

The IBM DB2 Analytics Accelerator

Version 3.1





Carlos Guardia
Executive IT/Specialist
IBM Software Group

© 2013 IBM Corporation
Information Management

DB2 Analytics Accelerator V3

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

2

DB2 Analytics Accelerator

Train-of-thought 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

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

DB2 Analytics Accelerator V3

Raising the Bar in Analytics



New System: "Striper"

Large analytic systems at dramatically lower costs

Business changing features

- 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 Savers
- Incremental Update
- zEnterprise EC12 Support
- Query Prioritization
- High Capacity
- UNLOAD Lite

Current situation around multi-temperature data

- Most of the data in an ODS or EDW is static
 - The large tables are partitioned by time
 - Older partitions are never changed, only the most recent partition is frequently changed
- Many DBMS vendors provide multi-temperature data solutions
 - Multi-temperature concept is not the same as archiving, but they have lots in common
 - Level of sophistication in implementing multi-temperature varies
 - the industry leading solutions are so called 'near-line storage servers' ('near-line' means 'near-online')
 - Multi-temperature's value proposition:
 - Move less frequently accessed data to cheaper storage
 - Improve performance for both queries and administrative operations accessing more recent data
- The drawback is degraded performance of queries that access old data

© 2013 IBM Corporation

High Performance Storage Saver

Reducing the cost of high speed storage

Store historic data on the Accelerator only

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

Applications → SQL → DB2 Table A, Accelerator Table A, Accelerator Table A

High speed indexed lookups, best for OLTP type queries (DB2 Table A)
High speed aggregate lookups, best for complex DSS type queries (Accelerator Table A)
Mixed workload type queries (Accelerator Table A)

Tables can be resident on:

- DB2 Only
- DB2 and Accelerator
- Accelerator Only

When data no longer requires updating, reclaim the DB2 storage

Special Registers control behavior

- CURRENT QUERY ACCELERATION
- CURRENT GET_ACCEL_ARCHIVE

Managed by zParsms

© 2013 IBM Corporation

Supplied Stored Procedure encapsulates Archiving Procedure

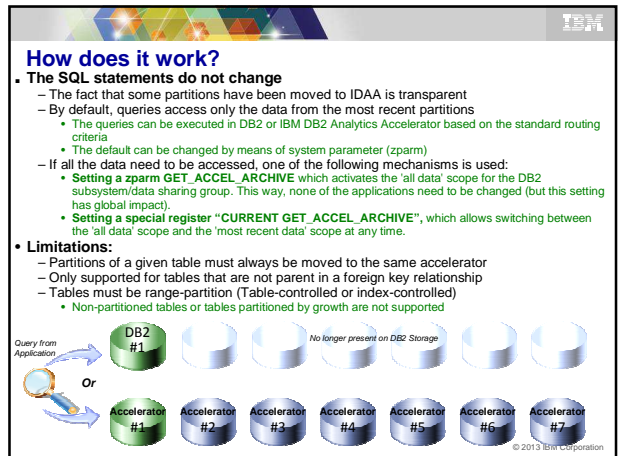
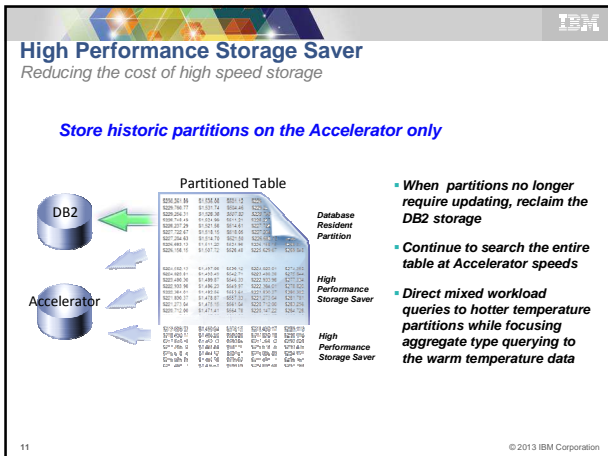
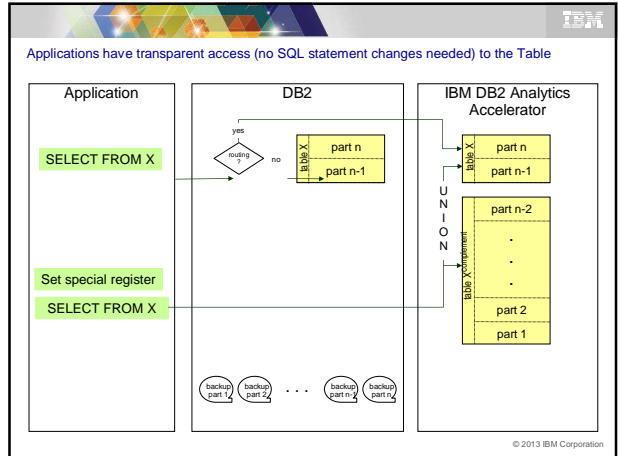
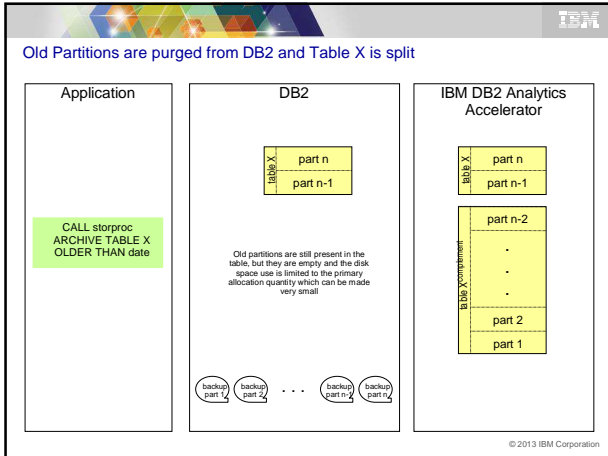
<p>Application</p> <p>CALL storproc ARCHIVE TABLE X OLDER THAN date⁽¹⁾</p> <p>(1)'date' is specified in terms of the partitioning key values</p> <p>in this particular example 'date' indicates that only last two partitions should remain in DB2</p>	<p>DB2</p> <p>part n part n-1 part n-2 . . part 2 part 1</p> <p>backup part 1, backup part 2, ..., backup part n-2, backup part n</p>	<p>IBM DB2 Analytics Accelerator</p> <p>part n part n-1 part n-2 . . part 2 part 1</p>
---	---	--

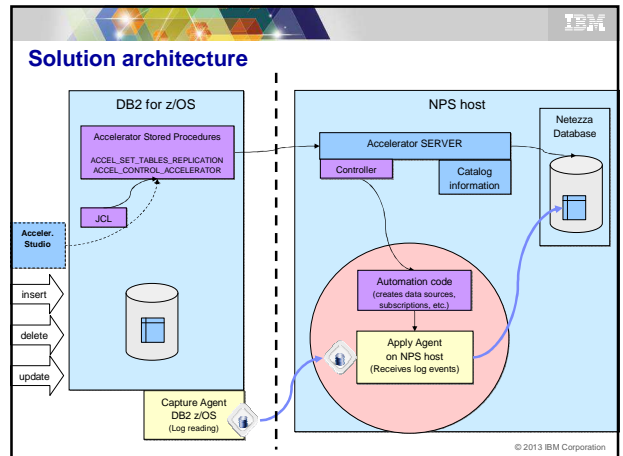
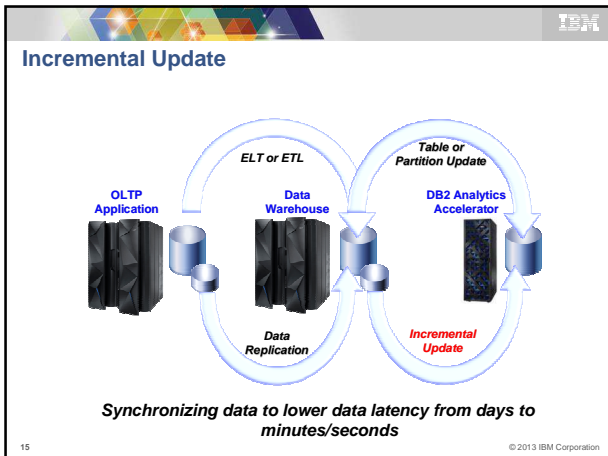
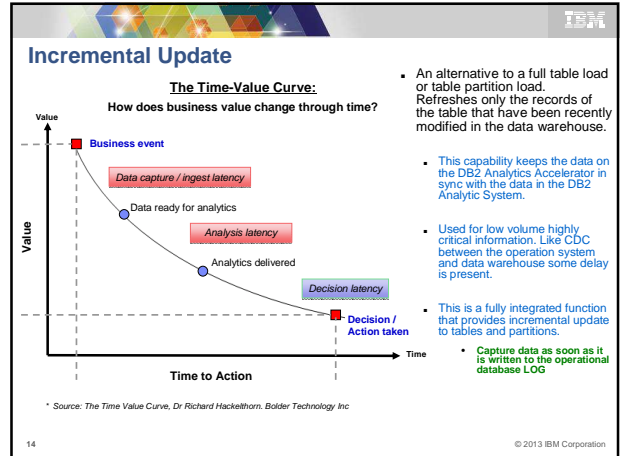
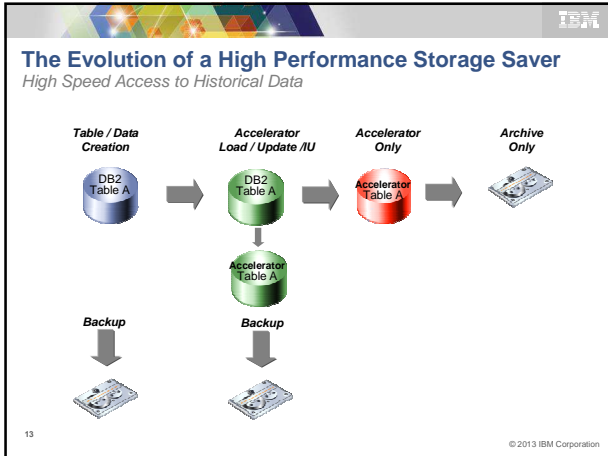
© 2013 IBM Corporation

Partitions to be Archived are first backed up

<p>Application</p> <p>CALL storproc ARCHIVE TABLE X OLDER THAN date</p> <p>in this particular example 'date' indicates that only last two partitions should remain in DB2</p>	<p>DB2</p> <p>part n part n-1 part n-2 . . part 2 part 1</p> <p>backup part 1, backup part 2, ..., backup part n-2, backup part n</p>	<p>IBM DB2 Analytics Accelerator</p> <p>part n part n-1 part n-2 . . part 2 part 1</p>
---	---	--

© 2013 IBM Corporation





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 from propagation of incremental changes
- Queries may continue during full table refresh for accelerator
- Full table refresh may be triggered through DB2 stored procedure (scheduled, integrated into ETL process or through GUI)

17 © 2013 IBM Corporation

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
- Maintains snapshot semantics for consistent reports
- Queries may continue during table partition refresh for accelerator
- Table partition refresh may be triggered through DB2 stored procedure (scheduled, integrated into ETL process or through GUI)

18

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

19

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

Name	Size	Rows	Acceleration	Last Load	Replication Since	Distribution Key	skew	C
TABLE_001	-	-	Enabled	3/15/12 11:26 AM	1/17/12 5:41 PM	Random	0.000	

Detect data "staleness"

Goal: Make data maintenance on the accelerator easier and reduce resource consumption

Situation in IBM DB2 Analytics Accelerator V2:

- Accelerator data can be refreshed from DB2 at table or partition granularity
- Admin must know explicitly which data needs to be refreshed (e.g. from ETL process)
- New feature: find out if DB2 data for a table or partition has changed since last load on accelerator**
- Avoid refreshing data that is already up-to-date without requiring explicit knowledge about data changes
- Based on DB2 real-time statistics about data updates*

© 2013 IBM Corporation

Interactive usage scenario

In the IBM DB2 Analytics Accelerator admin GUI:

- When re-loading tables, show which tables or partitions need to be reloaded and how much data needs to be loaded
- Allow to manually correct "in-doubt" cases where no automated decision is possible
- Implemented as a new load dialog with additional features
- There is also a new stored procedure behind the dialog that can be called directly to retrieve the information

From JCL or DB2 client

- Refresh one or more tables by loading only the smallest necessary amount of data (partitions or unpartitioned tables)
- If a table has not changed since last load, no data is transferred
- In-doubt cases will trigger a reload -- "false positives" are possible

© 2013 IBM Corporation

Now expandable to 960 cores and 1.28 petabytes

	002	005	010	015	020	030	040	060	080	100
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

PureData System for Analytics

Predictable, Linear Scalability throughout entire family

Capacity = User Data space
Effective Capacity = User Data Space with compression
* 4x compression assumed

Low Latency, High Capacity Update

© 2013 IBM Corporation

Query Prioritization

Extending the value of System z workload management

Various priority end-user applications

© 2013 IBM Corporation

Query Prioritization


- DB2 sends the importance level to IBM DB2 Analytics Accelerator with each query.
- IBM DB2 Analytics Accelerator 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

© 2013 IBM Corporation

UNLOAD Lite

Continuing to reduce overhead costs



Capitalizing on low cost MIPS on the Accelerator


26

© 2013 IBM Corporation

Additional enhancements

One table enabled in multiple accelerators

- Before
 - One table can only be enabled in one accelerator
- After
 - One table can be enabled in multiple accelerators
 - An accelerator is qualified for a query if it contains all the referenced tables that are enabled.
 - The first qualified accelerator is picked by DB2 (implementation-dependent)
- Use cases
 - Capacity
 - Workload balancing between multiple accelerators
 - High availability/disaster recovery



New Control for QUERY ACCELERATION

- CURRENT QUERY ACCELERATION:
 - NONE
 - ENABLE
 - ENABLE WITH FAILBACK
 - ELIGIBLE: skip heuristics & costing in the ENABLE option. Queries that are not eligible for acceleration are executed by DB2**
 - ALL: All queries to be executed in the Accelerator, error (-4742) for not eligible queries**

© 2013 IBM Corporation

Prerequisites

Prerequisites

- DB2 Analytics Accelerator V2 (To be sold concurrently with V3)
 - DB2 for z/OS V9 or V10
 - zEnterprise EC12, z196 or z114
- DB2 Analytics Accelerator V3
 - DB2 for z/OS V10 compatibility mode
 - DB2 for z/OS V10 new feature mode
 - zEnterprise EC12, z196, or z114

Pricing

- Customer can elect to use either V2 or V3 at no difference in price

28

© 2013 IBM Corporation

New model: PureData System for Analytics N2001

IDAA has now TWO models

- N1001 ("TwinFin") – economical, high performance and scalability
- N2001 ("Striper") – highest performance appliance to-date
 - Announced 2013, February the 5th

Accelerate Performance of Analytic Queries



- 3X faster performance¹ for Big Data analytics
- 128 GB/sec effective scan rate per rack² to tackle Big Data faster

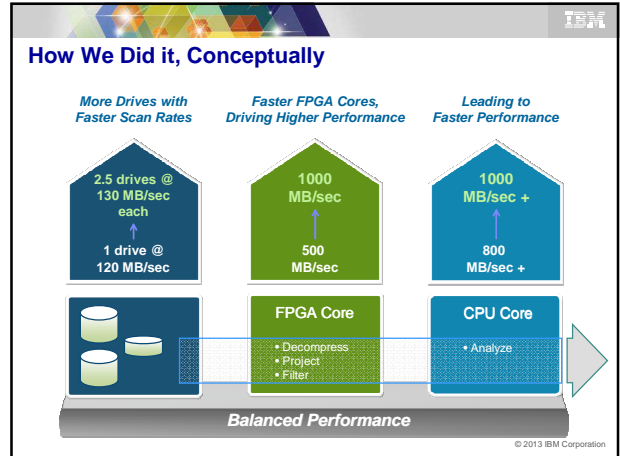
Increase Efficiency of your Data Center

- 50% greater data capacity per rack³ helps optimize data center efficiency
- More capacity and less power per rack than both Oracle and Teradata

Simplicity and Ease of Administration

- Improved system management and resilience to spend less time managing and more time delivering value
- 70% FEWER service calls with more spare drives and faster disk regeneration⁴



IDAA - PureData System for Analytics Models

	IDAA N1001	IDAA N2001
Blade Type	HS22	HX5
CPU Cores / Blade	2 x 4 Core Intel CPUs	2 x 8 Core Intel CPUs
# Disks	96 x 3.5" / 1 TB SAS (92 Active)	288 x 2.5" / 600GB SAS2 (240 Active)
Raw Capacity	96 TB	172.8 TB
Total Disk Bandwidth	~11 GB/s	~32 GB/s
S-Blades per Rack (cores)	14 (112)	7 (112)
S-Blade Memory	24 GB	128 GB
Rack Configurations	¼, ½, 1, 1 ½, 2 – 10	½, 1, 2, 4 (6 and 8 rack configs to follow)
FPGA Cores / Blade	8 (2 x 4 Engine Xilinx FPGA)	16 (2 x 8 Engine Xilinx Virtex 6 FPGA)
User Data / Rack *	128 TB	192 TB

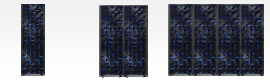
* Assuming 4x Compression

© 2013 IBM Corporation

N2001 Systems and Sizes

PureData System for Analytics N2001

New software NPST OS! Available for N1001 & N2001



	005	010	020	040	
Cabinets	1/2	1	2	4	Watch this space
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

© 2013 IBM Corporation



धन्यवाद
Hindi

多謝
Traditional Chinese

ขอบคุณ
Thai

多谢
Simplified Chinese

Спасибо
Russian

Thank You
English

Gracias!
Spanish

شكراً
Arabic

Obrigado
Brazilian Portuguese

Bedankt
Dutch

Merci
French

Danke
German

நன்றி
Tamil

ありがとうございました
Japanese

감사합니다
Korean

Carlos Guardia
cguardia@es.ibm.com

© 2013 IBM Corporation