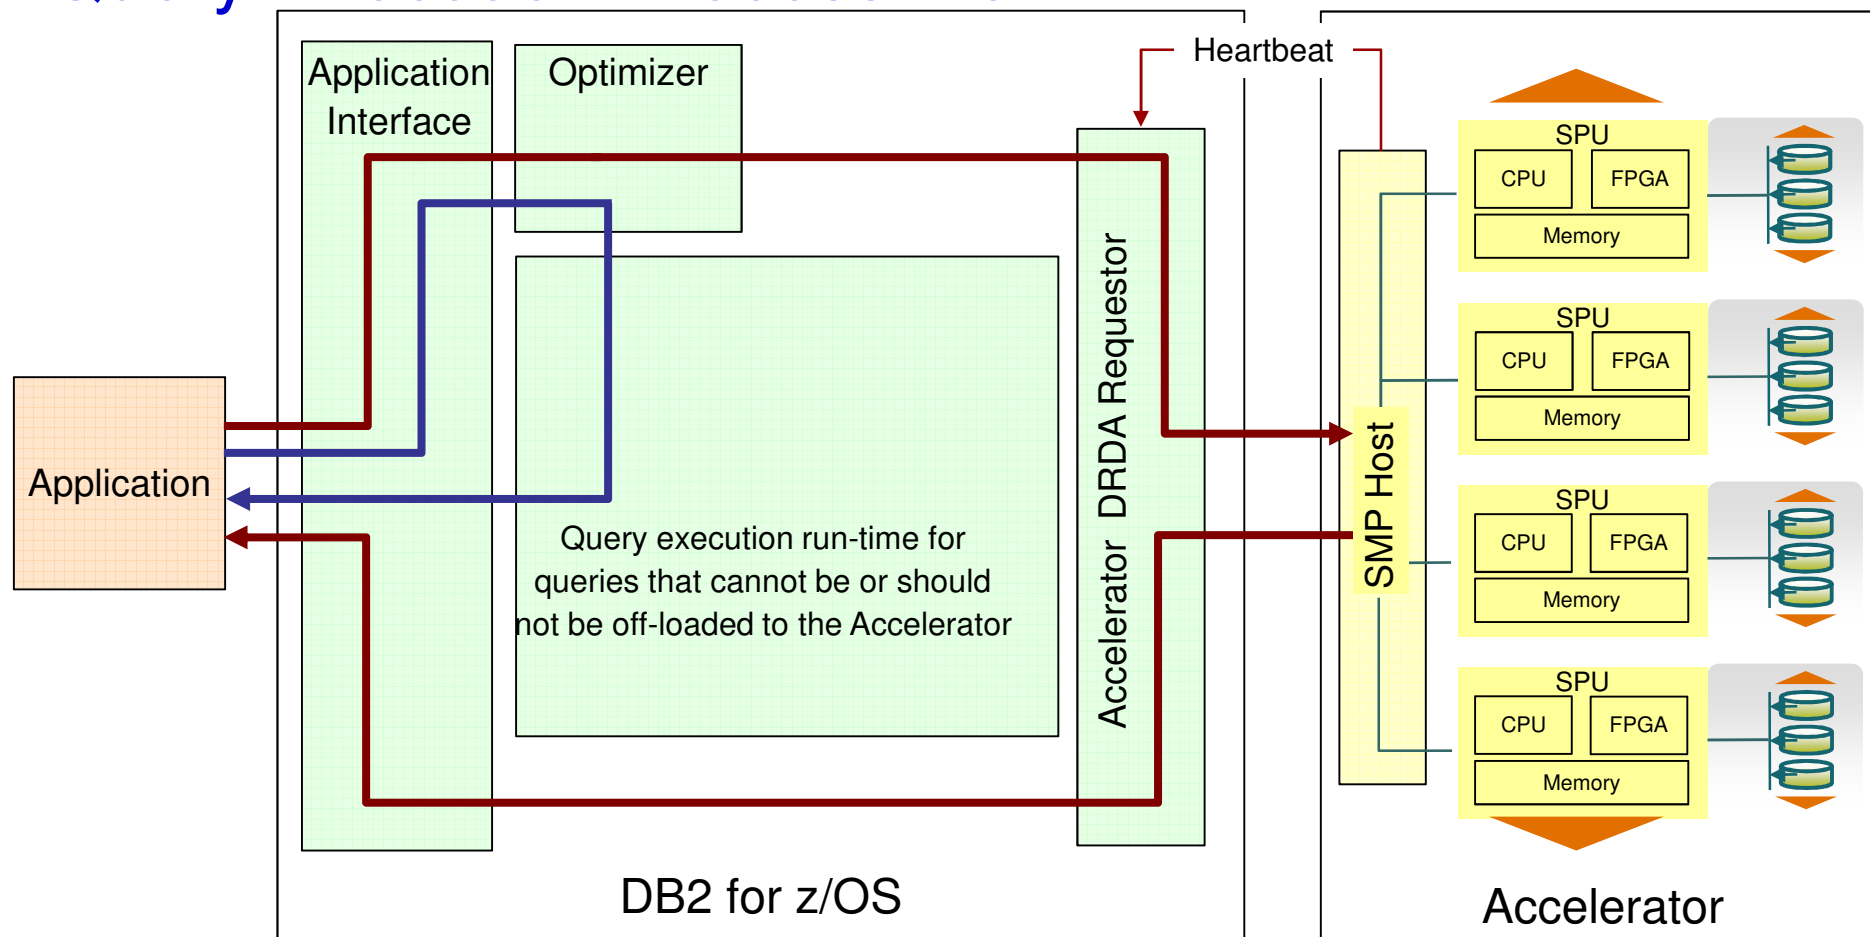
Workload Before IDAA - DB2 Address Spaces CPU activity



Workload After IDAA - DB2 Address Spaces CPU activity



Query Execution Process Flow



- Queries executed without DB2 Analytics Accelerator
- Queries executed with DB2 Analytics Accelerator
- Heartbeat (Accelerator availability and performance indicators)

Static SQL Support

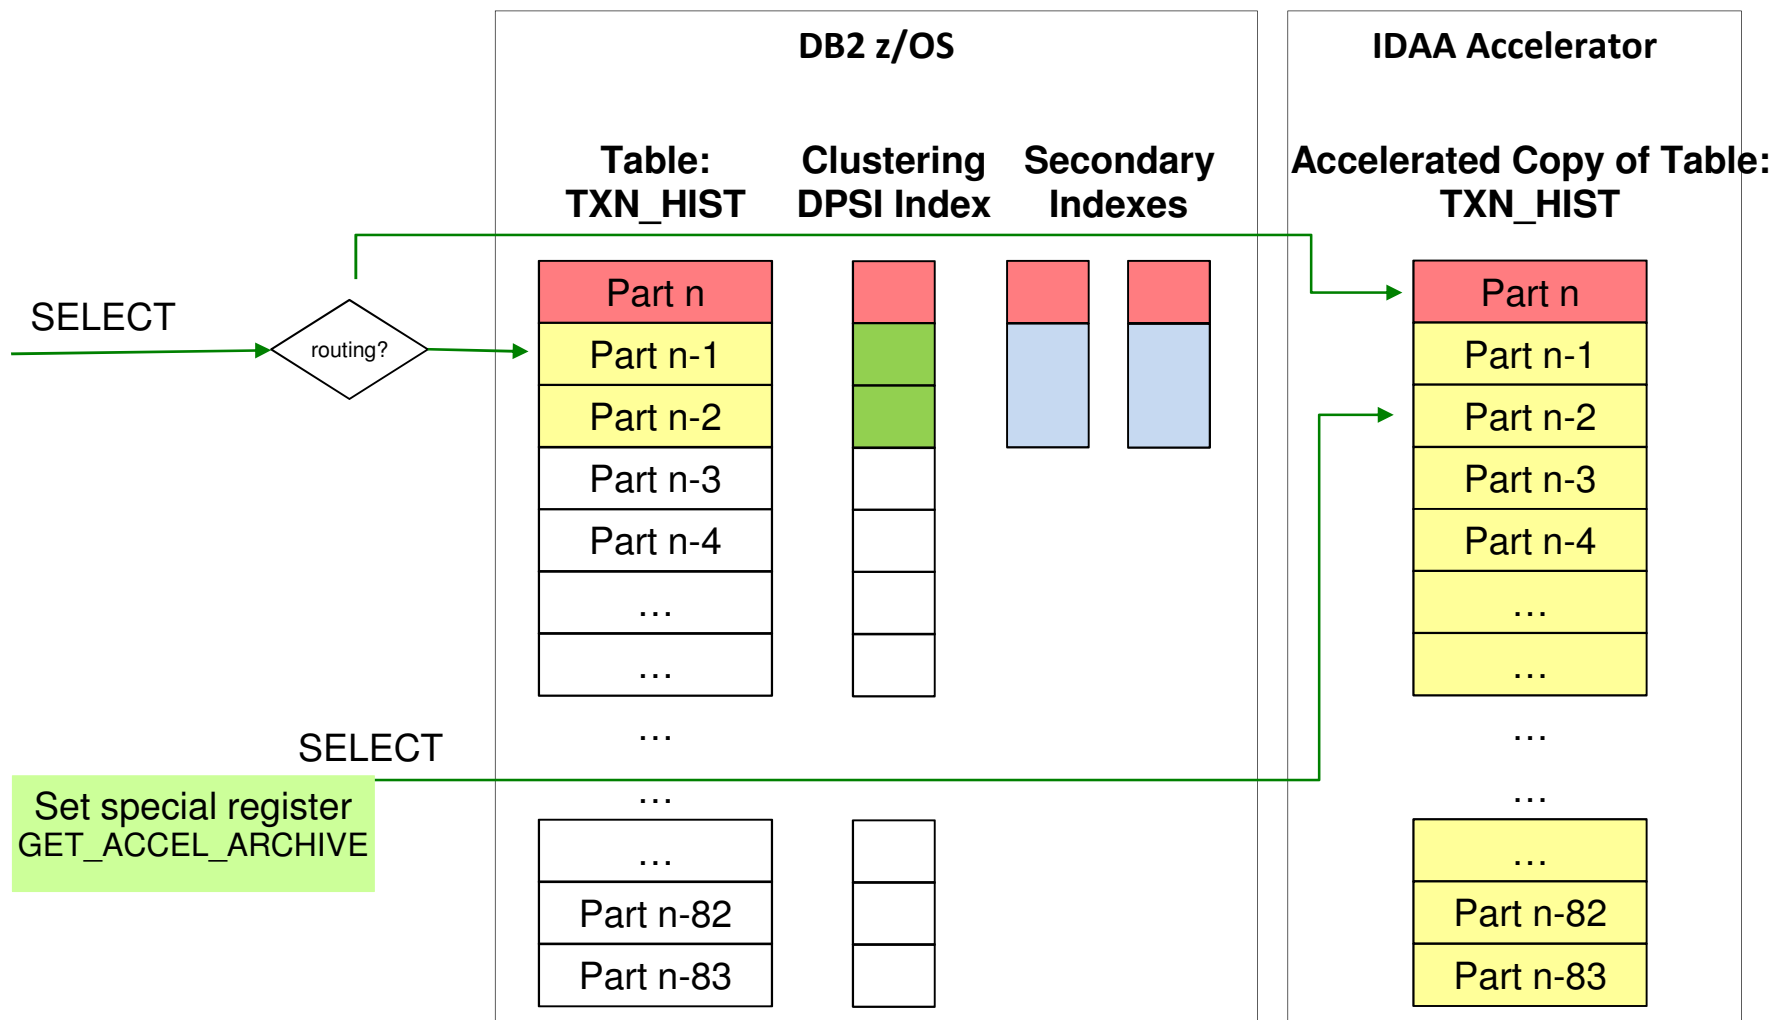
New bind options

**QUERYACCELERATION (NONE | ENABLE |
ENABLEWITHFAILBACK | ELIGIBLE | ALL)**

GETACCELARCHIVE (NO | YES)

- The new bind options are supported for
 - **BIND/BIND COPY/REBIND PACKAGE (local and remote binds)**
 - **BIND DEPLOY (for SQLPL procedure), REBIND TRIGGER PACKAGE**
 - **ALTER/CREATE PROCEDURE/FUNCTION for native SQLPL procedure and SQLPL scalar function**
- Default value for both is “option not specified”
- If the bind options are specified, they will set the initial values of the associated special registers.
 - **Note, that explicit SET statement and ZPARM setting not detected at BIND time**

High Performance Space Saver



Business innovation with zEnterprise Analytics



***Access to
operational data to
improve customer
satisfaction***

Swiss Mobiliar has achieved its objective of running its growing **transaction processing and analytics workloads side by side without increasing compute requirements.**

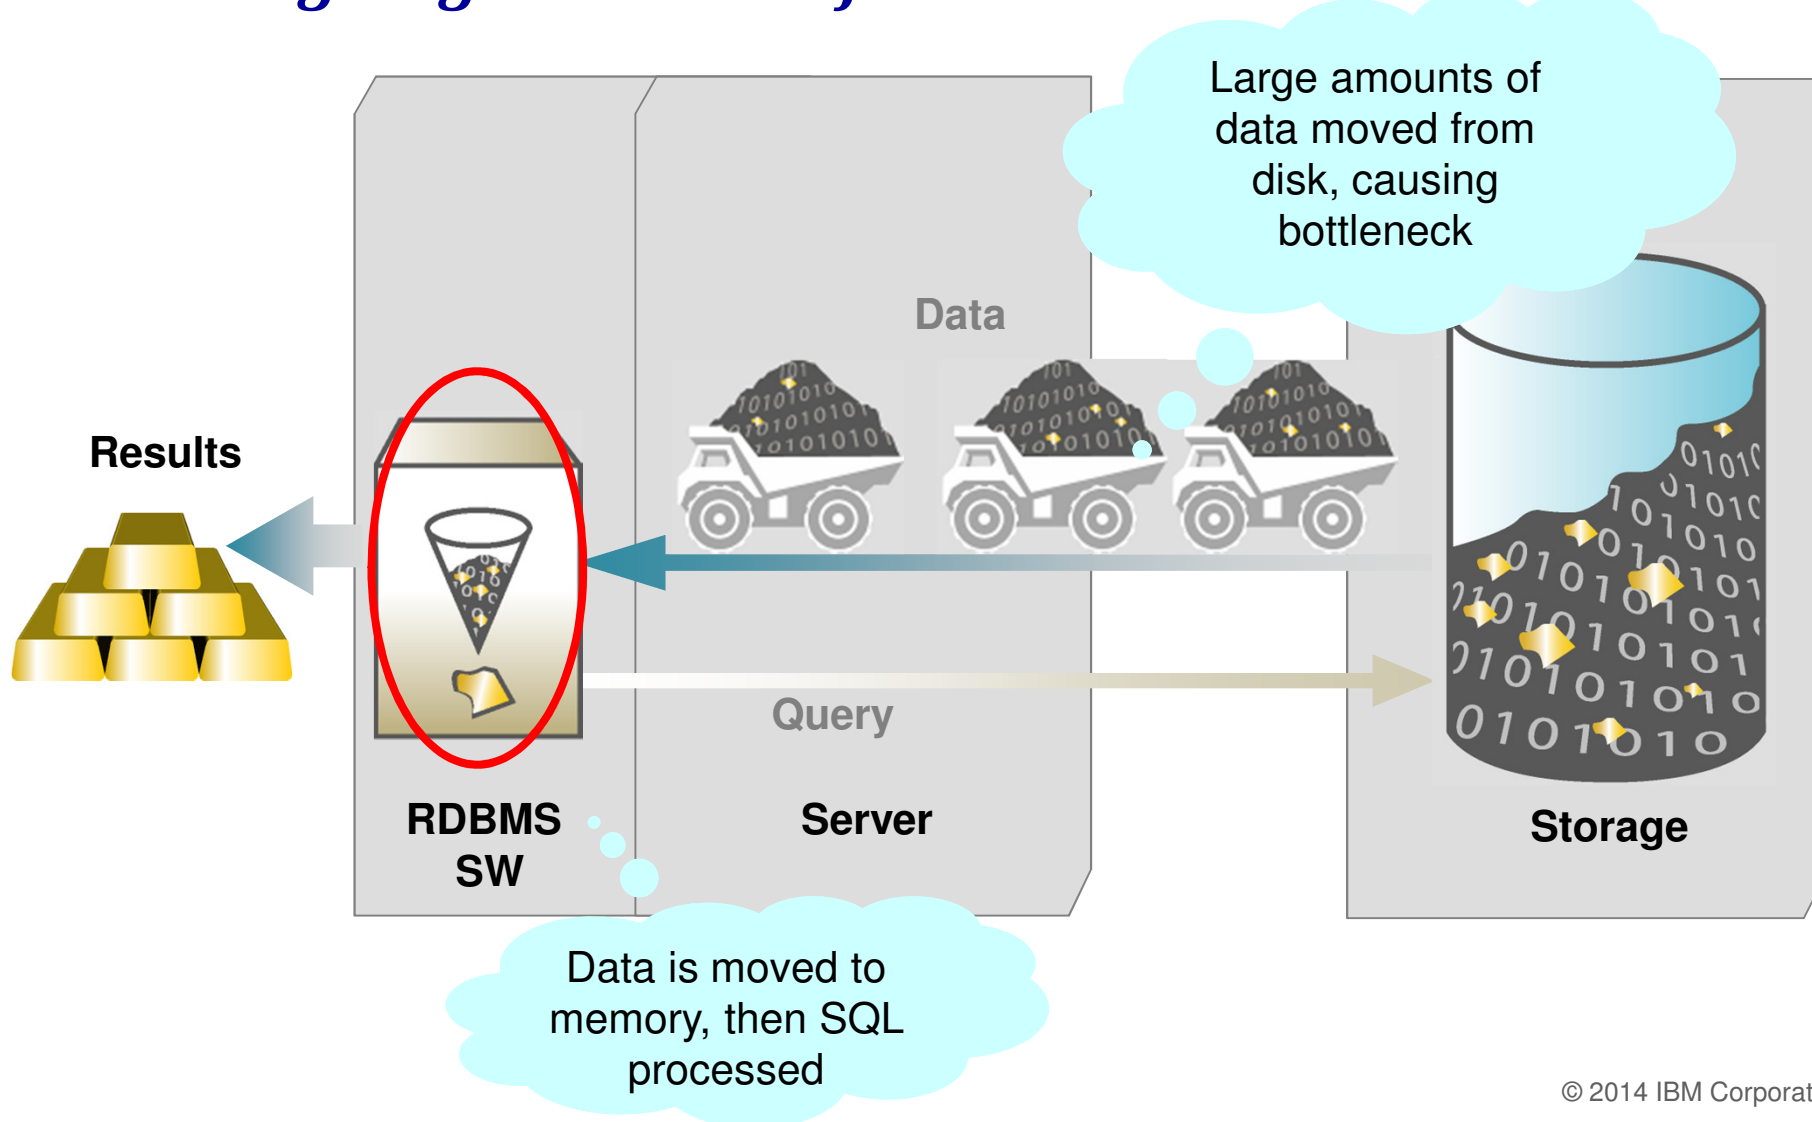
- Rapidly processes reporting queries with no increase in active z server cores
- Accelerates 50% of queries by a factor of 100
- Reduces transaction response times by 20%

Swiss Mobiliar
Insurance & Pensions

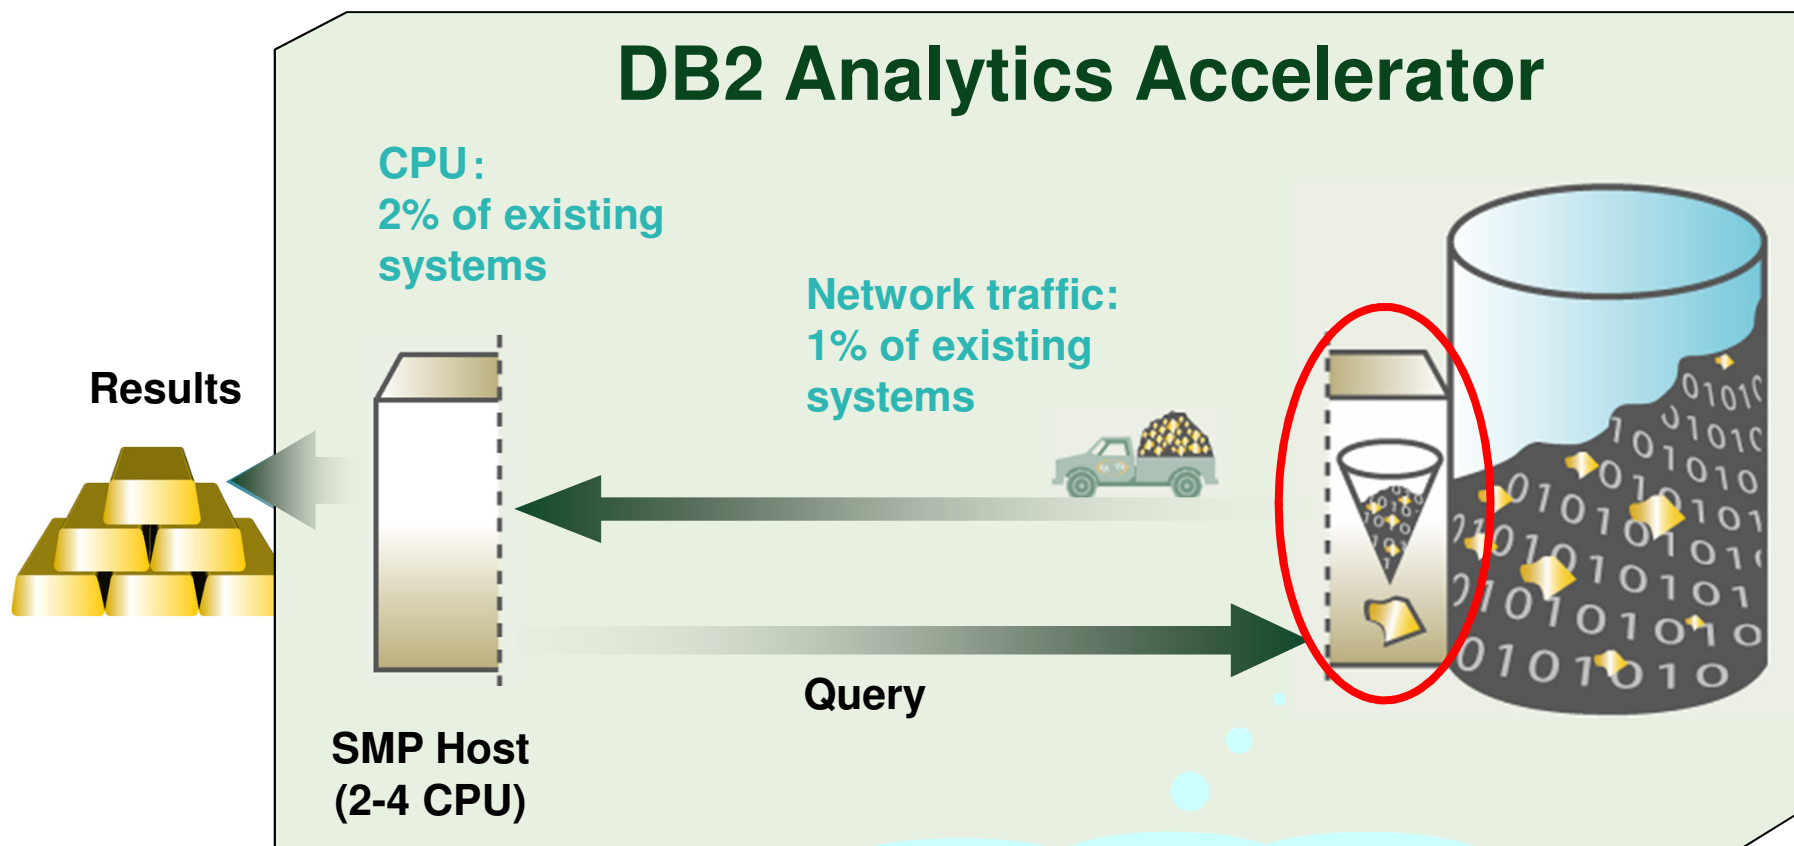
“Queries that used to take five hours to complete are now processed in just 20 seconds in the optimized mainframe environment—and we can run them any time, day or night, with no interruption to our production systems on the mainframe.”

Thomas Baumann, IT Performance Architect at Swiss Mobiliar

Legacy DWH Architectures: *Moving large amounts of data becomes Bottleneck!!*

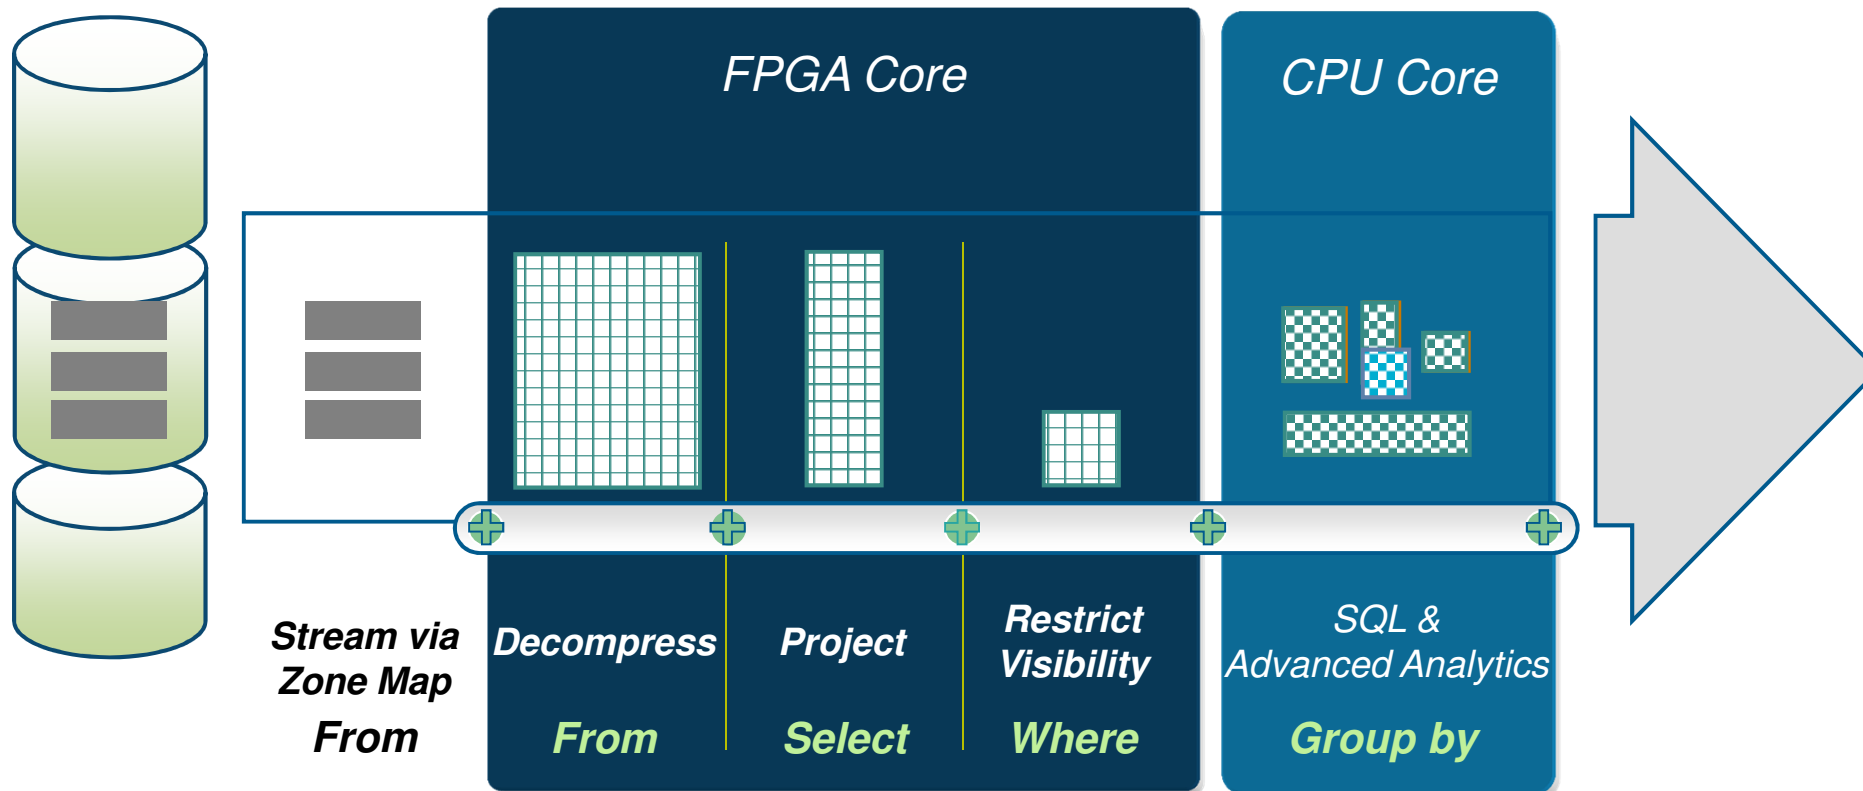


“Streaming Architecture”



Data processed as streams
from disk, before moved to
memory

S-Blade Data Stream Processing



```

Select State, Age, Gender, count(*) From MultiBillionRowBaseTable Where BirthDate <= '1960-01-01'
And State = 'CA' (And State in ('CA', 'NY')) Group By State, Age, Gender Order by State, Age, Gender
    
```


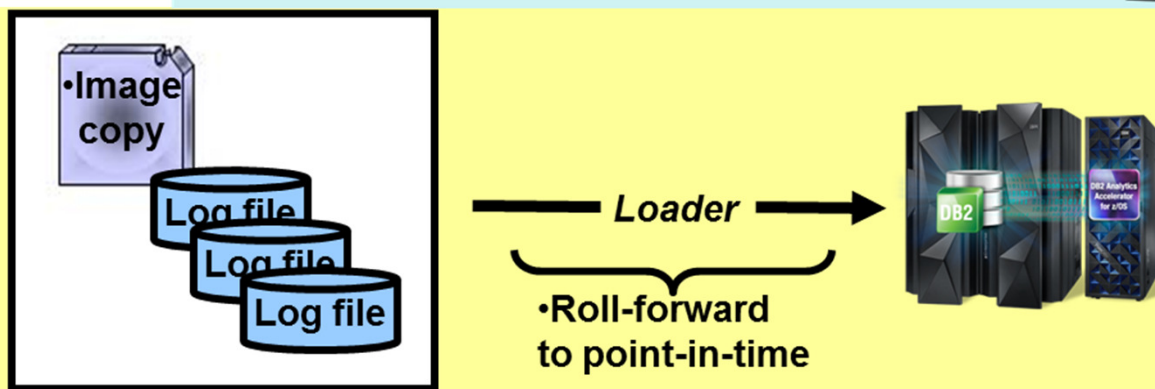
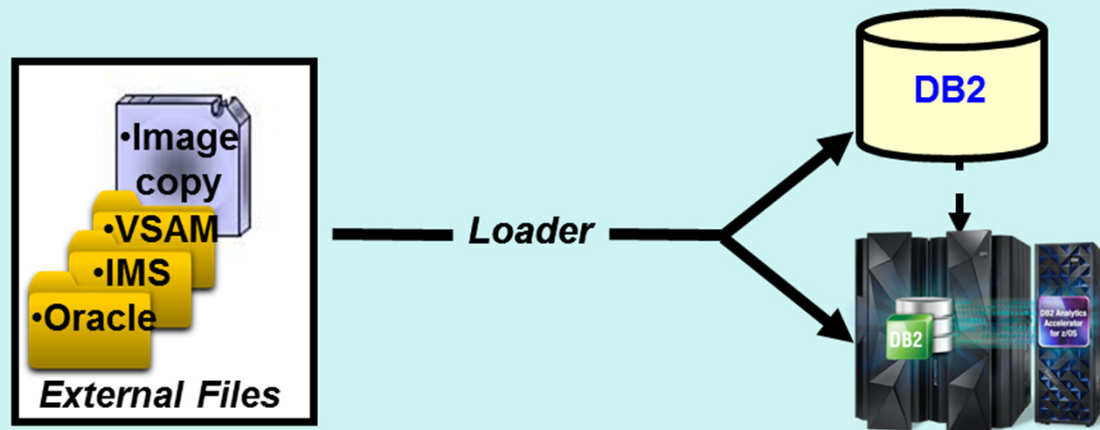
Synchronization Options with DB2 Analytics Accelerator

Synchronization options	Use cases, characteristics and requirements
<p>Full table refresh</p> <p>The entire content of a database table is refreshed for accelerator processing</p> <p>Supports Change Detection</p>	<ul style="list-style-type: none"> ▪ Existing ETL process replaces entire table ▪ Multiple sources or complex transformations ▪ Smaller, un-partitioned tables ▪ Reporting based on consistent snapshot
<p>Table partition refresh</p> <p>For a partitioned database table, selected partitions can be refreshed for accelerator processing</p> <p>Supports Change Detection</p>	<ul style="list-style-type: none"> ▪ 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
<p>Incremental Update</p> <p>Log-based capturing of changes and propagation to DB2 Analytics Accelerator with low latency (typically few minutes)</p>	<ul style="list-style-type: none"> ▪ Scattered updates after “bulk” load ▪ Reporting on continuously updated data (e.g., an ODS), considering most recent changes ▪ More efficient for smaller updates than full table refresh

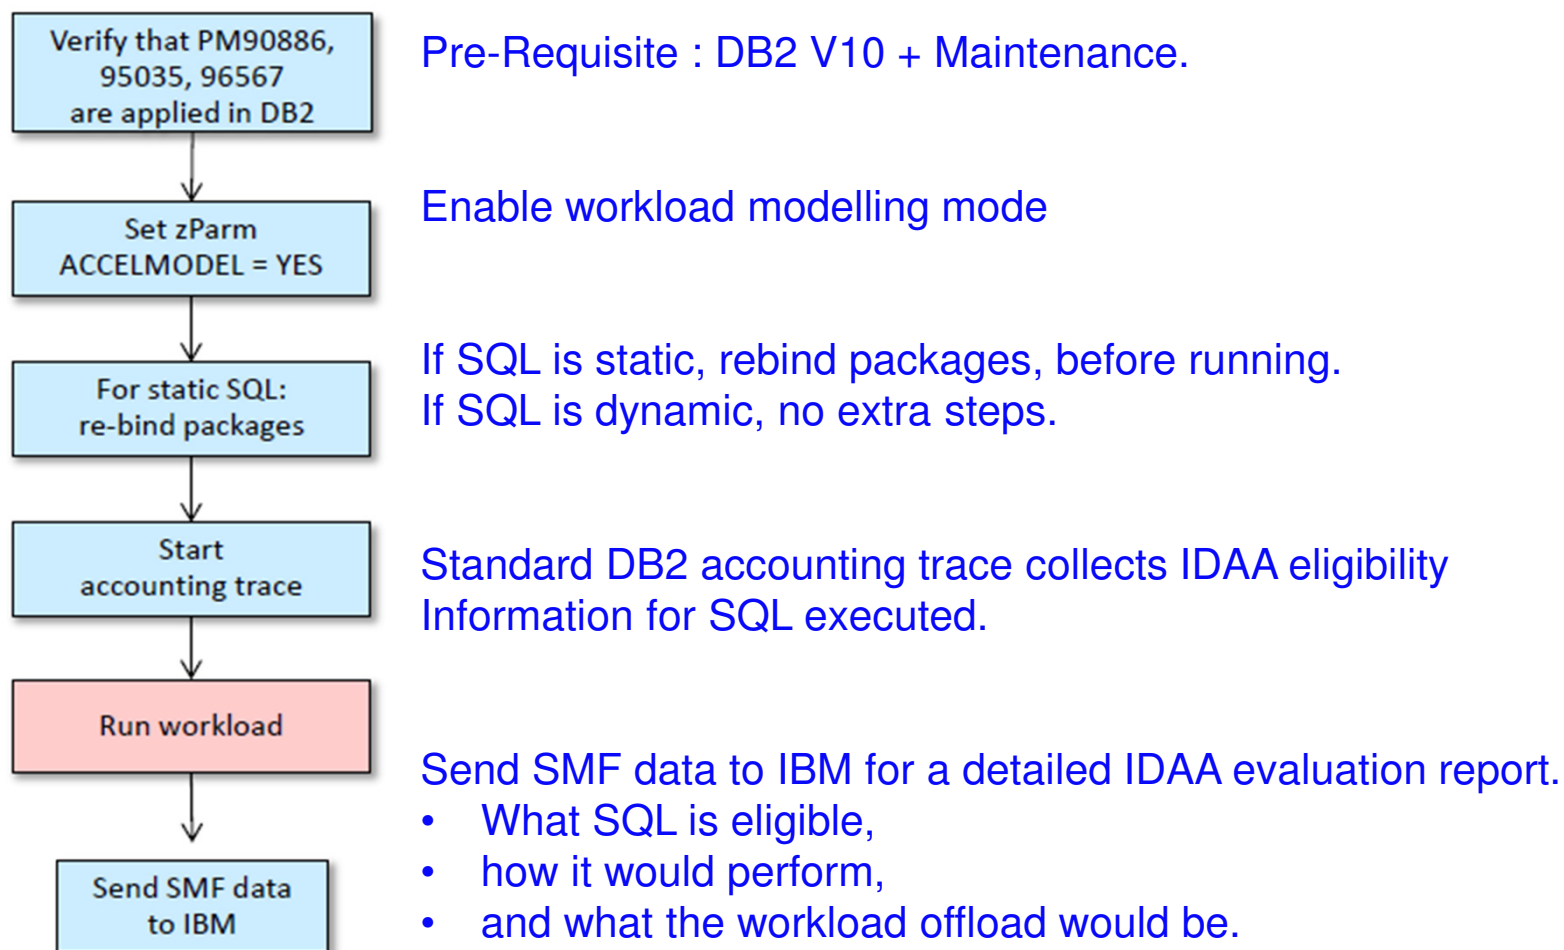
Change detection: DB2 automatically determines if table / partition was changed otherwise skips the table partition in the load request

- uses DB2 real time statistics to determine data changes

IDAA Loader



Simpler, Faster Workload Evaluation Method



OQWT : New Analytics Advisor

- Identify candidate queries and tables to be routed to the accelerator
- Identify candidate tables to be routed to the accelerator
- Implement advisor-based tuning recommendations for mixed workloads of accelerated and un-accelerated queries
- Diagram accelerated queries in Access Plan Graphs
- Shared Eclipse-based workspace and user interface with IBM DB2 Analytics Accelerator Studio
- Integrates with Query Monitor and OMPE for capturing query workloads for complete analysis
- Enable “what if” analysis

Benefits

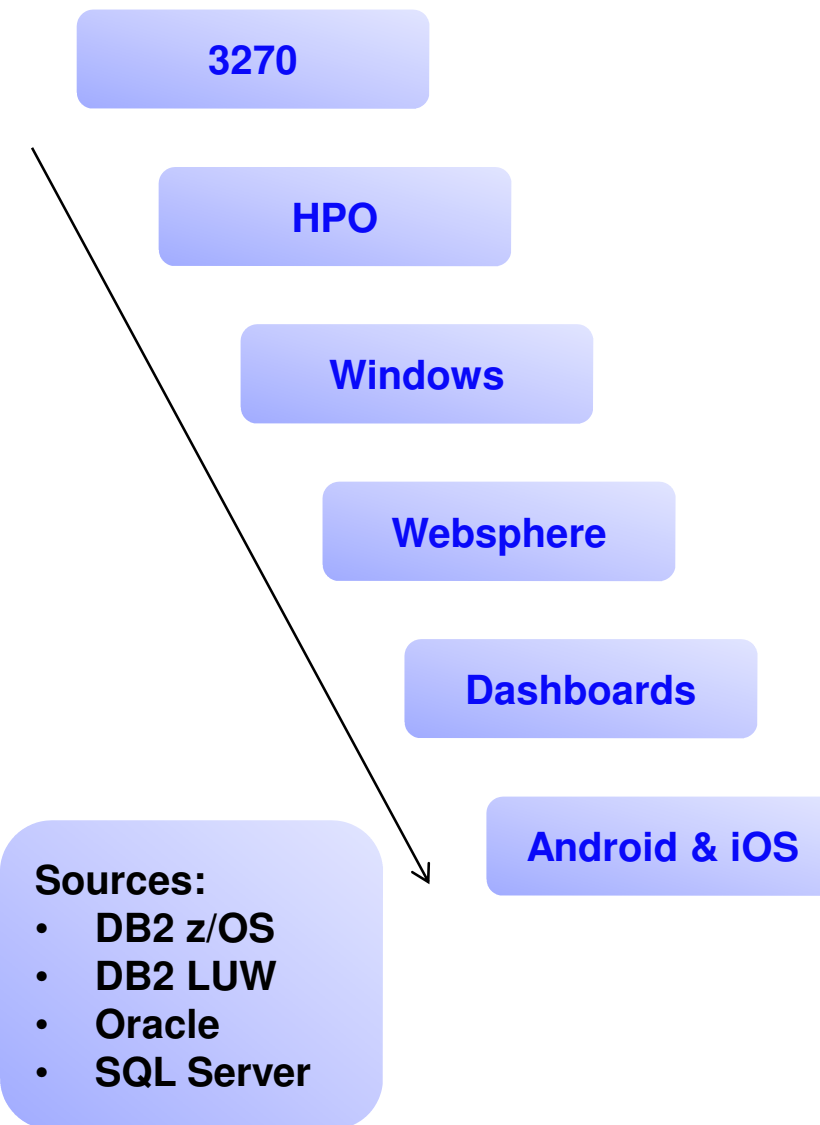
- Shorten the process of selecting tables to be accelerated
- Visualize access paths of accelerated queries
- Increase productivity by working with accelerated queries through a unified interface
- **Increase overall system capacity**

<http://youtu.be/pQYMRHJW7NU>

QMF V11

World-class set of features and functions for IBM DB2 queries, reporting, dashboarding and analytics with IBM DB2 Query Management Facility 11 for z/OS

- Expanded analytics for QMF TSO users with sophisticated new charting and statistical analysis capabilities
- Access to QMF dashboards and reports now on tablet devices
- Text analytics capabilities to process and extract key information



Agenda

- **IBM's DB2 Tools Portfolio Overview**
- **Using IBM DB2 Tools to Streamline V2V Migration**
- **DB2 V11 Utilities Enhancements**
- **Positioning of DB2 Archiving Options**
- **IDAA Operational Efficiencies**

DB2 z/OS V11 Migration Planning Workshop

Available since DB2 V8
 1 day workshop
 Identifies migration planning issues
 No charge

New Functions
 Version Packaging
 Pre-requisites
 Removed & Deprecated Functions
 New & Changed Parameters
 Migration Steps
 Catalog Evolutions
 Incompatibilities
 Application Compatibility
 RBS / LRSN Expansion

