


Extreme Performance for Complex Business Analysis

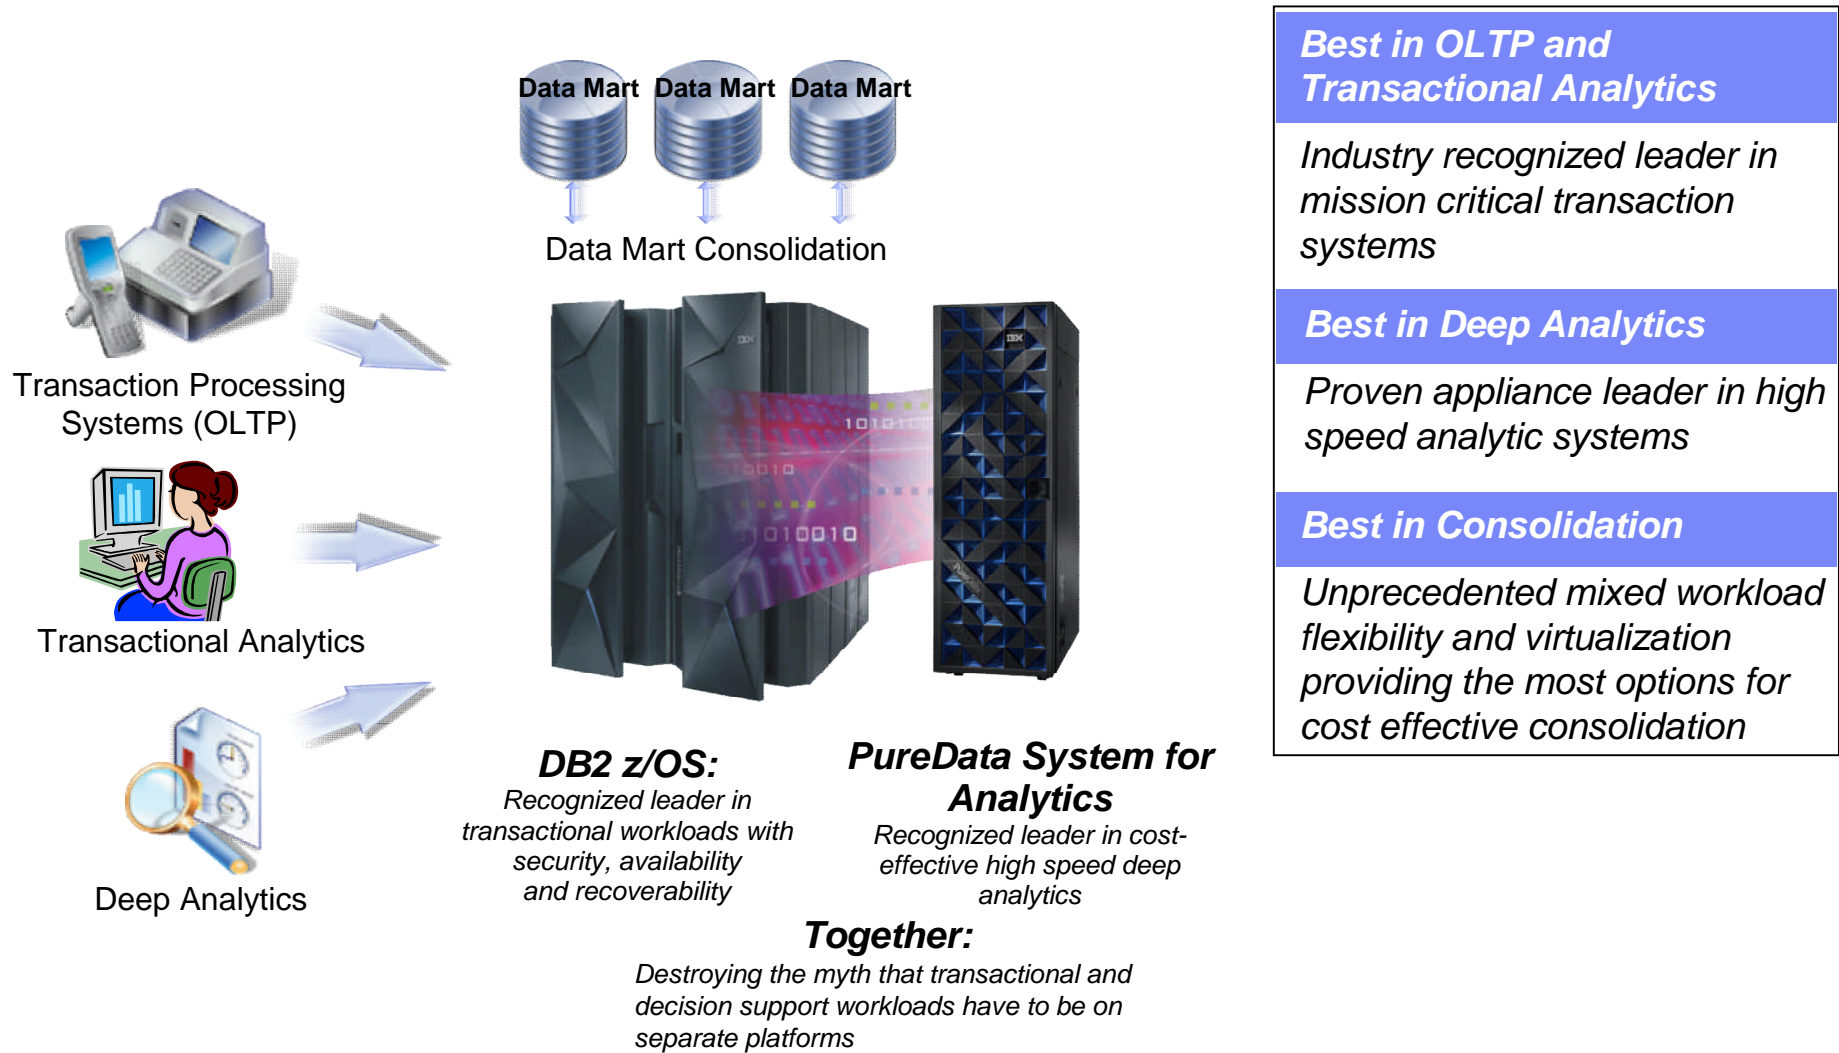
IBM DB2 Analytics Accelerator

...powered by Pure Data Systems for Analytics (Netezza technology)



Natalie Rogers, System z IT Specialist, nprogers@us.ibm.com

Creating the Hybrid Data Server – PureData System for Analytics (Netezza) and System z



DB2 Analytics Accelerator for z/OS Blending zEnterprise and Netezza technologies

A high performance analytics accelerator appliance for IBM zEnterprise, delivering dramatically faster complex business analysis transparently to all users.



Fast

Complex queries run up to 2000x faster while retaining single record lookup speed

Cost Saving


Eliminate costly query tuning while offloading complex query processing

Appliance

No applications to change, just plug it in, load the data, and gain the value

What's new in DB2 Analytics Accelerator V3

- ✓ *Lowering the cost of historical data*
- ✓ *Better decisions through lower latency of data*
- ✓ *Dramatic improvement in scale and growth opportunities*
- ✓ *Lowering the cost of analytic computing*



High Performance Storage Saver

- Significantly reduces the cost for storage resources
- Option to store data only once in the accelerator

Incremental Update

- Data changes are propagated to the accelerator as they happen
- Uses change data capture technology
- Extends the accelerator use to reporting on operational data

New optimization

- Tables or partitions refresh much faster and less resources intensive
- Optimized unloading data from DB2

High Capacity

- Capacity has been extended to 1.28 PB for a single Accelerator

New functions

- More queries eligible for acceleration

Fast Time to Value

- IBM DB2 Analytics Accelerator (N1001-010)
 - ➔ Production ready - 1 person, 2 days
- Table Acceleration Setup ... 2 Hours
- Initial Load Performance ...
 - ➔ 400 GB "Loaded" in 29 Min
 - 570 million rows (Loads of 800GB to 1.3TB/Hr)
- Actual Query Acceleration ... 1908x faster
 - ➔ 2 Hours 39 Minutes to 5 Seconds
- CPU Utilization Reduction
 - ➔ up to 35%

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”

“breakthrough technology allows us to free up MIPS that we can use for future growth”

Deep DB2 Integration within zEnterprise

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

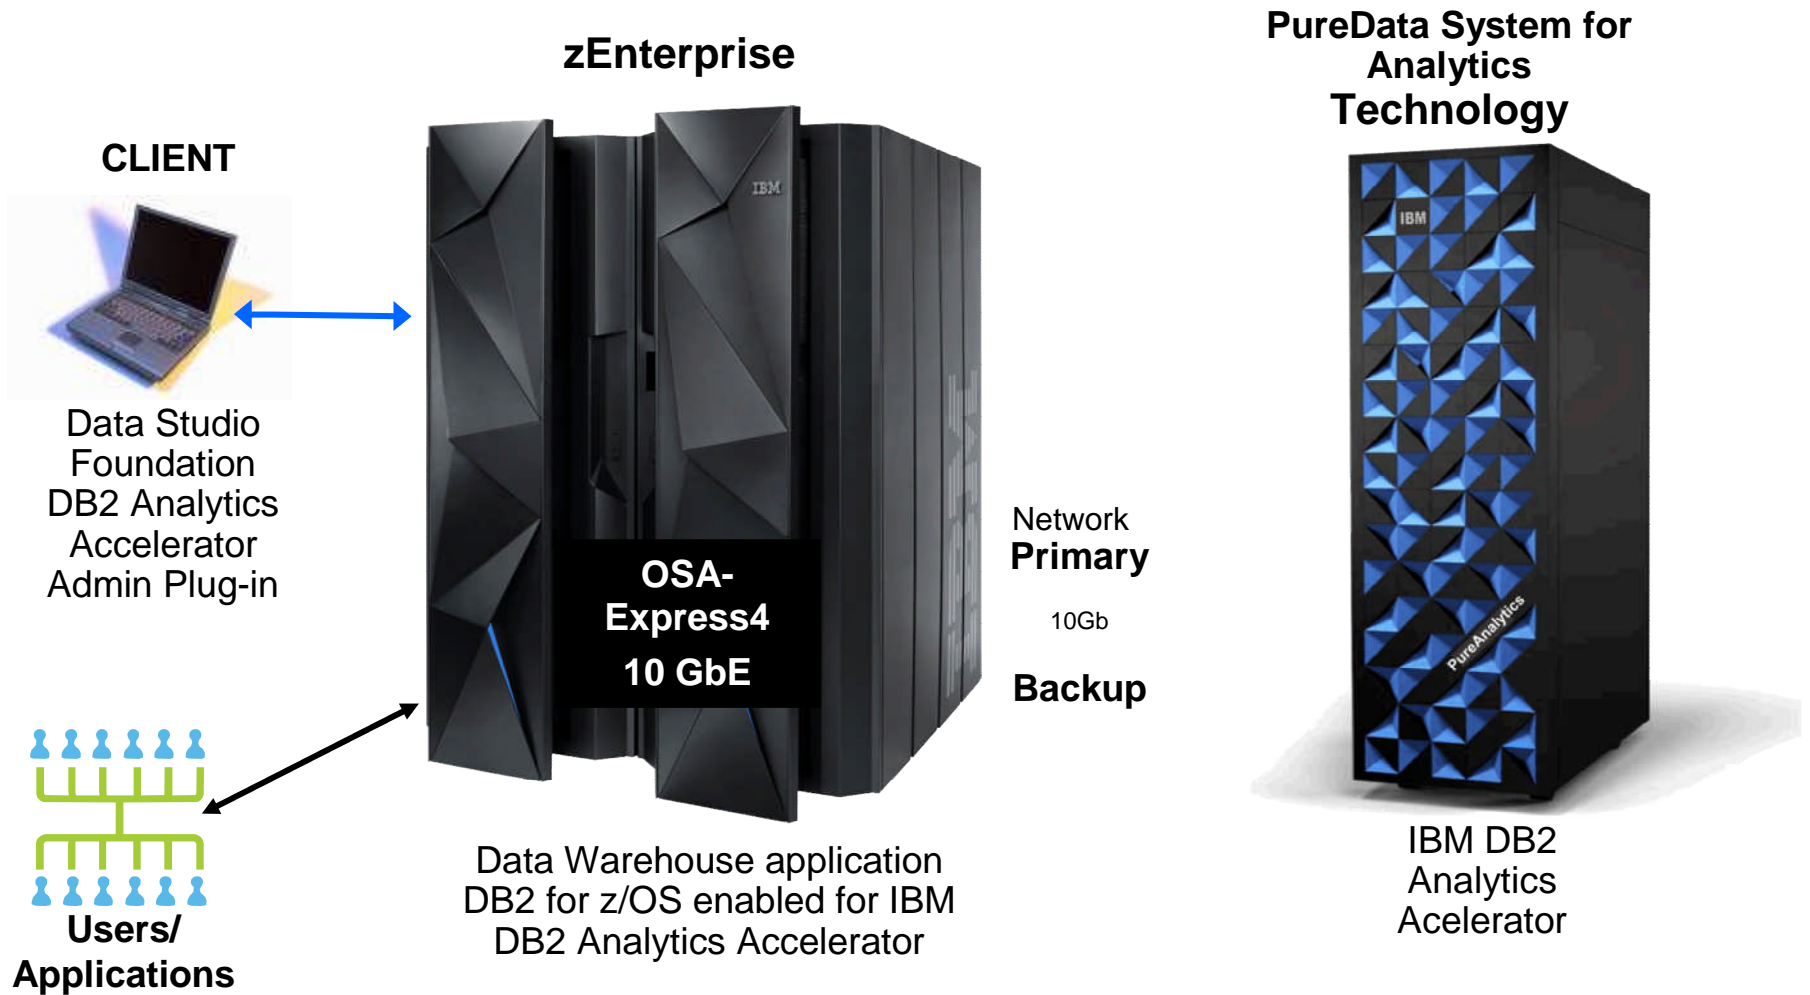
**z/OS on
System z**

IBM
DB2
Analytics
Accelerator

Superior
performance on
analytic queries

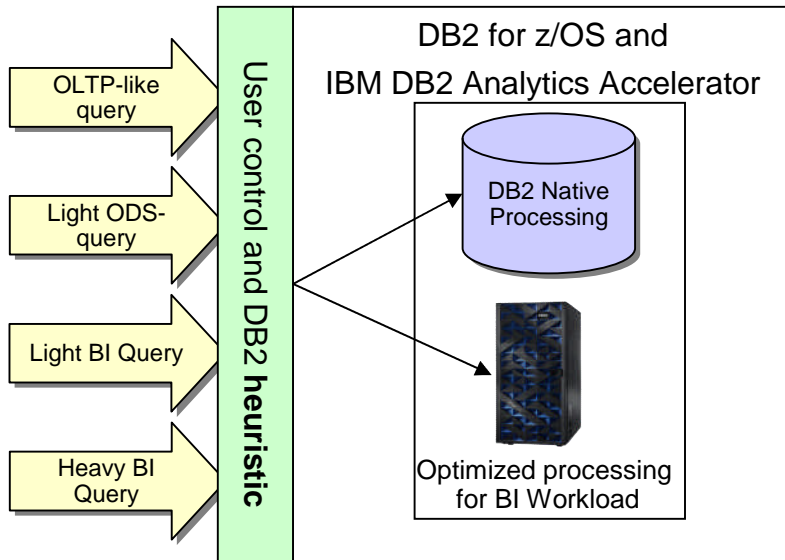
**PureData System
for Analytics**

IBM DB2 Analytics Accelerator Product Components



Query routing analysis

Values for CURRENT QUERY ACCELERATION

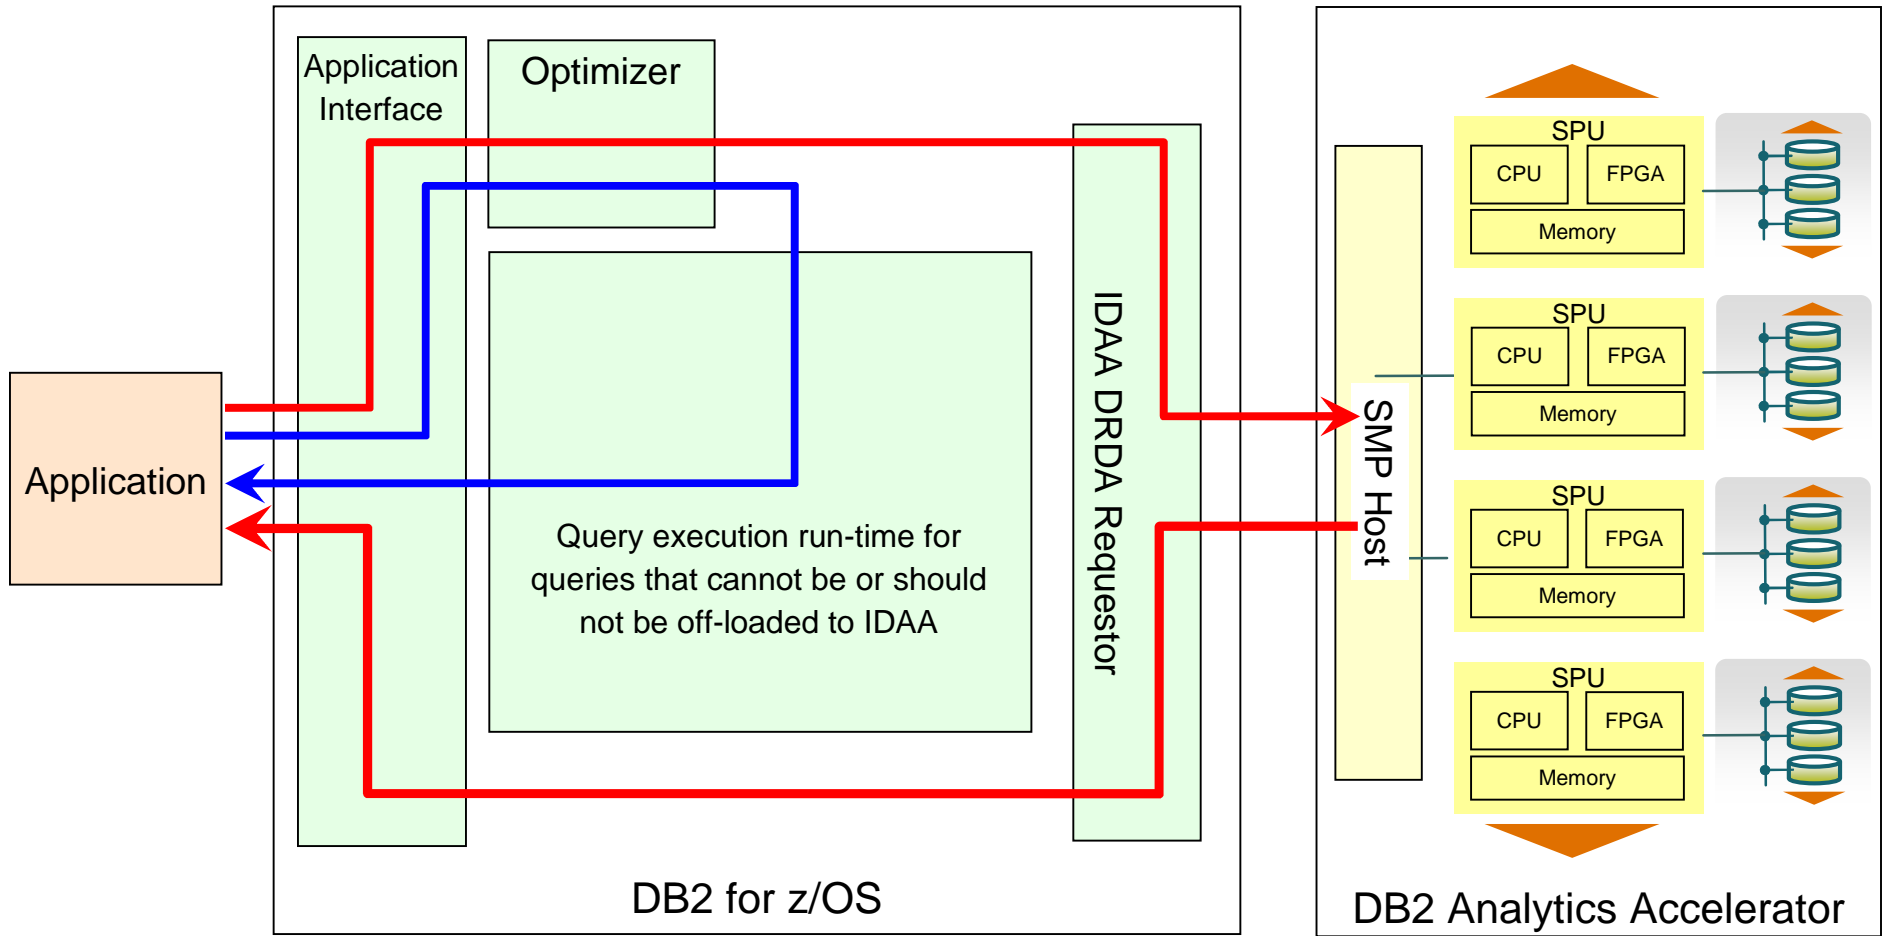


- Single and unique system for mixed query workloads
- Dynamic decision for most efficient execution platform
- New special register QUERY ACCELERATION
- New heuristic in DB2 optimizer

New

Value	Description
NONE	No query is routed to the accelerator
ENABLE	A query is routed to the accelerator if it satisfies the acceleration criteria including the cost and heuristics criteria. Otherwise it is executed in DB2. If there is an accelerator failure while running the query, or the accelerator returns an error, DB2 will return a negative SQL Code to the application
ENABLE WITH FAILBACK	A query is routed to the accelerator if it satisfies the acceleration criteria including the cost and heuristics criteria. Otherwise it is executed in DB2. Under certain conditions the query will run on DB2 after it fails in the accelerator. In particular any negative SQL code will cause failback to DB2 during PREPARE or first OPEN. No failback is possible after a successful OPEN of a query
ELIGIBLE	A query is routed to the accelerator if it satisfies the acceleration criteria irrespective of the cost and heuristics criteria. Otherwise it is executed in DB2
ALL	A query is routed to the accelerator, if it cannot execute the query fails and a negative return code is passed back to the application

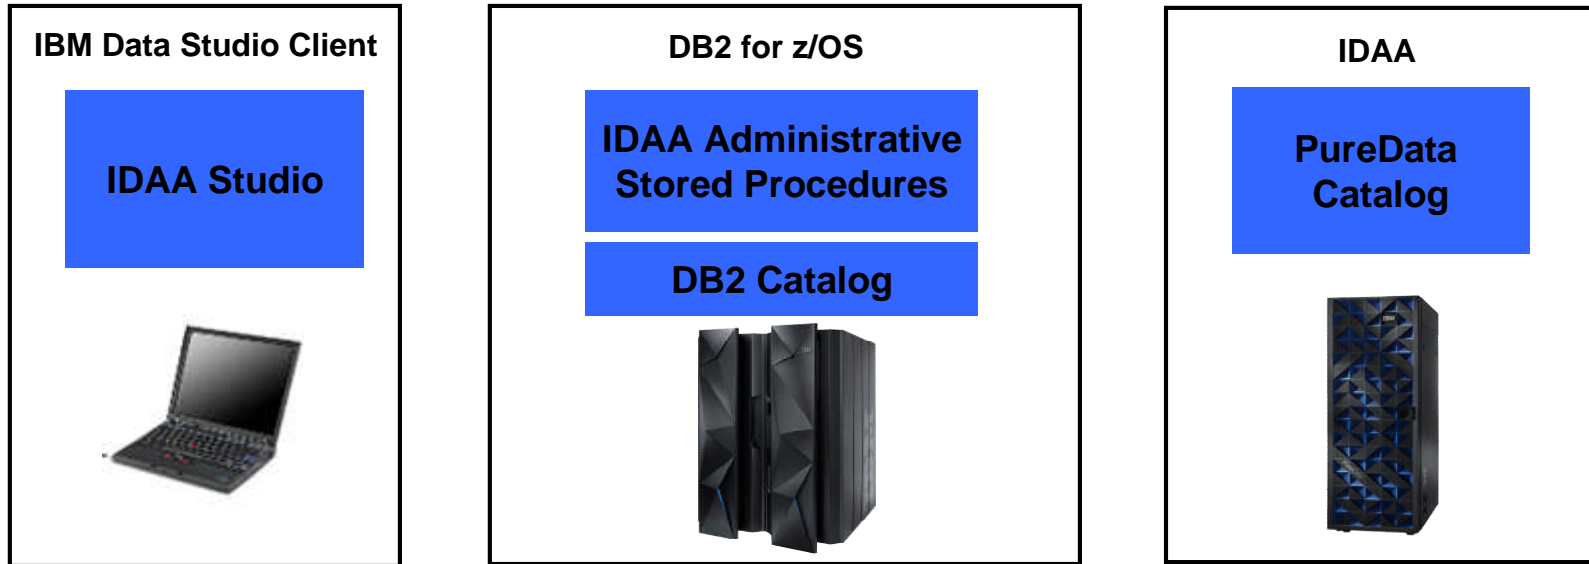
Query Execution Process Flow



Queries executed without DB2 Analytics Accelerator

Queries executed with DB2 Analytics Accelerator

Table Definition and Deployment

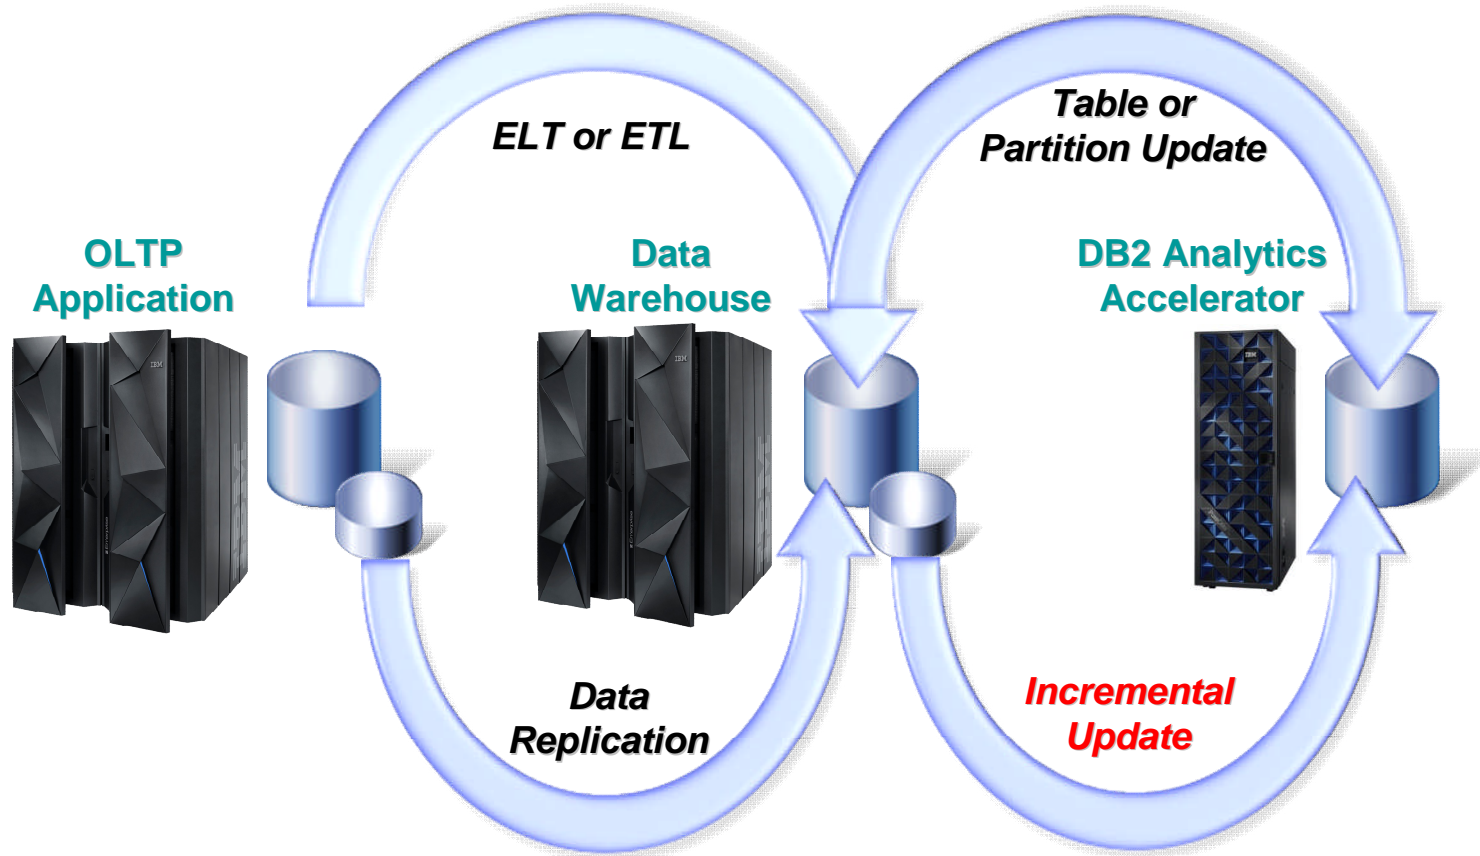


- The tables need to be defined and deployed to IDAA before data is loaded and queries sent to it for processing
 - Definition: identifying tables for which queries need to be accelerated
 - Deployment: making tables known to DB2, i.e. storing table metadata in the DB2 and PureData catalog
- IDAA Studio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

Synchronization Options with IBM 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>	<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>	<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 ▪ Changed partitions only through real-time statistics (*IDAA V3 and up, DB2 10 and up)
<p>Incremental Update (*IDAA V3 and up, DB2 10 and up)</p> <p>Log-based capturing of changes and propagation to IBM 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

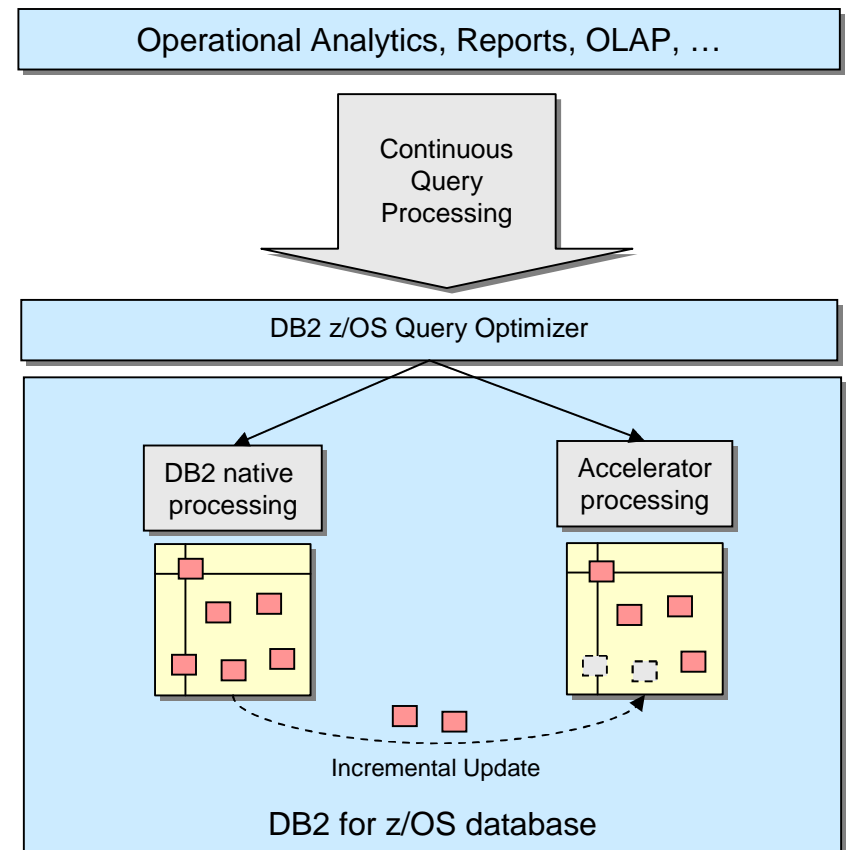
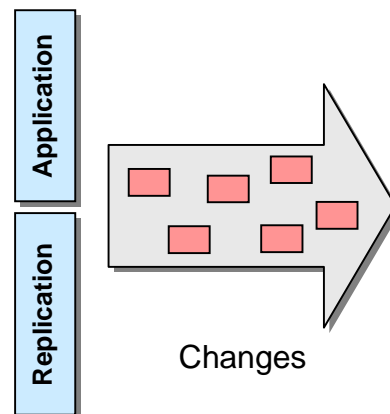
Introducing Incremental Update



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

Incremental Update – IDAA V3 for DB2 V10

- Changes in data warehouse tables typically driven by replication or an other update process
 - Corrections after a bulk-ETL-load of a data warehouse table
 - Continuously changing data (e.g. trickle-feed updates from a transactional system to the warehouse or ODS)
- Reporting and analysis based on most recent data
- May be combined with a full/partition table refresh
- Incremental update can be configured per table
- Initial full table load is required



EXPLAIN

- **DB2 EXPLAIN function is enhanced to provide information about accelerator usage**
 - Whether query qualifies for acceleration and reason why it is not qualified
 - The access path details associated with the query execution by IDAA are provided independently of DB2 EXPLAIN by the IDAA Studio.
- **For each query a row is inserted in the following tables:**
 - If Qualified for Acceleration**
 - PLAN_TABLE's ACESSTYPE column is set to a value of 'A'
 - DSN_QUERYINFO_TABLE's QI_DATA column shows the converted query text
 - If not Qualified for Acceleration**
 - DSN_QUERYINFO_TABLE - REASON_CODE and QI_DATA columns provide details

Note: The EXPLAIN tables can be populated with the above information even if there is no accelerator connected to DB2. Setting EXPLAINONLY on START ACCEL enables DB2 to consider the accelerators presence in the access path selection process.

DSN_QUERYINFO_TABLE



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

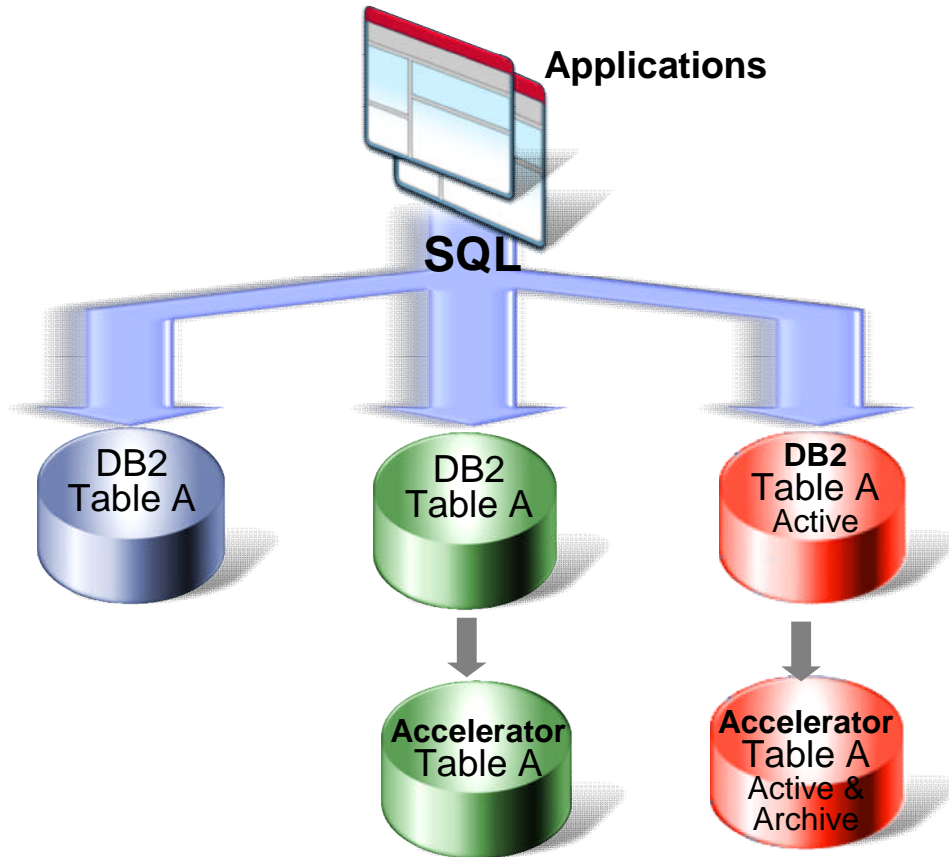
DSN_QUERYINFO_TABLE - Reason Codes

- 0 - The query can be accelerated.
- 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 offloading is not advantageous.
- 4 - The query is not read-only.
- 5 - The query is running under the private protocol (DB2 9 for z/OS only)
- 6 - The cursor is defined as scrollable or rowset cursor.
- 7 - The query references objects with multiple encoding schemes.
- 8 - The query FROM clause specifies a data-change-table-reference.
- 9 - The query contains a table expression with one or more correlated references to other tables in the same FROM clause.
- 10 - The query contains a reference to a recursive common table expression.
- 11 - The query contains an unsupported expression. The text of the expression is in QI_DATA.
- 12 - The query references a table that meets one of the following conditions:
 - The table is not defined in the accelerator server.
 - The table is defined in the accelerator server, but is not enabled for processing by an accelerator.
- 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.
- 15 - The query contains DB2 V10 new functionality. (DB2 9 for z/OS only)

High Performance Storage Saver

Reducing the cost of high speed storage

Store historic data on the Accelerator only



- **Best for OLTP**
- **High Speed Indexed queries**

- **Mixed Workload**

- **Active Only**
- **Archive Only**
- **Active & Archive**
- **Mixed Workload**

Tables can be resident on:

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

When data no longer requires updating, reclaim the DB2 storage

Managed by zPARMs

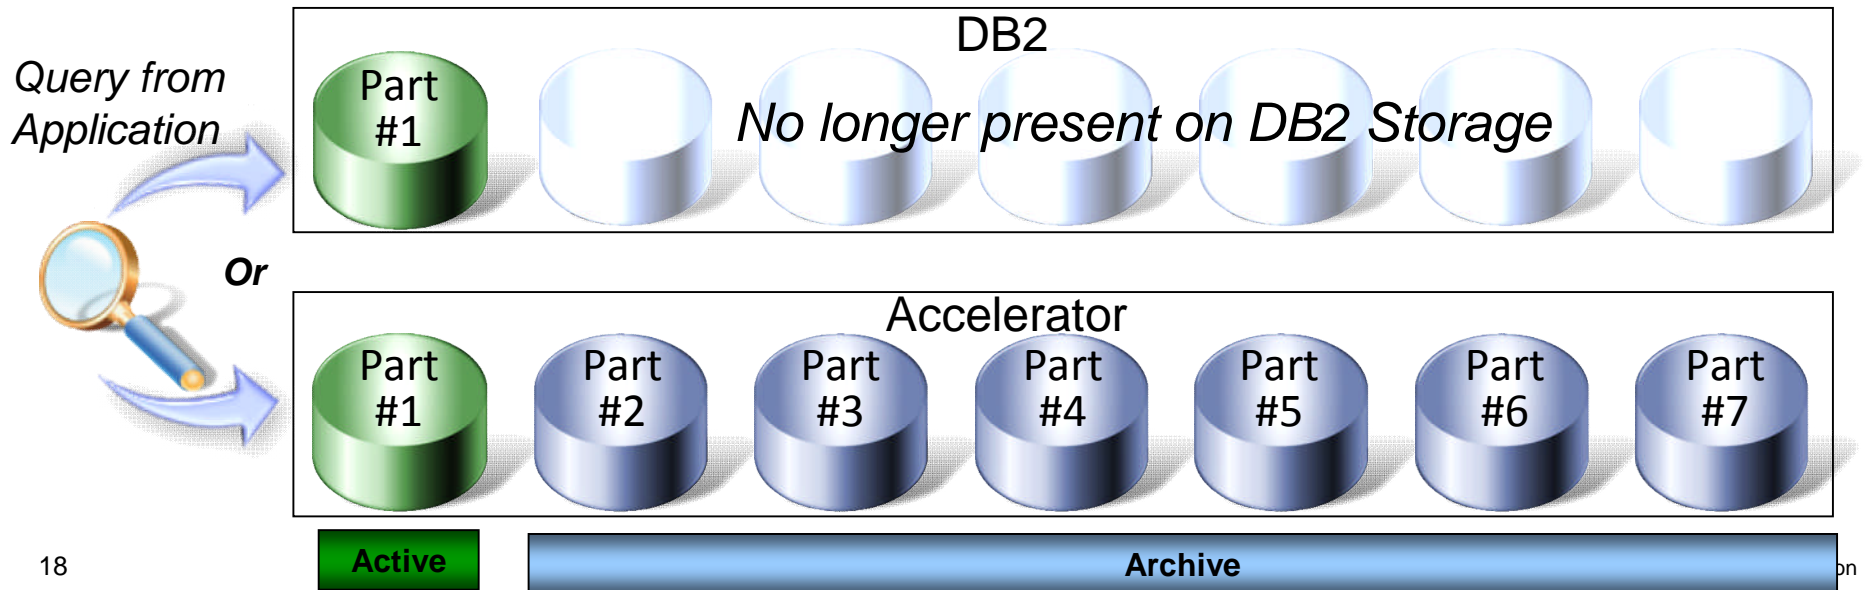
Controlled by Special Registers:

- CURRENT QUERY ACCELERATION
- CURRENT GET_ACCEL_ARCHIVE

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).
- **High Performance Storage Saver’s “Archive” Process:**
 - Data is loaded into Accelerator if not already loaded
 - Automatically takes Image Copy of Each Partition to be Archived
 - Automatically Remove data from DB2 archived tablespace partitions
 - DBA starts archived partitions as read-only



Work Load Management - Usage scenarios

- **Workload Isolation:**
Ensure that the workload of one DB2 subsystem doesn't monopolize the resources of a shared accelerator. A development subsystem, attached to the same accelerator as a production subsystem, should not be able to drain all accelerator resources.
- **Query Prioritization:**
More important queries should be executed before and faster than less important queries that are sent from the same DB2 subsystem against the accelerator.

Query Prioritization

- DB2 sends the importance level to IDAA with each query.
- IDAA maps the importance level to a Netezza priority and alters the session prior to query execution, using the corresponding priority. Also threads scheduled will have their priorities adjusted.
 - The changes in prioritization after query start are not reflected
- Netezza supports only 4 different priority levels, therefore multiple WLM importance levels have to be mapped against the same Netezza priority.

WLM Importance Level	Netezza Priority
System	Critical
Importance 1	Critical
Importance 2*	High
Importance 3*	Normal
Importance 4	Normal
Importance 5*	Normal
Discretionary	Low

Introducing the PureData System for Analytics N2001


Improved Performance, Energy Efficiency and Resiliency

Increased Performance & Capacity




- **Improved 3x faster performance¹**
- **Faster disk scan rate-128 GB/sec effective scan rate per rack**

Improved Energy Efficiency



- **50% greater capacity and faster performance, with no increase in floor space, power or cooling requirements²**
- **Uses less power than the competition**
- **Better capacity and power requirements than competition**

Improved Resiliency and Fault Tolerance



- **More spare drives per cabinet**
- **Faster disk regeneration due to smaller drives**
- **70% Fewer service calls³**

¹3x faster performance refers to the query times on both macro-analytic and mixed workload tests as conducted in IBM engineering lab benchmarks where the IBM PureData System for Analytics N2001 was shown to be an average of 3x faster than N1001. Individual results may vary.

² 50% greater capacity when compared to previous model PureData System for Analytics N1001. Power and cooling specifications within 97% of previous model PureData System for Analytics N1001.

³ Each N2001 rack contains 34 hot spare drives and 240 active drives for a ratio of 1 spare per 7 drives. Each N1001 rack contains 4 hot spare drives and 92 active drives for a ratio of 1 spare per 23 drives. The N2001 has 3.3x more spares per active drive. Frequency of disk related service calls expected to decrease by 70% assuming the same drive failure rates.

N2001 Hardware Overview – Full Rack

PureData System for Analytics N2001



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

N2001 Systems and Sizes

PureData System for Analytics N2001



	-	005	010	-	020	-	040			
Cabinets		1/2	1		2		4	...		
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

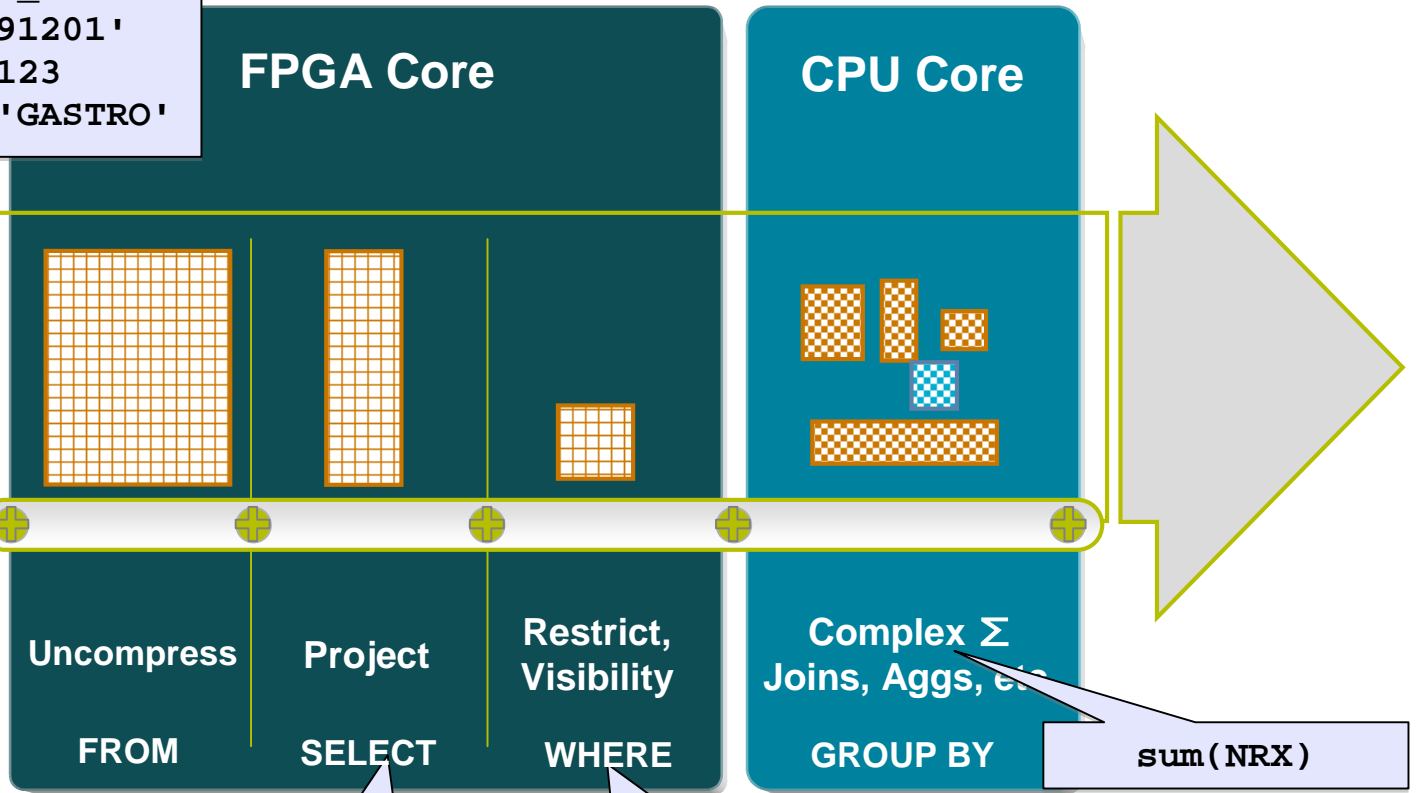
The Key to the Speed



```
select DISTRICT,
       PRODUCTGRP,
       sum(NRX)
from   MTHLY_RX_TERR_DATA
where  MONTH = '20091201'
and    MARKET = 509123
and    SPECIALTY = 'GASTRO'
```



Slice of table
MTHLY_RX_TERR_DATA
(compressed)



```
select DISTRICT,
       PRODUCTGRP,
       sum(NRX)
```

```
where MONTH = '20091201'
and    MARKET = 509123
and    SPECIALTY = 'GASTRO'
```


IDAA - Fast Time to Value

Reduce MIPS & CPU Consumption + Productivity Benefits + Delay Capacity Upgrades + Reduced ISV SW Costs = Fast Time To Value

Delay Capacity Upgrades

- DB2 workload that would benefit from IDAA would be offloaded
- Avoid storage costs
 - 1/3 less required Indexes of target tables
- Avoid ISV SW Costs
- Avoid Maintenance Costs

Productivity Benefits

- Reduced IT Operational Costs due to reduced complexity
- Over time, once initial IDAA environment is setup, DBA time freed up for additional projects
 - No indexes = less DBA time
 - Less database and storage administration
 - Less overall application tuning
 - Logical and physical design may not be required

Improved Service Levels on Key

Application and Reporting Environments

- Fast, predictable response times for “right-time” analysis
- Accelerated analytic query response times
- Improved price/performance for analytic workloads
- Highly secure environment for sensitive data analysis
- Transparent to the application

Customer Experience with IDAA visit:

http://www-01.ibm.com/software/success/cssdb.nsf/cs/JJAE-8SF863?OpenDocument&Site=swzseries&cty=en_us

No-Charge Workload Evaluation/Sizing

- ❑ Collect information from dynamic statement cache, supported by step-by-step instruction and REXX script (small effort for customer)
 - Optional masking available of sensitive data
- ❑ Upload compressed file (up to some MB) to IBM FTP server
- ❑ Analysis of file completed based on known Accelerator capabilities

IBM DB2 Analytics Accelerator -- Workload

Center of Excellence, Datawarehouse on System z, IBM Research & Development

Query Summary

	Total	DB2 natively	With potential	Uncertain	W/o potential
Queries	1856	781 (42%)	1072 (58%)	0 (0%)	3 (0%)
Elapsed Time (s) [1*]	48608.43	4324.99 (9%)	44191.64 (91%)	0.00 (0%)	91.80 (0%)
Elapsed Time (s)	60302.81	4522.30 (7%)	55688.70 (92%)	0.00 (0%)	91.81 (0%)
CPU Time (s) [1*]	5033.74	769.58 (15%)	4248.20 (84%)	0.00 (0%)	15.96 (0%)
CPU Time (s)	5782.50	772.47 (13%)	4994.06 (86%)	0.00 (0%)	15.97 (0%)

[1*] - Considers each query only once

Query classification	Unique Queries	Executed Queries	CPU %	CPU eligible %
Total	1856	9099	100%	86%
> 60 min elapsed time (info)	1	1	2%	2%
10-60 min elapsed time (info)	20	24	43%	37%
1-10 min elapsed time (info)	78	117	42%	37%
< 1 min elapsed time (info)	1757	8957	13%	11%


Reason breakdown for 0% queries with no potential and 0% of elapsed time with no potential

Reason	# Queries	% Queries	% Elapsed Time
Unsupported function	1	0%	0%
Unsupported encoding	1	0%	0%
Quant. predicates (temp. restr.)	1	0%	0%


Start trace time	End explain time	Min time stmt cached	Max time stmt cached
Mar 28, 2011 11:28 AM	Apr 4, 2011 10:11 AM	Mar 27, 2011 10:05 PM	Apr 4, 2011 10:06 AM

Report for a first assessment:

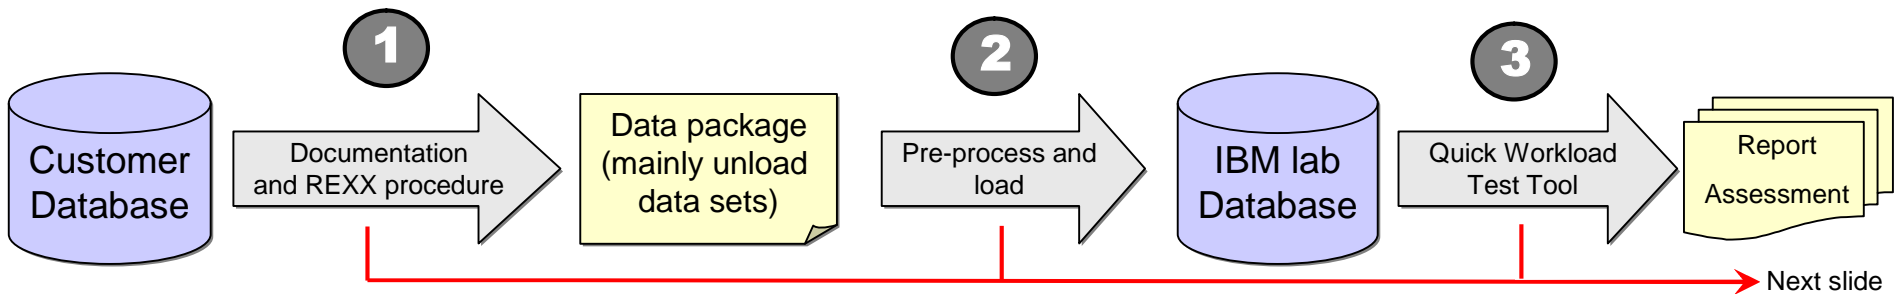
- Acceleration potential for
 - Queries
 - Estimated time
 - CP cost



Elapsed time for best DB2 native processing
 Elapsed time with acceleration potential
 Elapsed time with uncertain potential
 Elapsed time without acceleration potential



Query CPU time for best DB2 native processing
 Query CPU time with acceleration potential
 Query CPU time with uncertain potential
 Query CPU time without acceleration potential



IBM DB2 Analytics Accelerator for z/OS V4.1 at a Glance

More Query Acceleration	Enhanced Capabilities	Improved Transparency
<p>Static SQL</p> <p>DB2 Version 11 Support added</p> <p>Multi-row fetch from local applications</p> <p>EBCDIC & Unicode in same DB2 system & accelerator</p>	<p>Incremental Update</p> <ul style="list-style-type: none">• Greatly improved scalability• Better performance <p>Improved performance for large result sets</p> <p>HPSS</p> <ul style="list-style-type: none">• Archive to multiple accelerators• Better access control for archived partitions <p>Extend WLM support to local applications</p> <p>Richer system scope monitoring</p> <p>Report prospective CPU cost & Elapsed time savings</p> <p>Separation of duties for accelerator system administration operations</p>	<p>Automatic workload balancing over multiple accelerators</p> <p>New RTS 'last-changed-at' timestamp</p> <p>Automated NZKit installation</p> <p>HPSS</p> <ul style="list-style-type: none">• Built-in restore• Protection for image copies of archived partitions <p>Profile controlled special registers</p> <p>Improved continuous operations for Incremental Update</p>

Enabling new use cases

IBM DB2 Analytics Accelerator V4.1 - Highlights

More SQL Query constructs Eligible for Accelerator

▪ **Static SQL support**

- Static queries on active or archived data are now routed to the accelerator
- New bind options available (QUERYACCELERATION, GETACCELARCHIVE)
Same values for bind options available as for zparms QUERY_ACCELERATION and GET_ACCEL_ARCHIVE
- Acceleration for static queries is determined and fixed at bind package time
Tables must be added and enabled or archived on accelerator before binding the package
- Accelerator must be active and started when static query runs otherwise the execution fails

▪ **Rowset query offload and multi-row fetch support for local applications**

- CPU reduction for retrieving data from a local accelerated query returning a (very) large result set.
- For dynamic SQL in local applications:
 - Specify WITH ROWSET POSITIONING
 - To fetch use a FETCH NEXT ROWSET with FOR N ROWS clause

▪ **More DB2 functions and data types supported**

- BITAND(), TIMESTAMPDIFF()
- Enable comparison between different datatypes (e.g. VARCHAR and INTEGER)
- Retrofitted to IDAA V3.1 (PM88071)

IBM DB2 Analytics Accelerator V4.1 - Highlights

More Efficient use of the Accelerator Appliance

- **Efficient workload balancing for multiple accelerators**
 - If a query can be executed on more than one accelerator, DB2 will balance routing between qualifying accelerators considering accelerator utilization
- **Incremental update support of multiple DB2 subsystems**
 - Replicate from up to 10 different subsystems into a shared accelerator.
 - Based on provided IBM InfoSphere CDC 10.2.1
- **Support replication while tables are reloaded**
 - In IBM DB2 Analytics Accelerator V3.1 it was necessary to stop replication for all tables if one replication enabled table needs to be reloaded
 - Now reloading of a table can be done while replication of other tables continues

IBM DB2 Analytics Accelerator V4.1

Enhanced Accelerator Monitoring

Set of new and revised monitoring counters on system level exposed via IFI for better support of charge-back, capacity planning, monitoring and problem determination

System level statistic counters

- Exposed via DB2 statistics IFCID 2
- IBM Tivoli Omegamon Performance Expert can expose counts to users
- DISPLAY ACCEL(*) DETAIL command exposes some counts
- “Total” counters available in 2 flavors: for *this* DB2 system and for *all* DB2 systems

▪ Query Execution Statistics

- Total number of successful queries sent by all DB2 systems
- Total number of failed queries sent by all DB2 systems
- Number of currently executing queries
- Max number of concurrently executing queries

▪ CPU Costs

- Total CPU costs of executing queries sent by this or all DB2 systems
- Total CPU costs of maintenance operations from this or all DB2 systems (e.g. LOAD, ARCHIVE, RESTORE)

▪ Incremental Update Processing

- Total CPU cost associated to apply process
- Total number of log records read by capture agent
- Total number of log records processed by capture agent applicable to tables in the accelerator
- Total number of bytes processed by capture agent
- Total number of INSERTed/UPDATEd/DELETEd rows processed by capture agent applicable to tables in the accelerator

▪ System State

- Current replication latency
- Current replication state
- Timestamp of last status change of replication state

IBM DB2 Analytics Accelerator V4.1

More Efficient use of the Accelerator Appliance

WLM Support for local Applications

- Extend the already available WLM query prioritization support for remote queries now to local queries
- DB2 now detects the service class and importance level for local applications and send the information to the Accelerator via a special register prior to PREPARE
 - Service class is detected based on application's address space
 - ❑ TSO subsystem type – e.g. SPUFI
 - ❑ JES subsystem type – e.g. TEP3, host language apps run via TSOBATCH
 - ❑ STC subsystem type – any apps kicked off via START
 - For IMS/CICS applications the service class is detected based on a passed performance block
- Maps to Accelerator priority (same priorities as remote applications)

WLM Importance Level	Accelerator Priority
System	Critical
Importance 1	Critical
Importance 2	Critical
Importance 3	High
Importance 4	Normal
Importance 5	Low
Discretionary	Low

Note: Changes compared to V3

- Importance level **2** mapped to Accelerator Priority **High**
- Importance level **3,4,5** mapped to **Normal**.

IBM DB2 Analytics Accelerator V4.1 - Highlights

Enhanced Storage Saver Solution (HPSS)

▪ Archiving enhancements

- DB2 partitions are set to persistent read-only (PRO) status (after the partition was moved to the accelerator to prevent future INSERTs or UPDATEs to the partitions)
- Created image copies are protected by the PRO status of the DB2 partitions
 - No further image copies can be created
- Up to 4 image copies per partition are possible
 - Improves disaster recovery scenarios
- A flexible naming schema for the image copies can be defined
 - Simplifies the restart process or rearchiving since the image copies must no longer be deleted manually beforehand
- Ability to archive a table on multiple accelerators

▪ Restore archived partitions

- New stored procedure to automate the process to change one or more partitions from accelerated archived partitions to accelerated partitions (stored in DB2 and on the accelerator)
 - Restore is done from image copy data
 - Useful to update the data and archive again

IBM DB2 Analytics Accelerator V4.1 - Highlights

Easier Operation and Maintenance

- **Automated NZKit Install**

- The most common Netezza software update (NPS) can now be done through a stored procedure or Data Studio in the same way as accelerator software updates are done already

- **Fine grained access control for stored procedure
ACCEL_CONTROL_ACCELERATOR**

- SP offers several function to control an accelerator, e.g. Cancel task, starting replication, collecting trace
- For each of these functions now access can be controlled separately to control who can execute it

IBM DB2 Analytics Accelerator V4.1

Miscellaneous Topics

- **Multiple codepage support**
 - Mixed EBCDIC and UNICODE tables are now allowed on the accelerator for the same DB2 subsystem
 - Queries that combine both EBCDIC and UNICODE tables can still not be routed
- **Provide the IBM DB2 Analytic Accelerator Studio as a full install image based on Data Studio V4.1**
- **Subcapacity licensing – allows growth on demand**
 - Buy a full rack N2001 but pay only for 50 %, 62.5%, 75 % or 87.5 %
 - Licensee must not use more than defined capacity
 - Customer responsibility to monitor and document the actual usage
 - System provides interfaces to set/monitor resource limits

Thank You