




# *The DB2 Analytics Accelerator*

## *Version 3.1*

### *Overview*



Saso Prek, IBM SWG

© 2012 IBM Corporation

Information Management

## DB2 Analytics Accelerator

*Further extending the features*



*Blending System z and Netezza technologies to deliver unparalleled, mixed workload performance for complex analytic business needs.*

### *More insight from your data*

- Unprecedented response times for “right-time” analysis
- Complex queries in seconds rather than hours
- Transparent to the application
- Inherits all System z DB2 attributes
- No need to create or maintain indices
- Eliminate query tuning
- Fast deployment and time-to-value

*Introducing*

# ***DB2 Analytics Accelerator V3***

*Reducing the Cost of High Speed Analytics*

## **Improve Productivity**

- Eliminate query tuning
- Eliminate table indexing
- Minimize storage admin


## **Lower Host Costs**

- Reduce storage costs
- Offload query processing
- Defer system upgrades

## **Consolidate**

- Reduced complexity
- Reduced software costs
- Reduced hardware costs

# Customer Example:



**270 of the Mixed Workload Queries**



Executes in DB2 returning results in seconds or sub-seconds

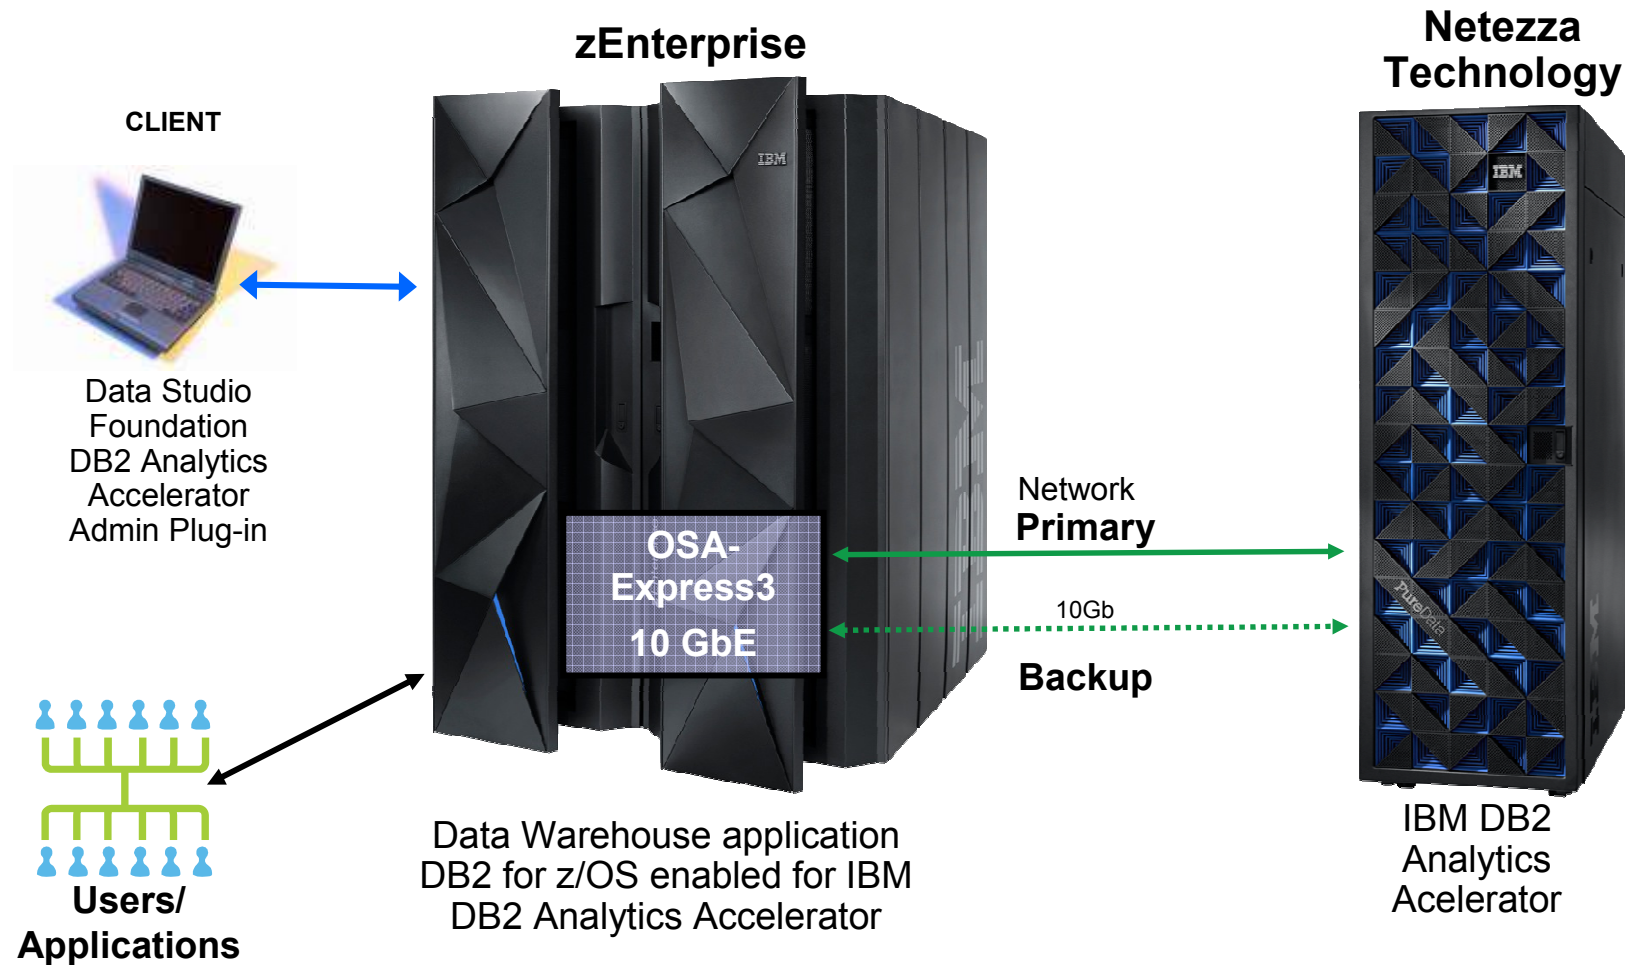
**30 of the Mixed Workload Queries took minutes to hours**

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

**Successfully accelerated the problem queries without affecting the rest**



# IBM DB2 Analytics Accelerator Product Components



*Note: There are several connection options using switches to increase redundancy*

## User Interface

Incremental update UI elements only visible if it has been enabled on the DB2 subsystem via IBM DB2 Analytics Accelerator configuration console

- Start / stop replication process (per subsystem-accelerator pair)
- Enable / disable replication (per table)
- Trace collection
- Information on replication latency and events

**Accelerator: VMNPS03 @ DWADA11** Refresh: Automatic off

Acceleration: Started [Stop](#) Software version: 2.1.3.201203261334 [Transfer new](#) [Apply other](#)

Status: Online Netezza version: 6.0.5-0.F-1.P-5.Bld-21415 [Transfer files](#)

Used space: 4.0 MB of 0.2 TB Credentials valid since: 3/28/12 6:45 PM [Update](#)

Active queries: 0 Trace: [REPLICATION / OFF](#) [Configure](#) [Save](#) [Clear](#)

Replication status: Started [Stop](#) Replication changes: 382,245 / 131,520 20s

▼ **Tables (1 of 1 loaded / 1 of 1 enabled for acceleration)**

[Add...](#) [Alter Keys...](#) [Remove](#) [Load...](#) [Acceleration](#) [Replication](#)

Name like:

Name	Size	Rows	Acceleration	Last Load	Replication Since	Distribution Key	Skew	C
▼ <b>RBE</b>	-	-	<b>1 of 1</b>	<b>1 of 1</b>	<b>1 of 1</b>	-	-	-
TABLE_001	-	-	Enabled	3/15/12 11:26 AM	4/17/12 5:41 PM	Random	0.000	

# Graphical User Interface for High Performance Storage Server Function

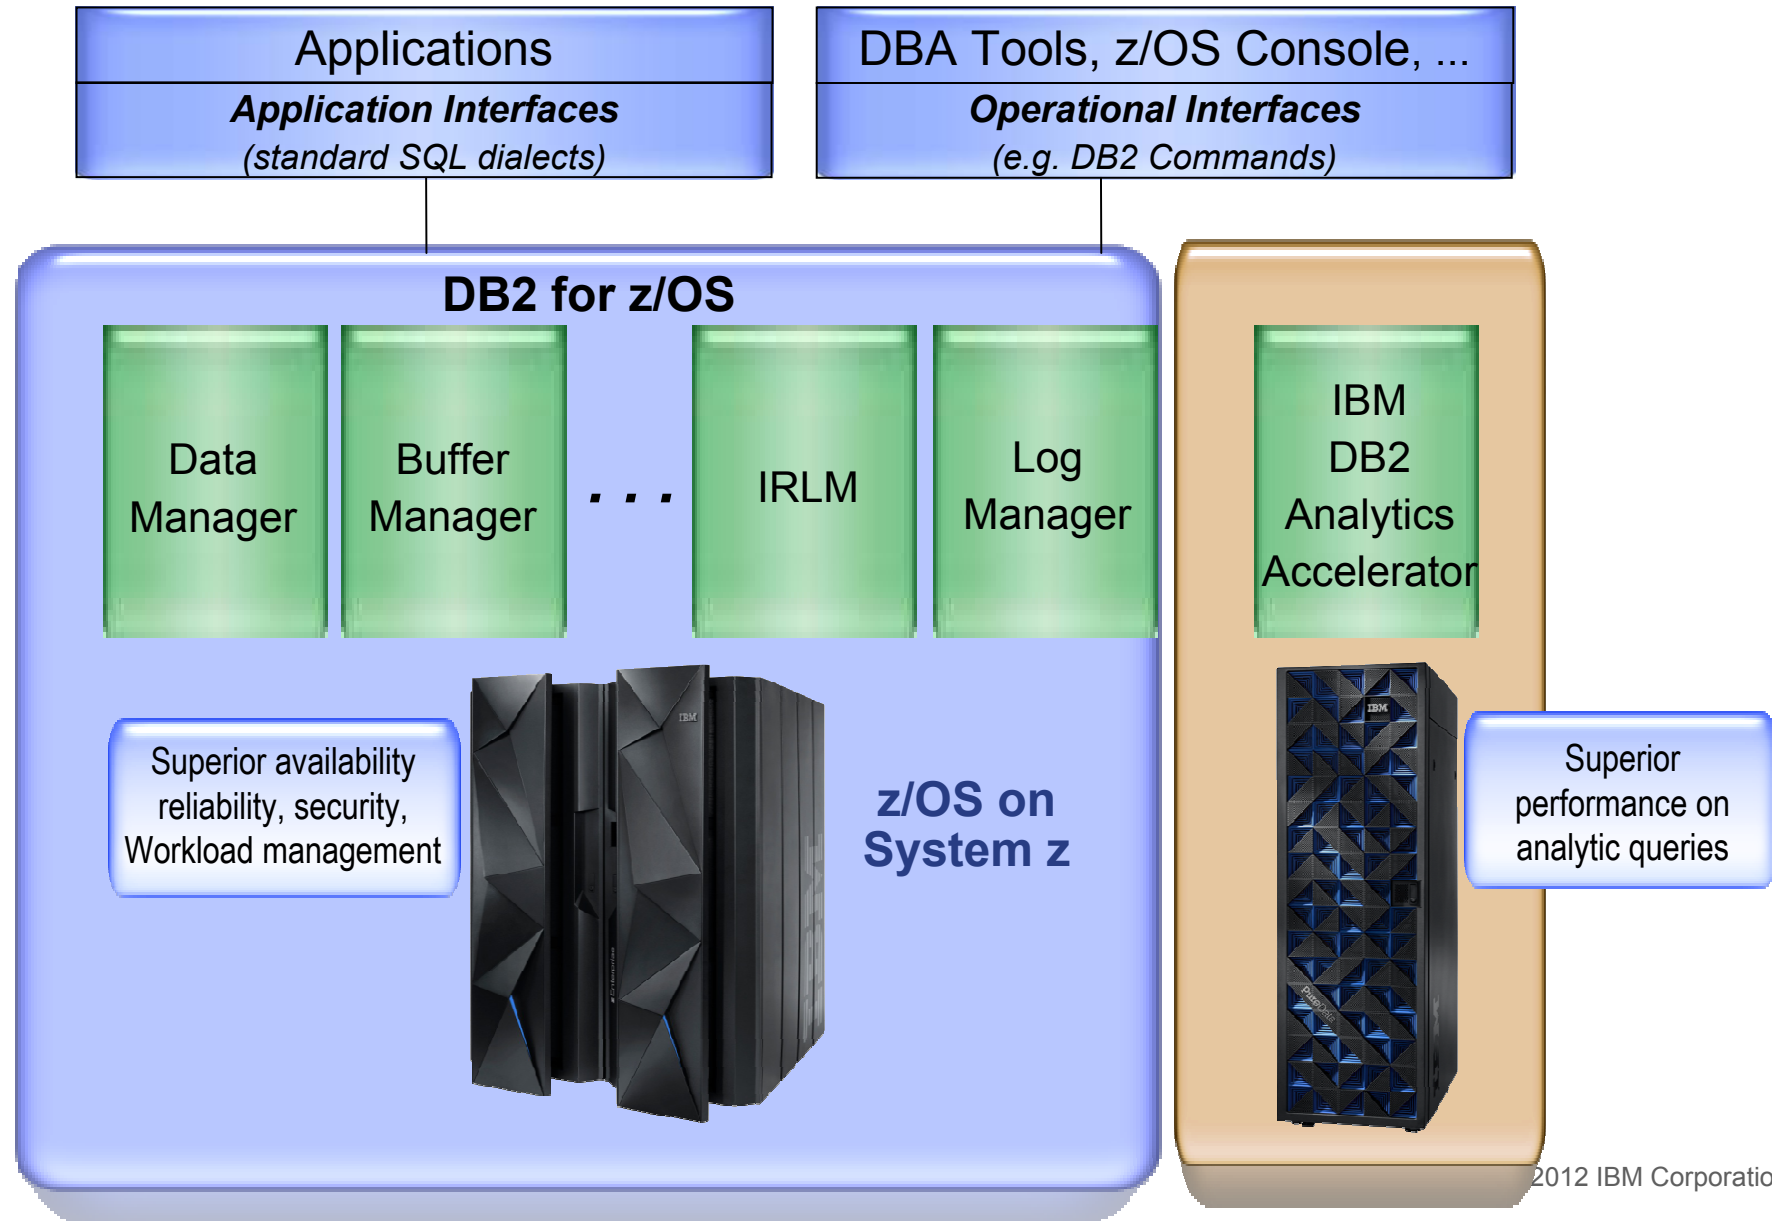
▼ Tables (8 of 8 loaded / 0 of 8 enabled for acceleration)

Name like: type filter text

Name	Size	Rows	Acceleration	Last Load	Moved Partitions	Distribution Key	Sk
<b>TPCH</b>	352 ...	-	<b>0 of 8</b>	<b>8 of 8</b>	<b>1 of 8</b>	-	
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

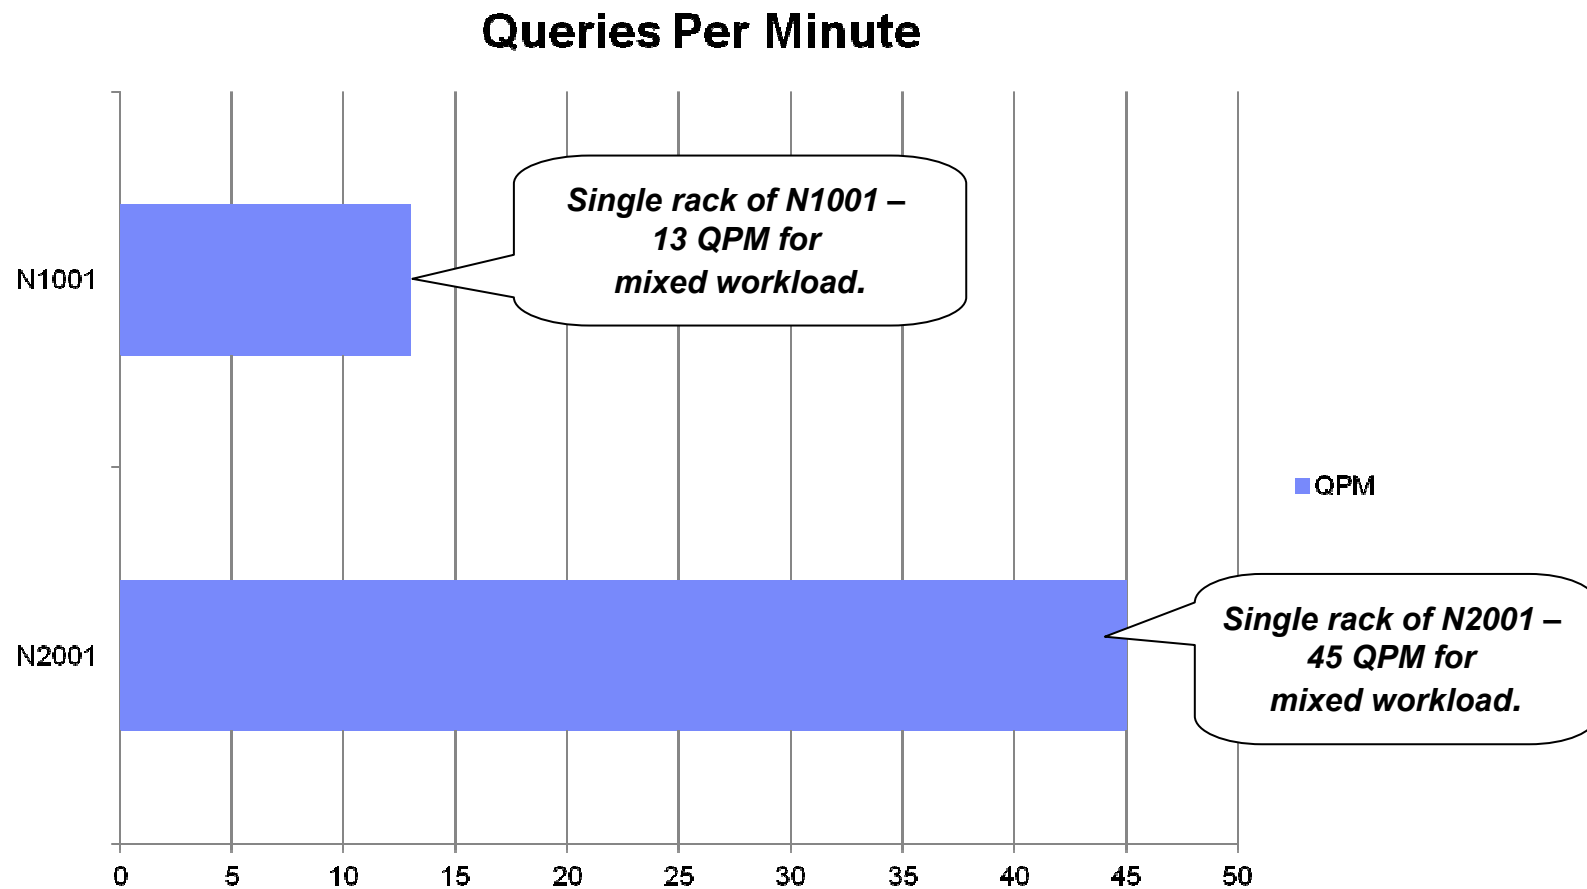
Move to Accelerator...  
 Restore in DB2...

# Deep DB2 Integration within zEnterprise



# Outstanding Mixed Workload Performance

*Benchmark of a mixture of light and power user queries*





## N2001 Hardware Overview



### 12 Disk Enclosures

- 288 600 GB SAS2 Drives
  - 240 User Data, 14 S-Blade
  - 34 Spare
- RAID 1 Mirroring

### 2 Hosts (Active-Passive)

- 2 6-Core Intel 3.46 GHz CPUs
- 7x300 GB SAS Drives
- Red Hat Linux 6 64-bit

### 7 PureData for Analytics S-Blades™

- 2 Intel 8 Core 2+ GHz CPUs
- 2 8-Engine Xilinx Virtex-6 FPGAs
- 128 GB RAM + 8 GB slice buffer
- Linux 64-bit Kernel

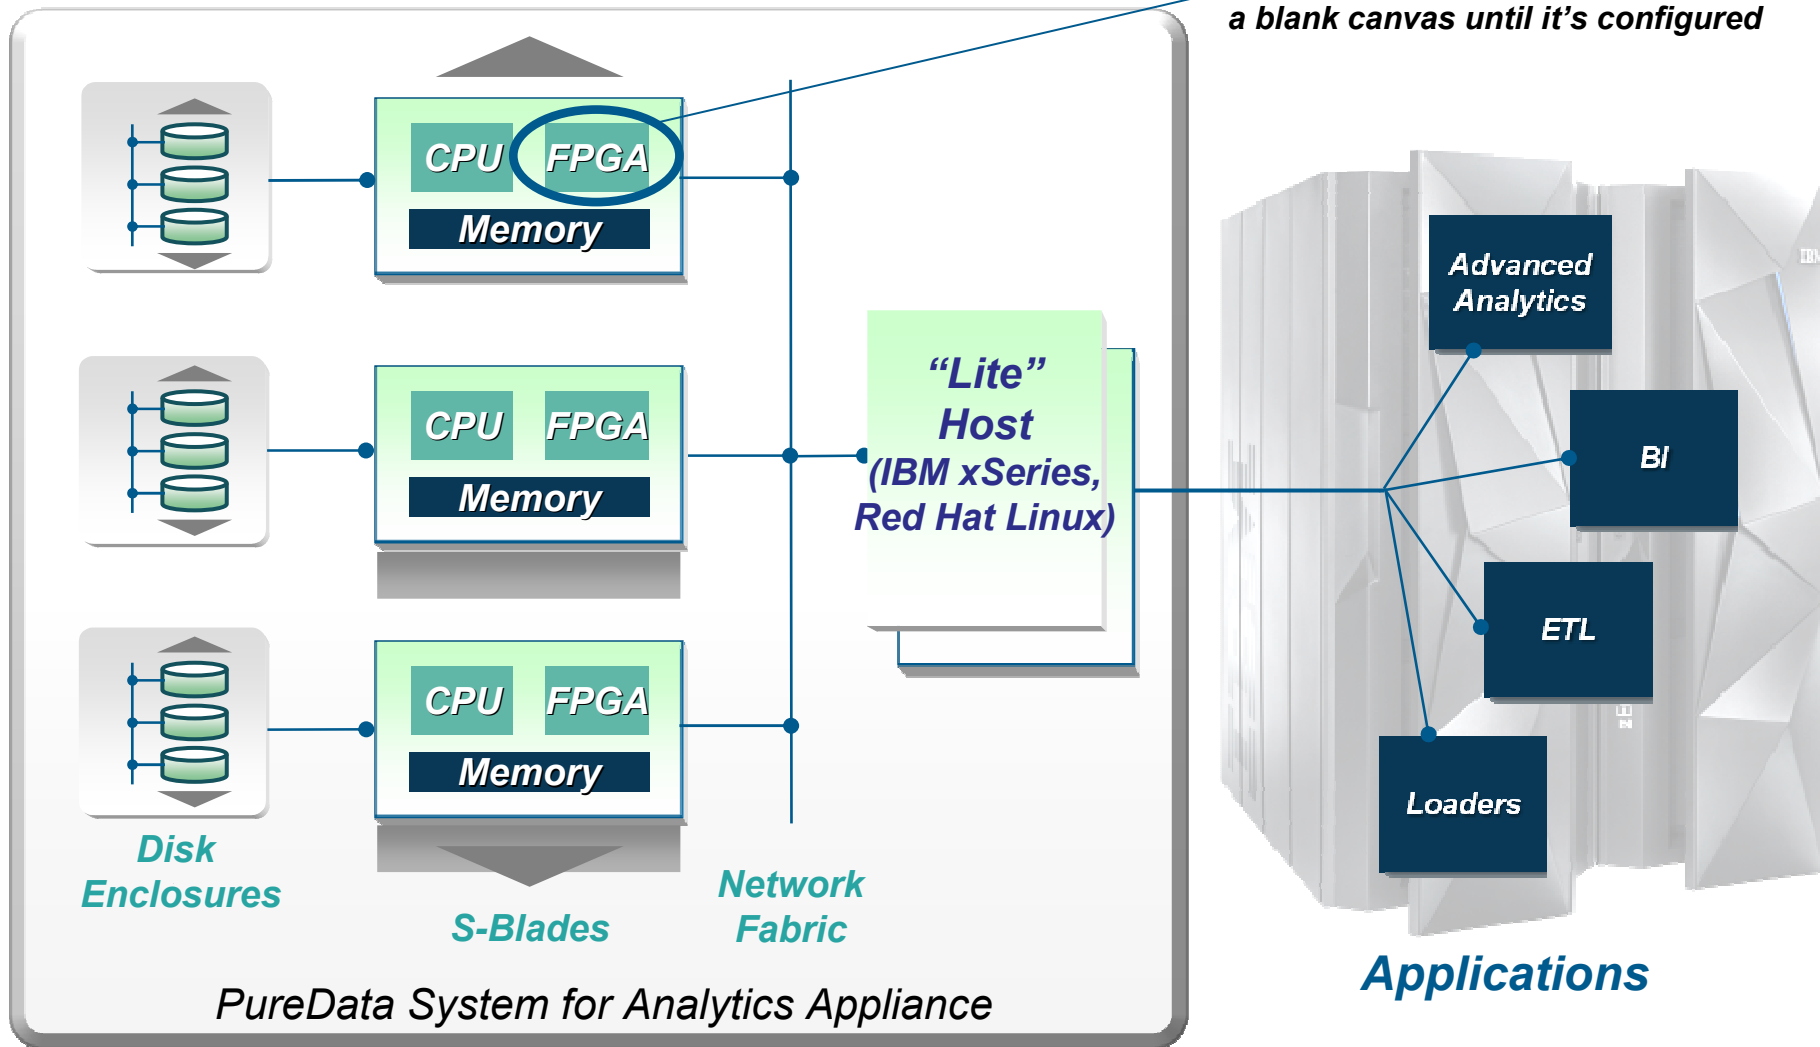
Scales from  
½ Rack to 4 Racks

- *User Data Capacity:* 192 TB\*
- *Data Scan Speed:* 478 TB/hr\*
- *Load Speed (per system):* 5+ TB/hr

- *Power Requirements:* 7.5 kW
- *Cooling Requirements:* 27,000 BTU/hr

# The PureData System for Analytics AMPP Architecture

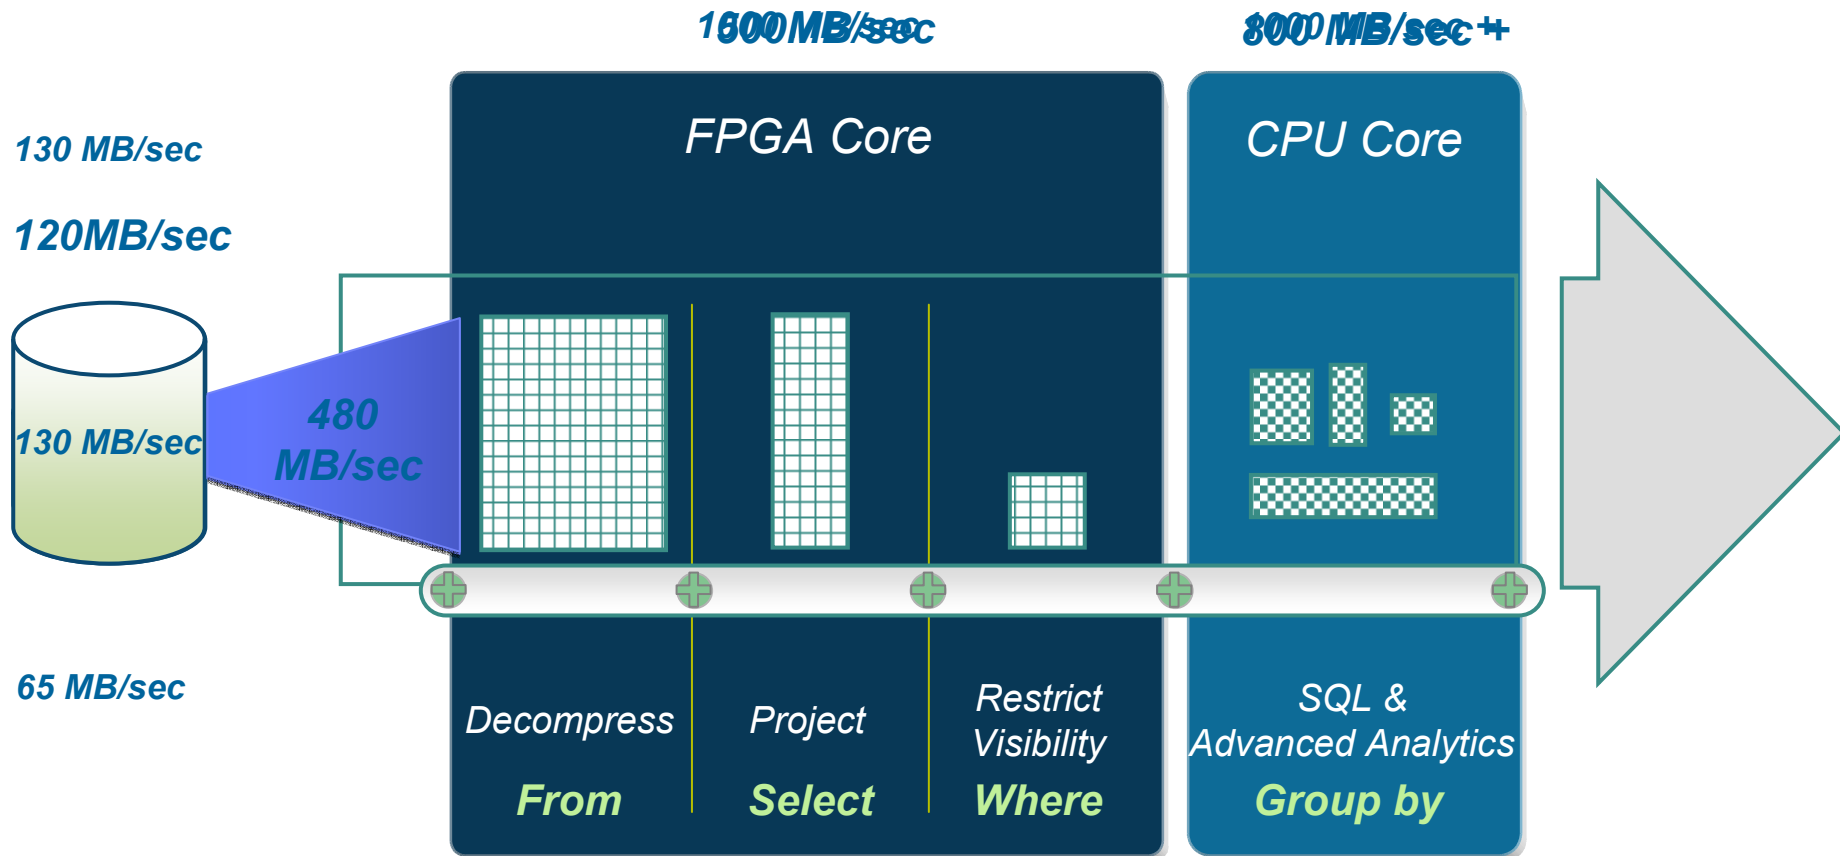
*Field Programmable Gate Array = a blank canvas until it's configured*



# Speed Through Hardware Acceleration

**N2001**

325 MB/sec  
(2.5 drives / core)





## Snippet-Blade™ (S-Blade) Components

- **HX5 Blade**
- **128 GB RAM**
- **16 Intel Cores**



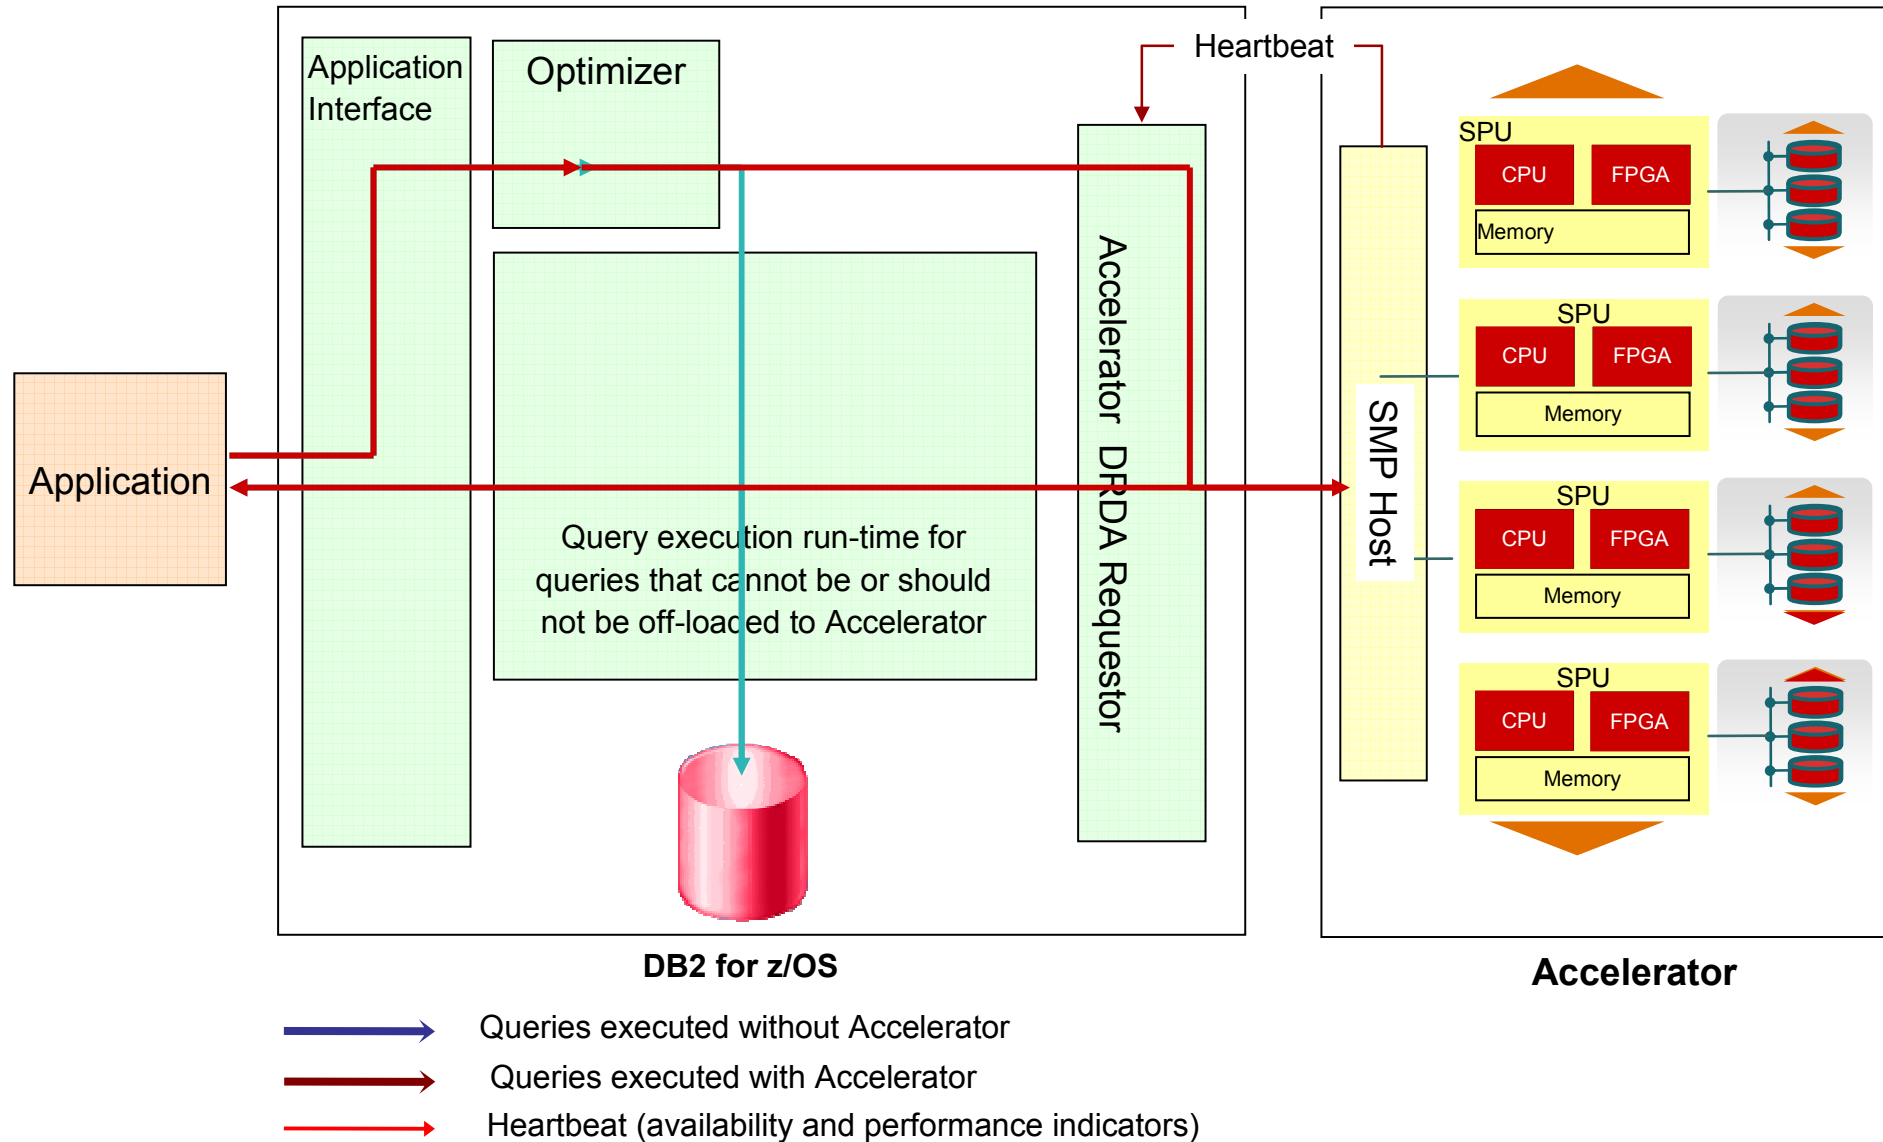
- **BPE4 Side Car**
- **16 GB RAM**
- **16 Virtex-6 FPGA Cores**
- **SAS Controller**

**IBM BladeCenter Server**

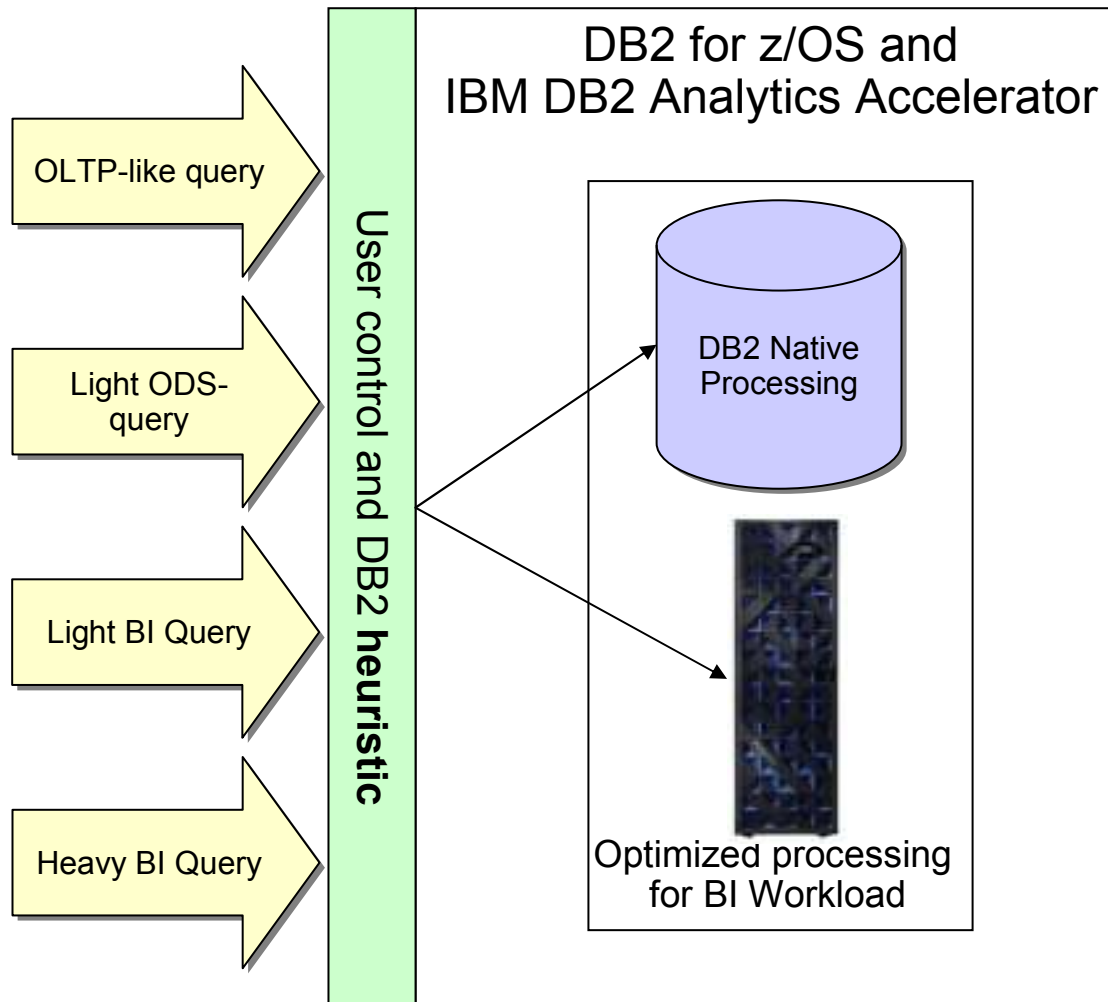
**Netezza DB Accelerator**



# Query Execution Process Flow

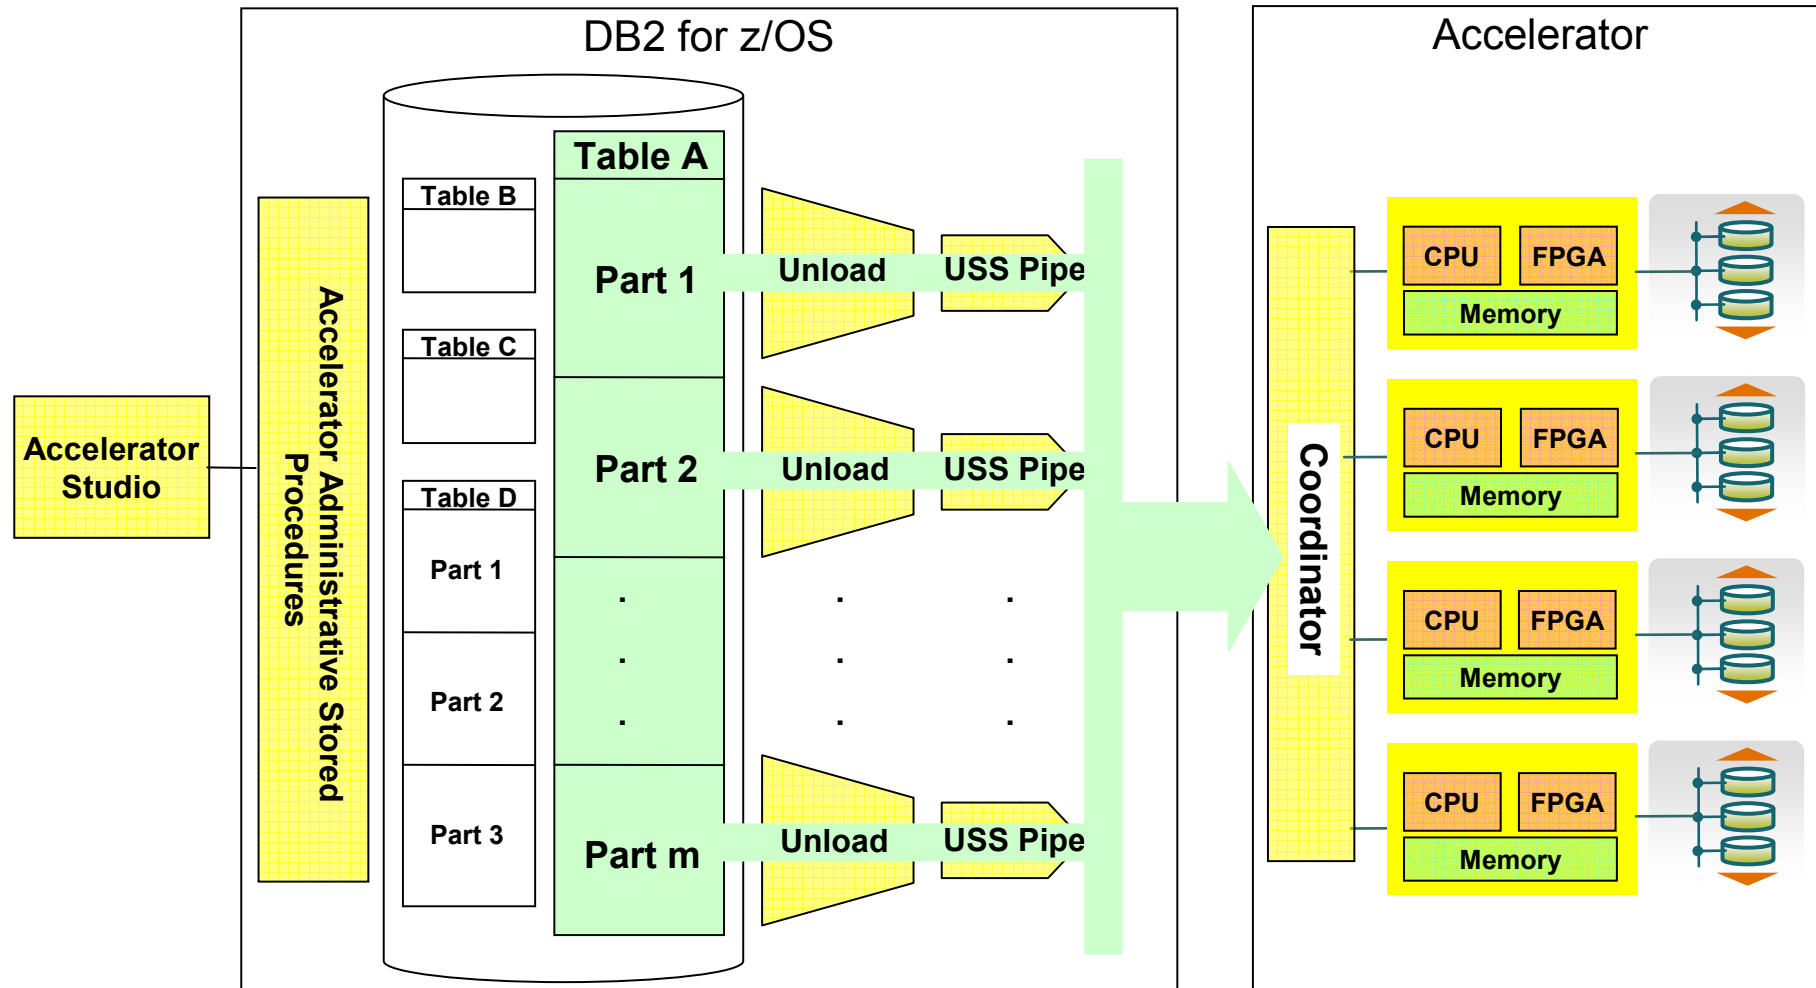


# Workload-Optimized Query Execution



- Single and unique system for mixed query workloads
- Dynamic decision for most efficient execution platform
- New special register  
QUERY ACCELERATION
  - NONE
  - ENABLE
  - ENABLE WITH FAILBACK
- New heuristic in DB2 optimizer

# Accelerator Data Load




- 1 TB / h – can vary, depending on CPU resources, table partitioning, ...
- Update on table partition level, concurrent queries allowed during load
- V2.1 & V3 unload in DB2 internal format, single translation by accelerator

# High Performance Storage Saver

*Reducing the cost of high speed storage*

**Store historic data on the Accelerator only**



**Tables can be resident on:**

1. DB2 Only
2. DB2 and Accelerator
3. Accelerator Only

**When data no longer requires updating, reclaim the DB2 storage**



**Special Registers control behavior**  
(enhanced)

- CURRENT QUERY ACCELERATION
- CURRENT GET\_ACCEL\_ARCHIVE

**Managed by zParms**

**Enhanced Heuristics**

# Save Over 95% of Host Disk Space for Historical Data



*Non historical ( still updating) data represents a small portion of stored data*

*One Quarter = 3.57% of 7 years of data*

*One Month = 1.12% of 7 years of data*

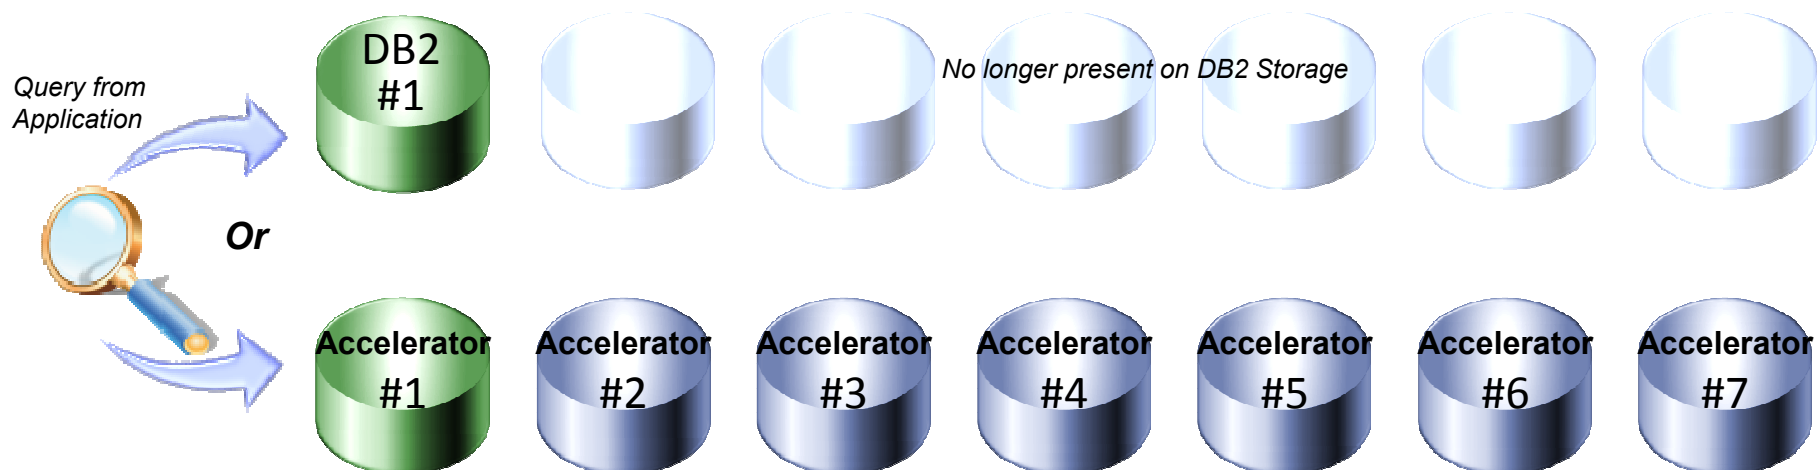
*One Month = 2.78% of 3 years of data*



# High Performance Storage Saver

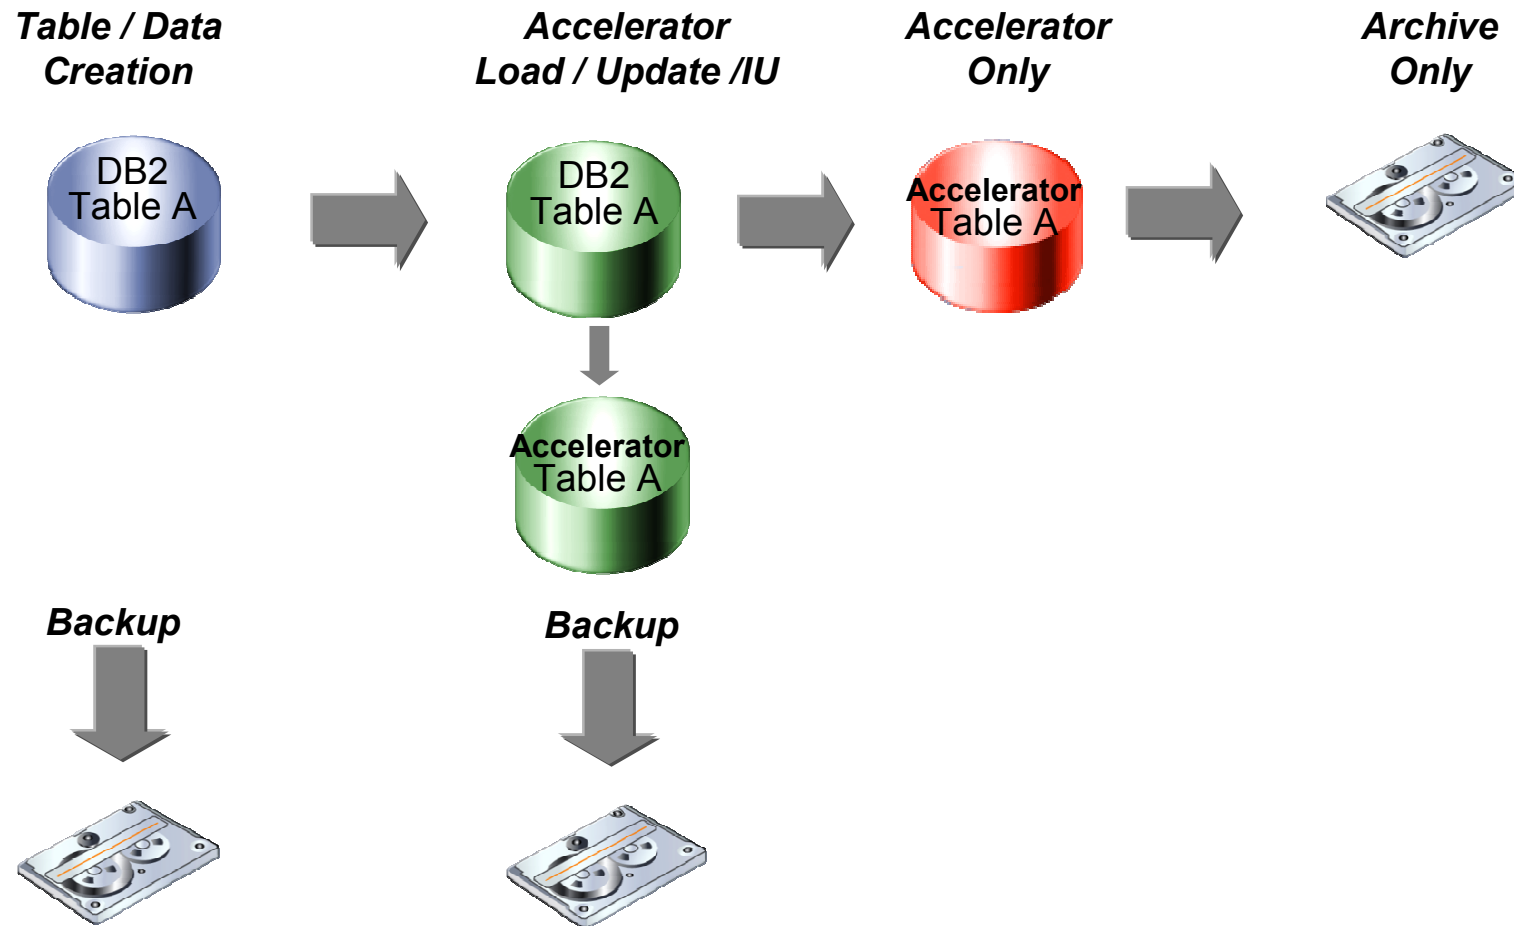
*Reducing the cost of high speed storage*

- **Time-partitioned tables where:**
  - only the recent partitions are used in a transactional context (frequent data changes, short running queries)
  - the entire table is used for analytics (data intensive, complex queries).
- **DB2 partitions are deleted after the being backed up in DB2 and moved to High Performance Storage Saver partitions on the accelerator**



# The Evolution of a High Performance Storage Saver


*High Speed Access to Historical Data*



# Storage options to match data needs

*Optimized in both price and performance for differing workloads*

High Performance Storage Saver	Database Resident Partitions
Accelerator Only	DB2 and Accelerator
<ul style="list-style-type: none"> <li>• Only stored on Accelerator storage (Less Cost)</li> <li>• Optimized performance for deep analytics, multifaceted, reporting and complex queries</li> <li>• Only full table update or full partition update from backup</li> <li>• Same high speed query access transparently through DB2</li> </ul>	<ul style="list-style-type: none"> <li>• Stored on both DB2 and Accelerator storage</li> <li>• Mixed query workload with transactions, single record queries and record updates with deep analytics, multifaceted, reporting and complex queries.</li> <li>• Full table, full partition update, Incremental update from DB2 data</li> <li>• Same high speed query access transparently through DB2</li> </ul>




# The zEnterprise Hybrid Solution

*Mixed Workloads for Next Generation Business Analytics*

## Operational Applications

Transaction Processing



Shared Everything DB

*High volume business transactions and batch reporting running concurrently*

## Analytic Applications

Data warehousing



Shared Nothing DB

*Low volume complex queries context switching*

## Mixed Workload Applications

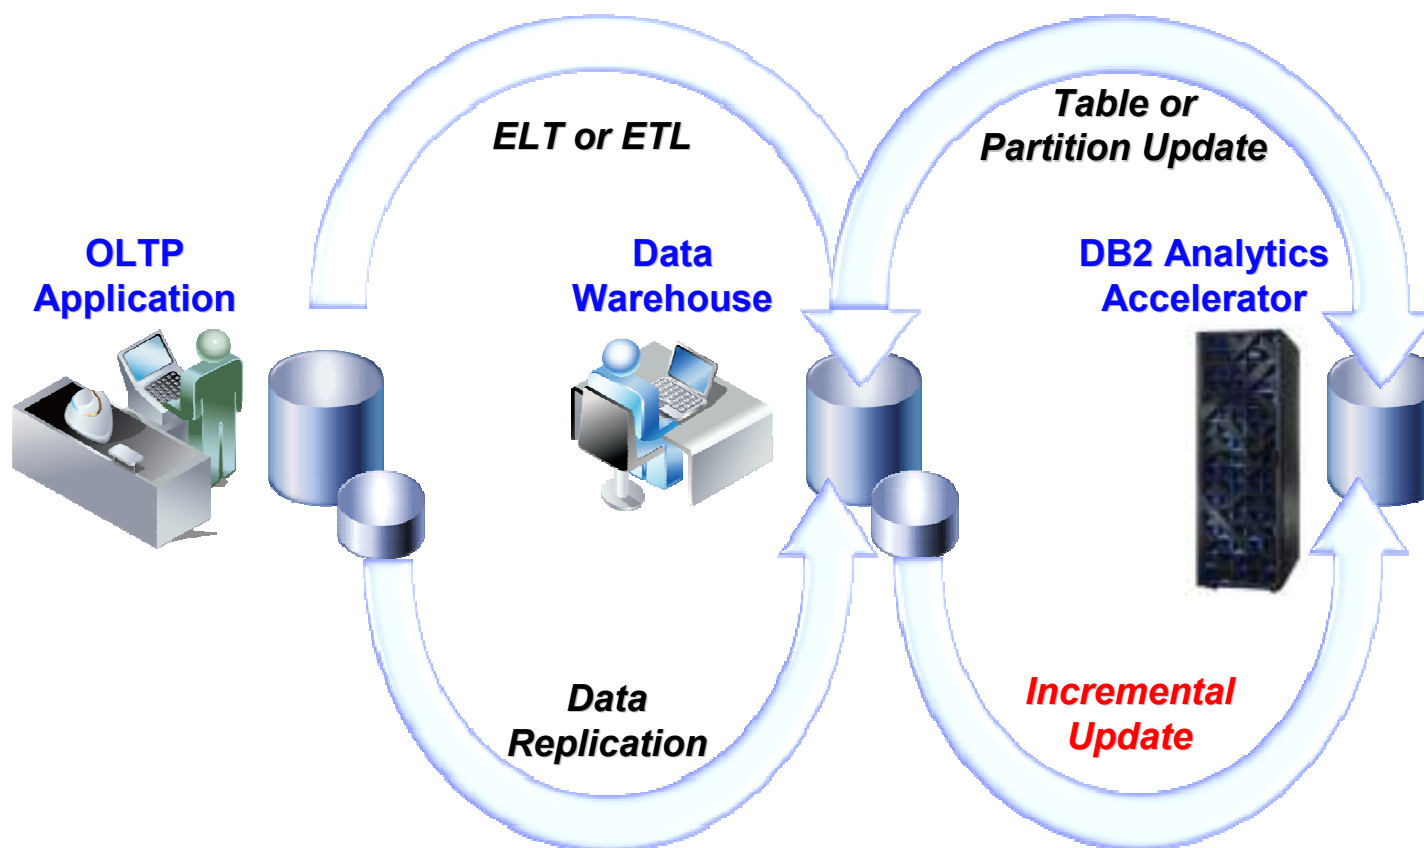
Operational BI



Hybrid DB

*High volume business transactions and batch reporting running concurrently with complex queries*

# Incremental Update

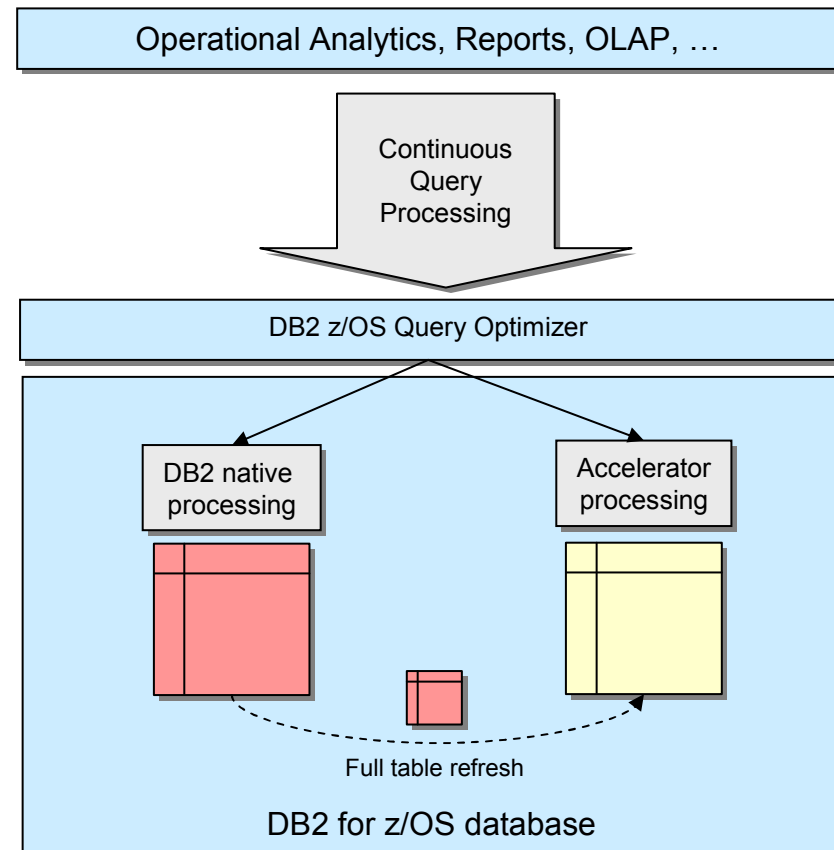



***Synchronizing data to lower data latency from days to minutes/seconds***



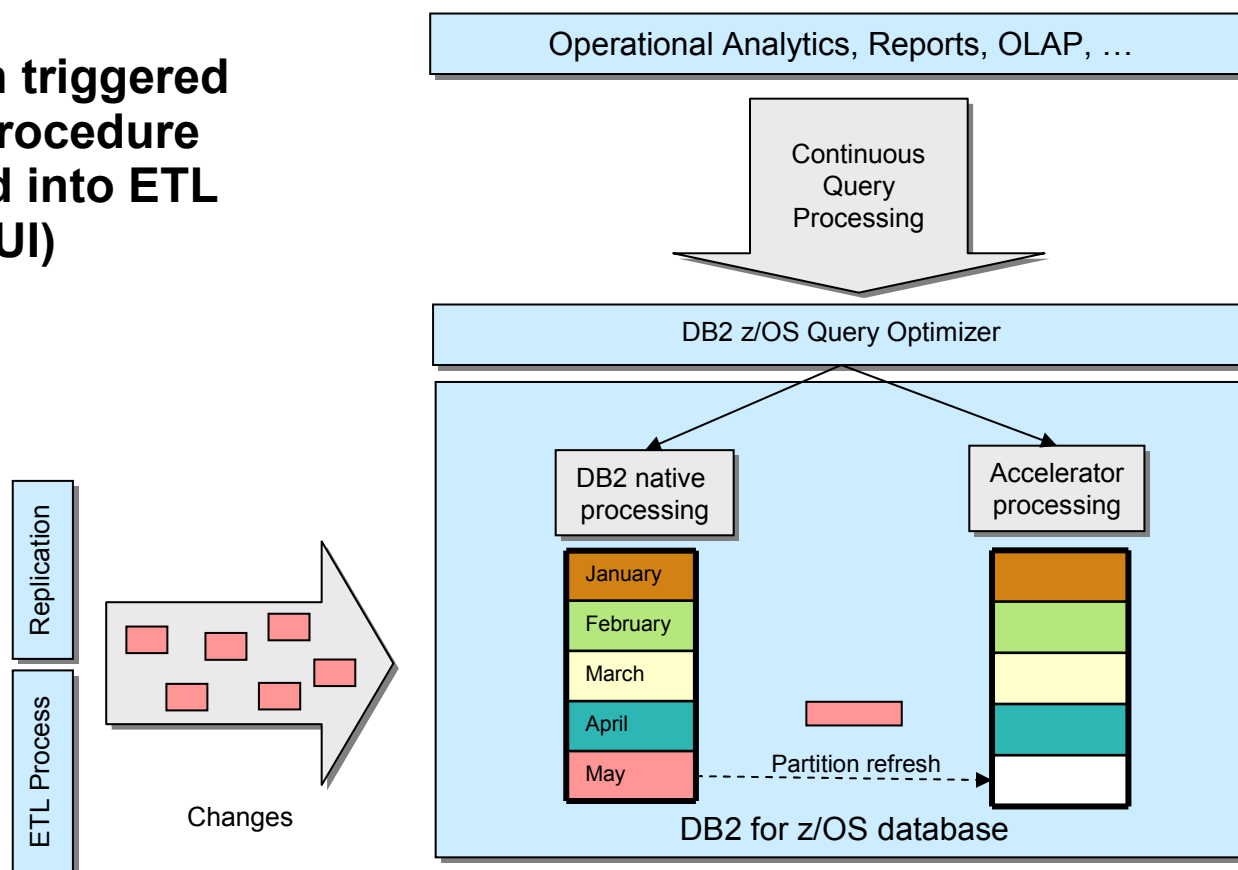
## Option 1: Full Table Refresh

- Changes in data warehouse tables typically driven by scheduled (nightly or more frequently) ETL process
- Data used for complex reporting based on consistent and validated content (e.g., weekly transaction reporting to the central bank)
- Multiple sources or complex transformations prevent propagation of incremental changes
- Full table refresh triggered through DB2 stored procedure (scheduled, integrated into ETL process or through GUI)
- Queries may continue during full table refresh for accelerator



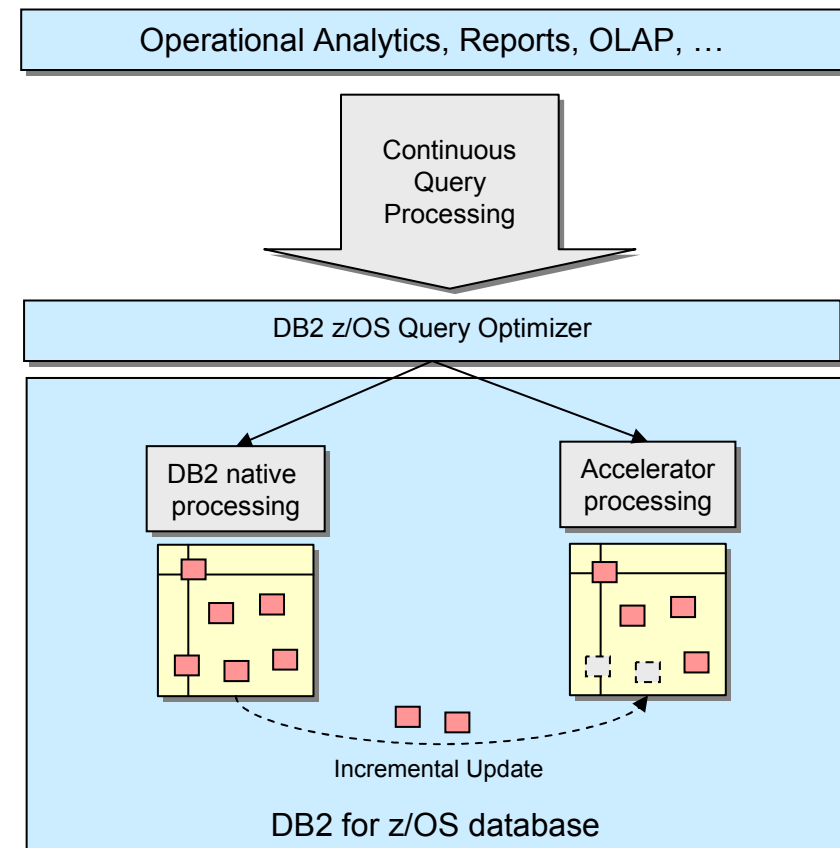

## Option 2: Table Partition Refresh

- Changes in data warehouse table typically driven by “delta” ETL process (considering only changes in source tables compared to previous runs) or by more frequent changes to most recent data
- Optimization of Option 1 when target data warehouse table is partitioned and most recent updates are only applied to the latest partition
- Table partition refresh triggered through DB2 stored procedure (scheduled, integrated into ETL process or through GUI)
- Maintains snapshot semantics for consistent reports
- Queries may continue during table partition refresh for accelerator



## Option 3: Incremental Update

- **Changes in data warehouse tables typically driven by replication or manual updates**
  - Corrections after a bulk-ETL-load of a data warehouse table
  - Continuously changing data (e.g. trickle-feed updates from a transactional system to an ODS)
- **Reporting and analysis based on most recent data**
- **May be combined with Option 1 & 2 (first table refresh and then continue with incremental updates)**
- **Incremental update can be configured per database table**



# N1001 Systems and Sizes

PureData System for Analytics N1001



	<u>002</u>	<u>005</u>	<u>010</u>	<u>015</u>	<u>020</u>	<u>030</u>	<u>040</u>	<u>060</u>	<u>080</u>	<u>100</u>
Cabinets	1/4	1/2	1	1 1/2	2	3	4	6	8	10
S-Blades	3	6	12	18	24	36	48	72	96	120
Processing Units	24	48	96	144	192	288	384	576	768	960
Capacity (TB)	8	16	32	48	64	96	128	192	256	320
Effective Capacity	32	64	128	192	256	384	512	768	1024	1280


**Predictable, Linear Scalability throughout entire family**

Capacity = User Data space  
 Effective Capacity = User Data Space with compression

\*: 4X compression assumed

# N2001 Systems and Sizes

PureData System for Analytics N2001



	-	<u>005</u>	<u>010</u>		<u>020</u>	-	<u>040</u>			
Cabinets		1/2	1		2		4	<i>Watch this space</i>		
S-Blades		4	7		14		28			
Processing Units		56	112		224		448			
Capacity (TB)		24	48		96		192			
Effective Capacity		96	192		384		768			

**Predictable, Linear Scalability throughout entire family**

Capacity = User Data space  
 Effective Capacity = User Data Space with compression

\*: 4X compression assumed

## Connectivity Options

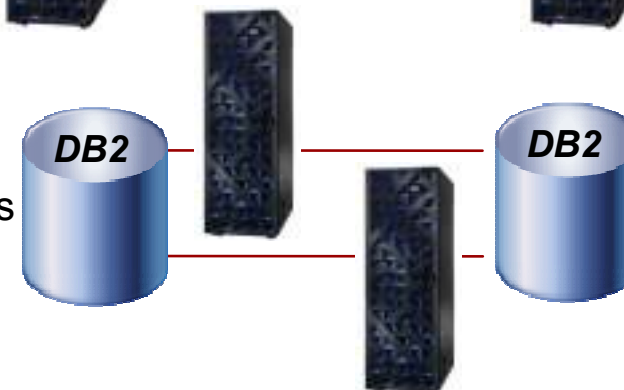
Multiple DB2 systems can connect to a single Accelerator



A single DB2 system can connect to multiple Accelerators



Multiple DB2 systems can connect to multiple Accelerators



The same table can be stored in the multiple Accelerators  
(except High Performance Storage Saver tables)

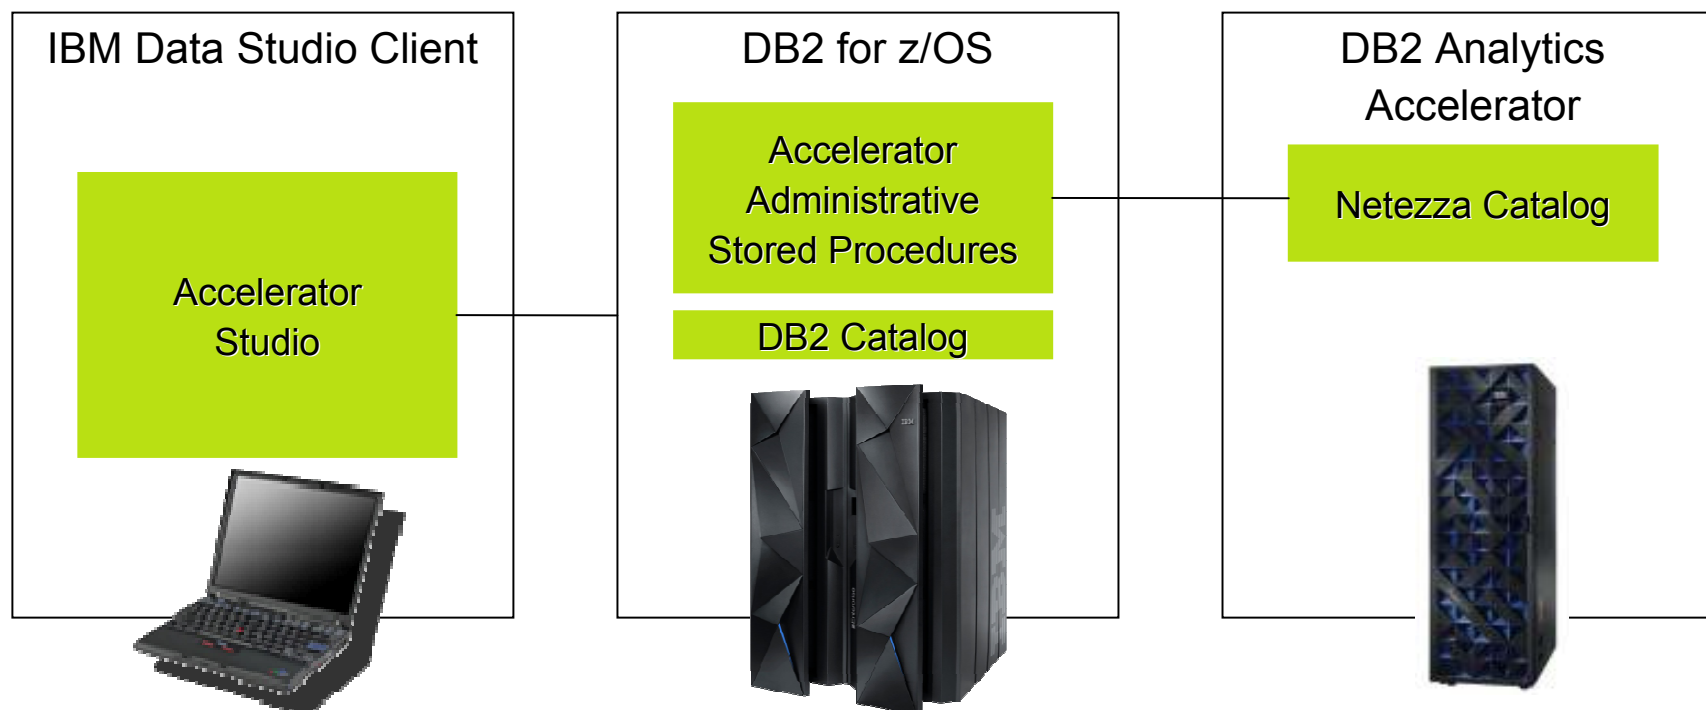
**Better utilization of Accelerator resources**  
**Scalability**  
**High availability**

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



# Analytics Accelerator Table Definition and Deployment



- The tables need to be defined and deployed to the Accelerator before data is loaded and queries sent to it for processing.
  - Definition: identifying tables for which queries need to be accelerated
  - Deployment: making tables known to DB2, i.e. storing table meta data in the DB2 and Netezza catalog.
- IBM DB2 Analytics Accelerator Studio guides you through the process of defining and deploying tables, as well as invoking other administrative tasks.
- IBM DB2 Analytics Accelerator Stored Procedures implement and execute various administrative operations such as table deployment, load and update, and serve as the primary administrative interface to the Accelerator from the outside world including Accelerator Studio.

# ***DB2 Analytics Accelerator***

*Raising the Bar in Analytics*

***Large analytic systems at dramatically faster speeds***

## ***PureData for Analytics N2001***

- **3x Faster than N1001**
- **Increased Throughput**
- **50% more storage /rack**
- **Improved Resiliency**

*Over 3 times the performance, 50% more storage capacity, in the same footprint, for a about a 40% increase in price*

## ***DB2 Analytics Accelerator V3.1***

- **High Performance Storage Saver**
- **Incremental Update**
- **zEnterprise EC12 Support**
- **Query Prioritization**
- **UNLOAD Lite**

***41 customers with 56 systems are experincing the speed of analytics on z***

# ***Build a System z Trusted Analytic System***

*Reduce the cost of host storage for historical data by 95%!*

## **Historical**

Most data in an analytic system is historical and not subject to change. Most data can be in a Storage Saver and maintain trusted performance and security

## **High Performance**

All aggregate queries run at the same high speed as any accelerator supported query

## **Low Latency Data**

Tables and partitions that require updating will be able to be updated by incremental update, table load or partition load

## Why Both?

*Marrying the best of both worlds*

### IBM PureData For Analytics N2001



**Focused Appliance**

### IBM System z



**Mixed Workload System**

*Capitalizing on the strengths of both platforms while driving to the most cost effective, centralized solution - destroying the myth that transaction and decision systems had to be on separate platforms*

**Very focused workload**

**Very diverse workload**

# Tailored to your needs

*A Hybrid Solution*

## IBM PureData For Analytics N2001

### Focused Appliance

- Appliance with a streamlined database and HW acceleration for performance critical functionality
- Price/performance leader
- Speed and ease of deployment and administration
- Optimized performance for deep analytics, multifaceted, reporting and complex queries

## IBM System z with IBM DB2 Analytics Accelerator

### Mixed Workload System

- Mixed workload system z with operational transaction systems, data warehouse, operational data store, and consolidated data marts.
- Unmatched availability, security and recoverability
- Natural extension to System z to enable pervasive analytics across the organization.
- Speed and ease of deployment and administration

**Simplicity**

*The right mix of simplicity and flexibility*

**Flexibility**



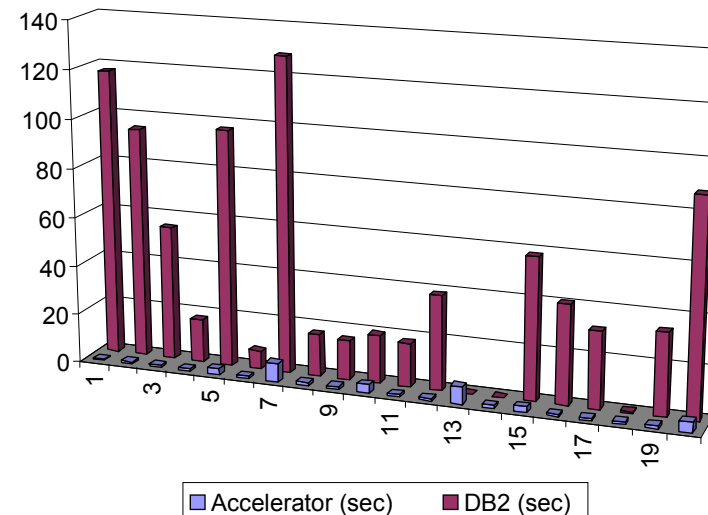
# Accelerating SAP for more business value



18 million records InfoCube, 20 dedicated queries

- Enhances the SAP with DB2 for z/OS
- Accelerates SAP NetWeaver BW
- Dramatic decrease in elapsed **time for SAP** BW ad-hoc reporting

No	Description	Records Read	Records Returned	DB2 (sec)	Accelerator (sec)	Acceleration Factor
1	Simple mass aggregation	17116647	21	117	0.78	150
2	Query #1 + 70% filter	11980812	21	94.2	0.86	110
3	Query #1 + 30% filter	5133708	21	54.8	0.82	67
4	Query #1 + 10% filter	1710293	21	17.6	0.87	20
5	Skewed data, low filtering	10790019	21	96.8	2.47	39
6	Skewed data, high filtering	24	14	7.28	0.83	9
7	Many restrictions	3805941	21	128	7.65	17
8	Navigational attributes	823646	21	17.1	1.27	13
9	Navigational attributes + selective condition	811	21	15.8	1.17	14
10	Open value ranges	2006	21	19.6	3.52	6
11	Hierarchy	1653981	21	17.6	0.97	18
12	Hierarchy + selective condition	55068	21	38.6	0.98	39
13	Restricted key figures on 2 dimensions	1314964	1948	207	7.22	29
14	Query #14 + hierarchy	132564	1499	> 1000	1.27	> 787
15	Calculated key figures (OLAP)	5321586	10	57.8	2.37	24
16	OR linked values	6212609	13	40.5	0.92	44
17	Non uniform data distribution	11016253	13	31.2	0.99	32
18	Selective line item	1724	1706	0.71	1.17	0.6
19	Non-selective line item	115481	68619	33.8	1.36	25
20	All together	3087692	468	87.7	4.42	20






# Traditional BI

*Data Warehouse, Data Mart, ODS*

Enhancing the analytics system  
by transparently accelerating  
complex, high row count  
queries



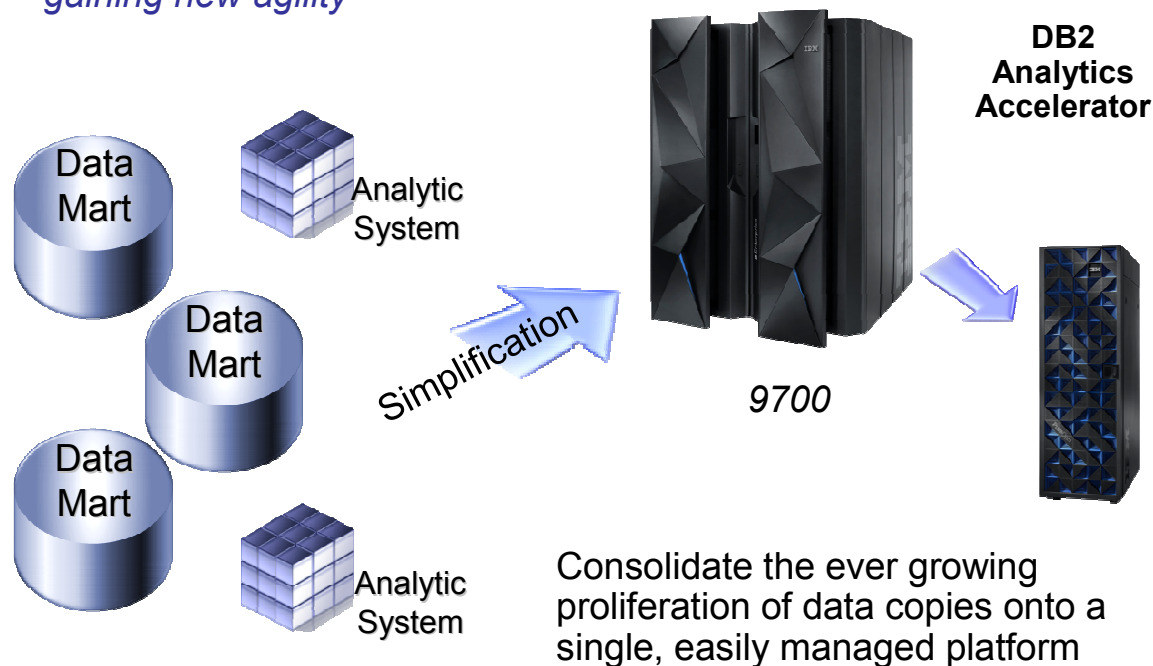
## Benefits

- Enable true mixed workload capabilities for analytic applications
- Realize significant increases in availability, security, recoverability, and virtualization
- Eliminate the need to move data from the operational platform to other platforms for analysis
- Accelerate long running DB2 for z/OS queries from minutes to seconds for greater business value
- The forgotten query: consider queries previously set aside due to performance challenges?
- Avoid costs and efforts to tune individual queries

# Data Mart Consolidation

*Utilize virtualization to optimize the use of resources while reducing costs and gaining new agility*

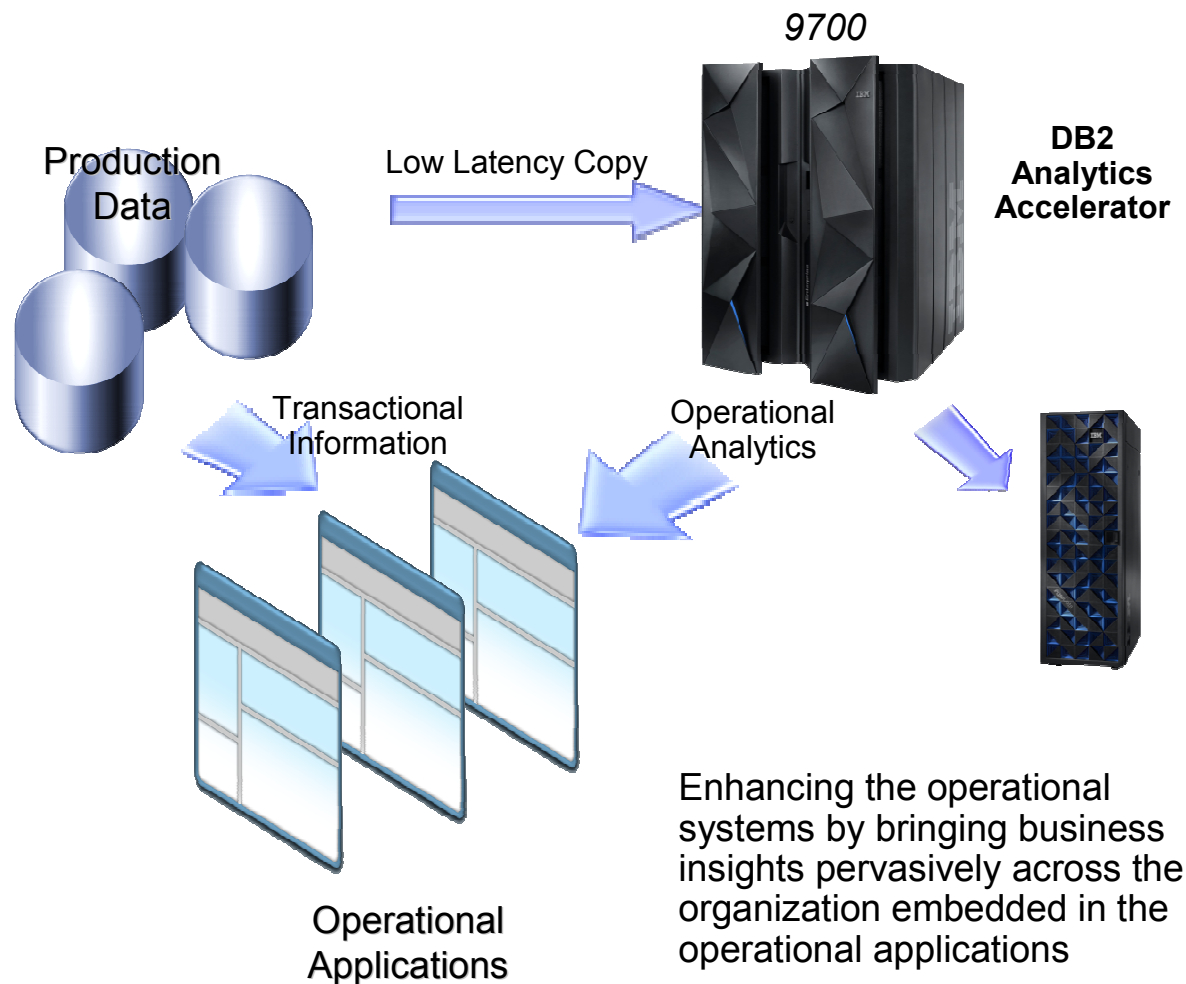
*A single platform to manage and administer*



## Benefits

- Consolidate isolated islands of data on one secure Hub
- Simplify management of costly/complex data copies while retaining the isolation benefits of individual platforms
- Provide consistency to informational data
- Increase time to value to deploy new Data Marts
- Enable application queries which would you prefer to run with more real-time data on System z

# Operational BI




## Benefits

- Provide analytic information at the point of decision enabling fact-based decisions
- Deliver new insights that help business users in an operational application
- Pervasively enable decision makers and other end users across the organization
- Enable customers with historical information to increase loyalty and sales
- Accelerate long running DB2 for z/OS queries from minutes to seconds for greater business value with the DB2 Analytics Accelerator

# Accelerator for Operational Reporting

Dramatically improve complex operational reporting. Reporting is no longer the simple sorting and summarizing of daily interactions, it now requires the analysis of the businesses key performance indicators



*Keep new reporting capabilities onto a single, easily managed platform*



## Benefits

- Time and agility gained through more timely delivery of complex information to the business
- Consolidate reporting where the majority of data being analyzed resides (z/OS)
- Business benefits of analytics for queries previously set aside
- Fast time to value with transparent integration into existing applications
- Easy-to-install appliance add-on to DB2 for database query acceleration

# Workload Assessment

- **Customer**
  - Collects information from dynamic statement cache, supported by step-by-step instruction and REXX script (small effort for customer)
  - Upload compressed file (up to some MB) to IBM FTP server
- **IBM**
  - Import data into local database
  - Quick analysis based on known DB2 Analytics Accelerator capabilities

**Key contact:** Data Warehouse System z/Germany/IBM

## Report for a first assessment:

- **Acceleration potential for**
  - **Queries**
  - **Estimated time**
  - **CP cost**

**IBM Smart Analytics Opt**  
Center of Excellence, Datawarehouse on

**Query Summary**

	Total	With potential	Uncertain	W/o potential
Queries	23	11 (48%)	5 (22%)	7 (30%)
Query Blocks	23	11 (48%)	5 (22%)	7 (30%)
Elapsed Time	144801.47	106821.61 (74%)	8150.21 (6%)	29829.66 (21%)
CPU Time	21300.25	11420.12 (54%)	1453.14 (7%)	8426.98 (40%)


Queries	23	100%
... no eligible blocks	7	30%
... not read only	0	0%
... with very large dim.	1	4%
... with quant. pred.	0	0%

Leaf Query Blocks	23	100%
... with aggregations	23	100%
... with UDFs	0	0%
... with unsupported functions	6	26%
... with unsupported join types	0	0%
... with very large dimensions	1	4%

Start trace time: Apr 2, 2010 9:41 AM  
End explain time: Apr 2, 2010 4:37 PM  
Min stmt cached: Apr 2, 2010 9:44 AM  
Max stmt cached: Apr 2, 2010 4:37 PM

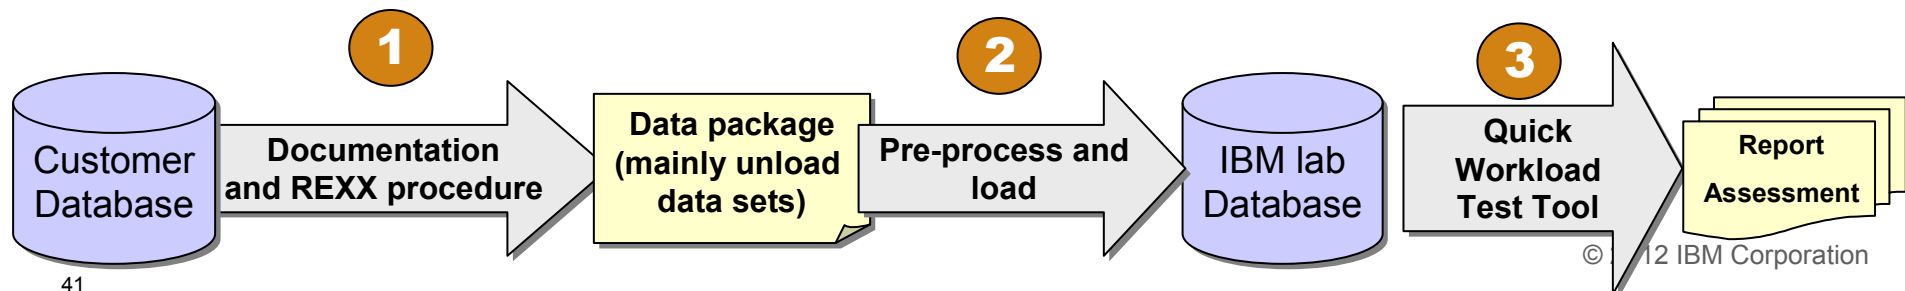
Disclaimer: Information provided in this document is for information purposes only and does not guarantee characteristics nor imply supported features of IBM products. V20100331  
Schema: IGBURG -- Date: 4/12/10 -- Page: 1



Query blocks with acceleration potential  
Query blocks with uncertain potential  
Query blocks without acceleration potential



Elapsed time with acceleration potential  
Elapsed time with uncertain potential  
Elapsed time without acceleration potential



**Thank You**