

Have you heard what's new with the Db2 Analytics Accelerator?

Karen Wilkins MBCS IBM UK Ltd

November 2018

Session II





Agenda

- Value Proposition
- Version 7.1 Architecture and Deployment Options
- Version 7.1 Functionality
- Version 7.1.2 and Version 7.1.3
- Version 7.1.4

-



Value Proposition



IBM Z Analytics

Keep your data in place – a different approach to enterprise analytics



- Keep data in place for analytics
 - Keep data in place, encrypted and secure
 - Minimize latency, cost and complexity of data movement
 - Transform data on platform
 - Improve data quality and governance
- Apply the same resiliency to analytics as your operational applications
- Combine insight from structured & unstructured data from z and non-z data sources
- Leverage existing people, processes and infrastructure



IBM Db2 Analytics Accelerator for z/OS

A workload optimized, appliance add-on to Db2 for z/OS that enables the integration of analytic insights into operational processes to drive business critical analytics & exceptional business value.

SPEED

- Dramatically improve query response up to 2000X faster to support time-sensitive decisions
- Right-time. Low latency. Trusted. Accurate.

SAVINGS

- Minimize data proliferation
- Lower the cost of storing and managing historical data
- Free up compute resources

SIMPLICITY

- Simplify infrastructure, reduce ETL and data movement off-platform
- Non-disruptive installation

SECURITY

Safeguard valuable data under the control and security of Db2 for z/OS
Protected. Secured. Governed.





Hybrid transaction/Analytical processing



Db2 Analytics Accelerator and Db2 for z/OS

A self-managing, hybrid workload-optimized database management system that runs every query workload in the most efficient way, so that each query is executed in its optimal environment for the greatest performance and cost efficiency



Db2 Analytics Accelerator Version 7.1

Delivering new flexible, integrated deployment options

High-speed analysis of your enterprise data for real-time insight under the control and security of IBM Z

- Introduces new flexible, integrated deployment options
 - Accelerator on IBM Integrated Analytics System
 - Fast, simple deployment on pre-configured hardware and software
 - Flexible and elastic data storage
 - Accelerator on IBM Z
 - Unified homogeneity of service, support and operations
 - Flexible Capacity
- Based on IBM's premier analytical engine, Db2 Warehouse software
- Transition easily between deployment options
 - One API
 - One database engine





Db2 Analytics Accelerator Version 7.1

Deployment Options



- Db2 Analytics Accelerator Version 7.1
- Two deployment options



- On an on-premises appliance
- On a SW appliance installed on the z14 mainframe
- Both new options will offer
 - the same functionality
 - the same API
 - the same implementation
- This provides:
 - Coexistence and combination of deployment options, fully transparent for Db2 applications
 - Flexibility in moving data for query acceleration as workload demands grow or change
 - Consistency and efficiency in managing different Db2 Analytics Accelerator environments



Db2 Analytics Accelerator Version 7.1

Long and short names





Version 7.1 Architecture and Deployment Options



The core of V7.1 Accelerator Deployments

	Docker container		docker
	Accelerator server	Database engine	Additional future functionality
	Systems Manager	Workload Monitoring	Authentication
Infrastructure Management	Linux	M	M
	Physical server (Power) CPU	Memory
	Storage (SAN) (Clustered) Filesystem		

Docker is a widely adopted, open-source project that automates the deployment of applications inside software containers.

It allows to bundle and preinstall components in a Docker image and then to launch the container from the image.

The accelerator Docker container includes

- the accelerator server that establishes the connection between the Db2 for z/OS subsystem and manages all accelerator tasks
- a database engine
- other components required for high availability, monitoring and security
- a Linux operating system

Outside of the container physical compute resources are needed

- a server that includes multi-core CPUs
- large memory
- shared filesystem to persist the data

On top of the physical hardware there is a Docker supported Linux operating system that is used to launch the Docker container and manage the HW resources.



Powered by Db2 with BLU Acceleration



- Huge potential for faster ingest for incremental updates, and thereby less HTAP query delay!
- IBM's premier analytics engine across many products
- Latest analytics technology innovations
- SQL compatibility across all IBM products
- High degree of concurrent users and queries

The new engine is replaced internally – external interfaces will stay the same. The same Db2 subsystem can be connected to the existing and new generation.



Deployment on IBM Integrated Analytics System (IIAS)





- New generation hardware appliance
- A full solution that provides all components out of the box including optimized hardware and software
- All components provided by IBM in a balanced, performance-optimized configuration
 - HW, which includes the rack, the physical servers and the storage
 - SW stack including the Linux operating system, the Docker software as well as the Docker container and the infrastructure management
- IBM Power hardware for the appliance, balanced and optimized for price/performance



IBM Db2 Analytics Accelerator on IIAS

Product components





IIAS Configurations



IBM Power 8 S822L 24 core

server 3.02GHz

	M4001-003 1/3 Rack	M4001-006 2/3 Rack	M4001-010 Full Rack	-	
Servers	3	5	7		
Cores	72	120	168		
Memory	1.5 TB	2.5 TB	3.5 TB		
Flash Storage Capacity	16 TB	32 TB	48 TB		
	M4002-003 1/3 Rack	M4002-006 2/3 Rack	M4002-010 Full Rack	M4002-020 Double Rack	M4002-020 Quad Rack
Servers	3	5	7	13	25
Cores	72	120	168	312	600
Memory	1.5 TB	2.5 TB	3.5 TB	6.5 TB	12.5 TB
Flash Storage Capacity	27 TB	54 TB	81 TB	162 TB	324 TB



Hardware Architecture Overview

Full Rack M4001-010

7 Compute Nodes in 1 rack containing

- IBM Power 8 S822L 24 core server 3.02GHz
- 512 GB of RAM (each node)
- 2x 600GB SAS HDD
- Red Hat[®] Linux OS

Up to 3 Flash Arrays in 1 rack containing

- IBM FlashSystem 900
- Dual Flash controllers
- Micro Latency Flash modules
- 2-Dimensional RAID5 and hot swappable spares for high availability

2x Mellanox 10G Ethernet switches

- 48x10G ports
- 12x40/50G ports
- Dual switches form resilient network

IBM SAN64B 32G Fibre Channel SAN

- 16Gb FC Switch
- 48x 32Gb/s SFP+ ports





Deployment on IBM Z

Db2 Analytics Accelerator on IBM Z



- An appliance running on IBM Z
- Packages the SW stack into an IBM Secure Service Container to deliver a fully self-managed appliance running in a SSC LPAR that can be deployed in minutes
- Integrates seamlessly into the customer's Z environment and leverages known LPAR-, memory and CPU management
 procedures, including call home support for enterprise hardware components.
- Uses customer-provided storage to hold the accelerator-side data

Accelerator Image Deployment into a Secure Service Container (SSC) LPAR

Customize Image Profiles: M135:LP15 : LP15 : General				
 	Profile name: Description: Partition identifier: Mode: • Clock Type Assignment • Standard time of day • Logical partition time of Ensure that the image p	LP15 Yoda Deploytest 15 General ESA/390 TPF Coupling facility LINUX only z/VM SSC	ns to the curr	Assigned for activation



SSC Partition & Deployment Service

- New SSC partition mode for SW appliances
- SSC partition is not for general purpose usage, only appliances can be deployed
- Speeds up deployment with an integrated installer
- Ensures integrity of downloaded appliance

SSC Appliance Runtime

- Operating system incl. Docker runtime Network and storage control incl. encryption
- System monitoring and FFDC

SSC Appliance Management Services

- Admin/User controls
 - View messages, events
 - Network management
 - User management
 - Disk management
 - View Appliance status
 - Create/obtain dumps/logs
 - Apply service, updates
- All management interactions with an appliance are via Web UI / **REST APIS**
- Software appliances based on SSC technology include zAware, zVSE, Blockchain, ...

GUIDE SHARE

FUROPE UK REGION

Hardware Considerations



- 1 Accelerator = 1 LPAR but multiple Accelerator LPARs on a single system possible
- Each LPAR requires IFLs, RAM, and storage

IFLs	Dedicated z14 Drawer	Storage
Use existing IFLs and memory For small production test/dev or getting- started use cases Minimum suggested configuration: • 4 IFLs, 256 GB memory for test/dev • 8 IFLs, 512 GB memory for production	 Order one drawer, comprised of 35 IFLs and up to 2.56 TB memory, with your z14 order You can convert your z14 models M01, M02, M03 servers by adding a dedicated drawer containing IFLs and memory at a very attractive price Also available on M04, M05 as a new build inclusive of the drawer, as you cannot add another drawer to M04 or M05 machines 	 Customer provided storage (actual size depends on workload) FCP or FICON attached Flash storage is not required, but e.g. IBM FlashSystem 900 is an effective option



IBM Db2 Analytics Accelerator on IBM Z

Product components





Network Connection to Db2 for z/OS



OSA 10 Gbps recommended Shared OSA is possible

A defined TCP/IP network interface from one Accelerator to Db2 subsystem(s)

- HiperSocket (CEC-internal) or
- OSA (internal or external)

M:N connectivity

- Multiple Db2 subsystems can share an Accelerator
- Multiple Accelerators can connect to a Db2 subsystem



Deployment Considerations

Accelerator Instances

- Db2 Analytics Accelerator on IBM Z instances can be deployed quickly (download & go)
 - Instances are currently limited to CPU resources in one physical drawer
- Higher instance granularity is possible than on IIAS based systems
 - Multiple independent applications / Db2 subsystems may share a bigger IIAS system while Accelerator on Z characteristics may lead to more, dedicated and smaller instances
 - Beneficial for (many) smaller test/dev instances which could be temporarily (re-) provisioned





One API – One implementation – Two deployment options



- Uniform experience, simultaneous use, and easy transition between different implementations
 - Common accelerator engine across all the platforms:

Db2 Warehouse



z based Appliance

POWER hardware and storage integrated in a self-contained workload-optimized system for analytics	Software appliance deployed on customer's z hardware and storage infrastructure
Out-of-the-box experience	Download & go experience
POWER-grade QoS	Z-grade QoS
Scalability by scale-out	Scalability by using Z: expansion of IFLs and memory
HA support within appliance	CA/DR support based on GDPS (active/passive and active/active)
Analytics for largest data volumes and highest performance	Flexible, smaller, elastic deployment option
Integrated storage and management	integration into existing z environment: hardware and storage management, CA/DR infrastructure, support processes, organizational structures, – no new infrastructure needed 24



Connectivity options



Full flexibility for Db2 systems and deployment options:

- 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

Represents all deployment options

Db2 Analytics Accelerator

for z/OS



Version 7.1 Functionality

Functionality overview as of V7.1.2



Db2 Analytics Accelerator Version 7.1	Db2 Analytics Accelerator Version 5.1			
Query acceleration for	dynamic and static SQL			
More query acceleration through more SQL native support				
Full/Partition tabl	e load and refresh			
Incremental Update				
Db2 Analytics Accelerator Loader Support (separate licensed tool)				
In-database transformation using Accelerator-only tables (AOTs)				
~ ·	High Performance Storage Saver (HPSS)			
~ ·	In-database analytics			
+	Federation			
*	- "True HTAP" (as Technical Preview)			
*	- Add column support			

The blue arrow denotes that IBM intends to deliver these functions in Version 7.1 in 2018/2019.



Query execution process flow





Routing criteria



- Dynamic and static queries can be accelerated
- Db2 Optimizer decides if query should be sent to Accelerator
 - Dynamic: At execution time
 - Static: At BIND time
- Whole query, <u>not</u> parts of query are accelerated
- Only read queries are considered for acceleration
- Queries within INSERT statements can be accelerated

- Prerequisites for query routing:
 - Accelerator is started
 - All used tables are available on Accelerator
 - Query routing option is specified
 - Via special register, BIND option or ZPARM



Query routing options

Dynamic

Values for special register CURRENT QUERY ACCELERATION

• Needs to be set prior to query execution to enable or suppress routing of queries

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 SQLCODE will cause a 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 satisfies the acceleration criteria irrespective of the cost and heuristics criteria. Otherwise the query fails and a negative return code is passed back to the application.



Query routing options

Static

Values for QUERYACCELERATION bind option

• Needs to be set at bind time to either bind a query for acceleration or for execution in DB2

а.
a.
QLCODE uery.
iteria.
e of the
; i



Comparing Accelerator V7 with V5

Data Types

Data Type support	Accelerator V7 (on IIAS or Z)	Accelerator V5 (on PDA)
EBCDIC MBCS, GRAPHIC	Supported natively	Converted to UTF-8
TIMESTAMP(12)	Supported natively	Truncated to precision 6
FOR BIT DATA subtype	Supported natively for EBCDIC, UNICODE, ASCII	Supported for EBCDIC only
DECFLOAT, BINARY, ROWID	With V7.1.2 for Db2 11	Not supported

V5: <u>https://www.ibm.com/support/knowledgecenter/en/SS4LQ8_5.1.0/com.ibm.datatools.aqt.doc/gui/references/r_idaa_supported_data_types.html</u> V7: <u>https://www.ibm.com/support/knowledgecenter/en/SS4LQ8_7.1.0/com.ibm.datatools.aqt.doc/gui/references/r_idaa_supported_data_types.html</u>



Comparing Accelerator V7 with V5

SQL Coverage

SQL support	Accelerator V7 (on IIAS or Z)	Accelerator V5 (on PDA)
correlated subqueries	All types supported including table expressions with sideway references	Only a small subset supported
Recursive SQL	Supported	Not supported
TIMESTAMP value 24:00:00	Supported natively	Mapped to 23:59:59.999999
Scalar functions	Improved support	Some not supported when using specific datatypes, e.g. MIN/MAX, DAY, LAST_DAY, BIT*, TIMESTAMP_ISO, VARIANCE/STDDEV/ with UNIQUE clause,
HEX() function	Supported	Not supported
Mixed Encodings	Adding EBCDIC tables when UNICODE tables already added is supported (SQL joins with data in Unicode and EBCDIC not possible)	Only supported to add UNICODE tables after EBCDIC table has already been added (required to set AQT_ENABLE_MULTIPLE_ENCODINGS environment variable)
CURRENT_TIME, CURRENT_TIMESTAMP, CURRENT_DATE	Supported with improved accuracy (no longer dependent on time synchronization)	Supported
Local Date Exit format	Not supported	Supported

Data load and update options



Synchronization options	Use cases, characteristics and requirements	Technical aspects
Full table load/refresh The entire content of a database table is loaded/refreshed	 Source table data is entirely replaced Smaller, un-partitioned tables Reporting based on consistent snapshot 	 Scope: Table or Partition ACCEL_LOAD_TABLES stored procedure Data Studio provides options to Load/Refresh a table/partitions
Table partition load/refresh <i>For a partitioned database</i> <i>table, selected partitions can</i> <i>be loaded/refreshed</i>	 More efficient than full table refresh for larger tables Reporting based on consistent snapshot Optionally: automatically load changed partitions only 	 Indicate changed partitions Queries can be routed while load is in progress
Incremental Update Log-based capturing of changes and propagation to Accelerator with low latency (typically few minutes)	 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 	 Scope: Row Based on Change Data Capture (CDC) of IBM InfoSphere Data Replication Management integrated into stored procedures and Data Studio to: Enable/Disable tables for replication Start/Stop replication



Loading data with Version 7.1

Smart Load Governor

- Data is loaded or refreshed using ACCEL_LOAD_TABLES stored procedure
 - Non-partititioned tables and table partitions are loaded in parallel
 - Max number of tables and partitions loaded in parallel determined by AQT_MAX_UNLOAD_IN_PARALLEL environment variable (load streams)
- AQT_MAX_UNLOAD_IN_PARALLEL applies to each stored procedure call
- The Accelerator checks the number of all load requests coming in in parallel (from one subsystem or multiple connected subsystems) and dependent on resource consumption of load requests, the percentage of resource consumption per subsystem and the max number of parallel load streams per subsystem will be adjusted
- Advantages
 - All tables or partitions of tables specfied in one ACCEL_LOAD_TABLES call are effeciently loaded in parallel. Load streams are efficiently used
 - Reduces the need for customers to manage load steams effeciently by calling ACCEL_LOAD_TABLES in parallel
 - Ensures that Accelerator manages the resources used for loading data proactively to prevent resource bottlenecks and overloading



Accelerator data load





Incremental update architecture





Db2 Analytics Accelerator Loader for z/OS

- Significant cost & time reduction by eliminating manual ETL processes of non-Db2 data
- Insight into more data types such as IMS, VSAM, sequential files, Oracle, Adabas, etc.
- Ensure High Availability for critical analytics applications
- **Protect critical analytics data** with backup & restore strategy
- Support cascading changes made to replicated tables to reduce data latency
- Enhance use of Accelerator-only tables with one-step load and load resume

Extend analytic capabilities by bringing non-Db2 data to IBM Db2 Analytics Accelerator and IBM Z

*Separately licensed option





Extends Capabilities of the Accelerator beyond Db2 data





Automates loading of non-Db2 data

Ξ	Accelerator Loader -	Accelerator Loader/Accelerator/loader.jcl - IBM Data Studio		
File Edit Navigate Search	h Project Data SQL Run Window Help			
💿 🖄 📩 👻 Activity:	Other Activity ▼ 🛤 👔 📓 💁 • 🖋 • 🗄 • 🗟 • 🖘 🗇 •	· · · · · ·	Quick Access	Process to identify and map
Studio N 20 Common Looks	Client Server Network Favorites	<pre>C Task Launcher</pre>		 source data Entire process is automated in one job Create Db2 table Table added to Accelerator Specified source data extracted In-memory conversion to Db2 format Data loaded to Accelerator Table is ready for acceleration



Loads to multiple Accelerators are **performed serially**

High Availability Load



PI59666: HA Load from Db2 for z/OS or external source **PI67510:** Specify accelerator group name



Loads to multiple

High Availability Load



PI59666: HA Load from Db2 for z/OS or external source **PI67510:** Specify accelerator group name



Accelerator Backup and Recovery

Expanded Data Load Capability Plus Data Protection

Makes sense for

- Data loaded to AOT or accelerator only
- Data that has been changed or altered in AOT
- Data loaded with LOAD RESUME over time

Allows

- Protection of investment spent in loading data
- Protection for data changed in an AOT
- Ability to potentially reduce backup of data at its source

Features

- Included with Accelerator Loader product
- Integrated backup, and fast recovery when needed
- Familiar DBA Functionality
- Fits into disaster recovery scenarios
- Fast Restore to copy point





Introducing Accelerator-only table type in Db2 for z/OS

Creation (DDL) and access remains through DB2 for z/OS in all cases

Non-accelerator DB2 table

• Data in DB2 only

Accelerator-shadow table

• Data in DB2 and the Accelerator

Accelerator-archived table / partitionEmpty read-only partition in DB2

• Partition data is in Accelerator only

Accelerator-only table (AOT)

- "Proxy table" in DB2
- Data is in Accelerator only



Not available for V7 yet



Accelerator-only tables

Supporting in-database transformation and multi-step processing



- Can be used to store a set of data in Db2 Analytics Accelerator only, not on Db2 for z/OS, without using the High Performance Storage Saver functionality
- Accelerate in-database data transformations and data movement processes
- Reduced need of data movement processes to other platforms for data transformation purposes
- Enables multi-step reporting on the Accelerator
- Saves disk space and CPU cost on IBM Z currently used for transformations and reporting steps
- Allow data preparation steps for data mining and other advanced analytics to execute on the Accelerator



Accelerator-only tables

Technical basics

- AOTs are created and dropped using Db2 DDL statements (CREATE; DROP)
 - Accelerator must be started
 - QUERY ACCELERATION behavior may have any value during CREATE/DROP
 - Syntax:
 - CREATE TABLE MYTABLE (...) CCSID ccsid IN ACCELERATOR <ACCEL1>;
 - DROP TABLE MYTABLE;
- Recommended to create a database in Db2 to be used for the AOTs
 - CREATE TABLE MYTABLE (...) CCSID ccsid IN ACCELERATOR <ACCEL1> IN DATABASE MYDB;
 - Usual authorization necessary to create objects in database
- SELECT and INSERT/UPDATE/DELETE operations using AOTs can only run on the Accelerator
 - QUERY ACCELERATION behavior must be set to ENABLE/ELIGIBLE/ALL
 - Accelerator-shadow tables and other AOTs can be used in the same statement
- Dynamic and static SQL can be used with AOTs



Multi-step reporting applications

With Accelerator-only tables: Temporary objects and processing on the Accelerator

Reporting Application



Data for transactional and analytical processing



Accelerator-only tables and ELT logic

In-database transformation



Data for transactional and analytical processing



Version 7.1.2 and Version 7.1.3



IBM Db2 Analytics Accelerator Version 7.1

Continuous delivery – fixes and new features every other month (on avg.)



Accelerator on IBM Z



Accelerator on IIAS (Appliance)

Enhancements in 7.1.2 and 7.1.3 (September) levels

- Incremental Update
- New datatypes: BINARY, DECFLOAT, ROWID
- Recursive SQL support
- Support for Db2 Analytics Accelerator Loader
- Accelerator monitor in Data Studio
- Configuration Console access via ssh
- Support for additional IIAS hardware sizes and configurations
 - M4002 series including M4002-003 GoD
- Additionally supported IBM Z hardware
 - z14 ZR1
 - z13 and Linux One (check firmware levels with IBM)

Incremental Update



- Keeps data in DB2 and the Accelerator in synch in near real-time
- Based on the Change Data Capture (CDC) component of IBM InfoSphere Data Replication
- INSERT/UPDATE/DELETE statements captured from DB2 log data and replicated to the Accelerator
- Tables enabled for incremental update require either an enforced uniqueness (primary key, unique index) or a defined informational constraint (via ACCEL_ADD_TABLES stored procedure)
 - Required for DELETEs
- Continuous replication
 - Base table not locked while table initially loaded to the Accelerator
 - Replication not stopped if replication subscription is changed (tables added, removed, loaded, reloaded)
- Accelerator V5 configurable options are all enabled per default and are not configurable any more
 - Continuous replication, Parallel Apply, Suspending faulty tables



Support for DECLOAT, BINARY, ROWID

- BINARY, VARBINARY
- DECFLOAT(16) and DECFLOAT(34)
- ROWID
- Add/Load table and query acceleration
- Scalar functions with these data types are supported
- Required Db2 APARs:
 - Db2 11: PI96928 (ROWID), PI92359 (DECFLOAT and BINARY)
 - Db2 12: PI99602 (ROWID), PI98408 (DECFLOAT and BINARY)
- Limitations
 - DECFLOAT constants that are outside the DOUBLE range are not supported. <u>https://www.ibm.com/support/knowledgecenter/SSEPGG_11.1.0/com.ibm.db2.luw.sql.ref.doc/doc/r0000731.html</u>
 - No ROWID support for accelerator-only tables



Recursive SQL – now fully supported

- Queries with recursive common table expressions
- Db2 APAR PI99670
 - Db2 11 PTF:UI57572
 - Db2 12 PTF:UI57573
- Example:

```
WITH temp (n, fact) AS
(SELECT 0, 1 -- Initial Subquery
    UNION ALL
    SELECT n+1, (n+1)*fact FROM temp -- Recursive Subquery
        WHERE n < 9)
SELECT * FROM temp;</pre>
```



Accelerator Monitoring in Data Studio



SSH Configuration Console



ssh configuration-console@vfb -p 2222 Last login: Tue May 8 09:11:36 2018 from dyn-9-152-204-199.boeblingen.de.ibm.com
To exit the console press CTRL+] followed by CTRL+d
Trying 10.101.13.48 Connected to 10.101.13.48. Escape character is '^]'.
Enter userid: acceladm
Enter password (in TSO, use PF3 to hide input): Licensed Materials - Property of IBM 5697-DA7

- Up until V7.1.2 the configuration console was only accessible via telnet
 - Telnet is an unencrypted protocol
 - Many companies require all network traffic to be encrypted
- SSH access with user "configuration-console" encrypts configuration console connections between the Accelerator and z/OS
 - Usage: ssh configuration-console@<accelerator-ip> -p 2222

Hardware Considerations for Accelerator on Z

- Accelerator software supported on IBM Z starting with z13
 - IBM z14 Type 3906 Models M01, M02, M03, M04 and M05
 - IBM z14 Type 3907 Model ZR1
 - IBM z13 Type 2964 Models N30, N63, N96, NC9 and NE1
 - IBM z13s Type 2965 Models N10 and N20
 - IBM LinuxONE Emperor II Type 3906 Models LM1, LM2, LM3, LM4 and LM5
 - IBM LinuxONE Emperor Type 2964 Models L30, L63, L96, LC9 and LE1
 - IBM LinuxONE Rockhopper II Type 3907 Model LR1
 - IBM LinuxONE Rockhopper Type 2965 Models L10 and L20
 - Check with IBM for required firmware levels
- Accelerator runs in SSC LPAR
 - 1 Accelerator = 1 LPAR but multiple Accelerator LPARs on a single system possible
 - Each LPAR requires IFLs, RAM, and storage
 - IFLs can be shared
 - No other software needed (no z/VM, no Linux distribution,...)
 - Not under z/VM or KVM



Version 7.1.4

Why "True HTAP"?



• Problem Statement

When a table changes on Db2 for z/OS, the changes are not immediately visible to queries routed to IDAA. However, some applications require the guarantee that any change committed on Db2 will be seen by a subsequent query

The HTAP feature provides this guarantee, **presenting the next step in the journey towards full query offload transparency** for accelerated queries

Analytics is performed against COMMITTED data that is known to be current relative to the SQL Latency no longer impacts SQL result consistency

V5 PTF6: True HTAP Technical Preview available January 2018 V5 PTF6a: True HTAP Production Ready available February 2018 V7.1.4: True HTAP available October 2018

What is True HTAP?



DB2 Analytics Accelerator

for z/OS

- Single System, Single Interface ٠
- Data coherency between Db2 and Accelerator ٠
- Don't worry about latency of committed data •
 - Queries run first when data available
- No longer need to explain to users why committed data not seen. •

Fransactional

Processing

Analytical processing

- Data used by SELECT contains corresponding commits ٠
- e.g. Use case "Fraud Detection"
 - Detection query works on most recent data ٠
- e.g. Use case "Multi form consistency"
 - Enter data in a field on a panel. Application commits data. ٠ => Just entered data shown on following panels









HTAP enablement

Customization of WAITFORDATA

- Two new options in the replication submenu of the configuration console
- Both settings are appliance-global Apply to all queries



(14) - Change the allowed table types for WAITFORDATA queries

How is HTAP implemented?



- Introducing a new ZPARM + Bind option + Special register
- CURRENT QUERY ACCELERATION WAITFORDATA = n.m
 - n.m = 0.0 3600.0 (seconds)
 - Default: 0.0 = No wait
 - Important: Can be set differently for each query
- WAITFORDATA = 0.0
 - Immediately execute in accelerator (Current behavior, no delay)
- WAITFORDATA > 0.0
 - Wait for committed changes to be applied via asynchronous replication
 - If wait time is exceeded check **Delay Expiration** rule (Option 13):
 - If "CONTINUE", immediately execute in accelerator and return SQLCODE +904
 - If "FAIL", check CURRENT QUERY ACCELERATION special register
 - If "WITH FAILBACK" is specified, execute query in Db2
 - Else FAIL query with SQLCODE -904



Monitoring HTAP

- IFCID 2, 316 and 401 SMF records have new fields
 - Query requests, Query requests expired, Replication velocity
 - Statement wait time, Statement max wait time expired
- -DIS ACCEL DETAIL includes new monitoring counters

CURRENT REPLICATION LATENCY FOR THIS DB2 SYSTEM=2000 MSNUMBER OF SUCCESSFUL QUERY REQUESTS WITH DELAY PROTOCOL FOR ALL DB2 SYSTEMS=270NUMBER OF EXPIRED QUERY REQUESTS WITH DELAY PROTOCOL FOR ALL DB2 SYSTEMS=15REPLICATION VELOCITY (DB2 LOG SECONDS APPLIED PER SECOND)=1

- ACCEL_GET_QUERIES shows a delayed query as "QUEUED"
- No related changes to Data Studio
 - Wait time shown in Data Studio does not include HTAP wait. Info only in SMF so far.







Links

- Release Notes V7.1.2
 - http://www-01.ibm.com/support/docview.wss?uid=swg27051088
- Release Notes V7.1.3
 - <u>https://www-01.ibm.com/support/docview.wss?uid=ibm10716267</u>
- Release Notes V7.1.4
 - http://www-01.ibm.com/support/docview.wss?uid=ibm10732149
- Prerequisites and maintenance
 - http://www-01.ibm.com/support/docview.wss?uid=swg27050440
- Knowledge Center
 - <u>https://www.ibm.com/support/knowledgecenter/SS4LQ8_7.1.0/com.ibm.datatools.aqt.doc/idaa_kc_welcome.html</u>



Join ibm.biz/db2analyticsaccelerator



IBM Bluemix	Develop in the cloud at the click of a button!				
IBM					Sign in
IBM developerWorks _®	Technical topics	Evaluation software	Community	Events	Search devel
	My home Forums	Blogs Communities	Profiles Poo	lcasts Wikis	Activities
	IBM Champion program	n			
Communities			(1)	This Community	 Search

IBM DB2 Analytics Accelerator Community

IBM DB2 Analytics Accelerator Community

- Community Description



Welcome to the IBM DB2 Analytics Accelerator Community!

IBM DB2 Analytics Accelerator is a workload-optimized appliance that integrates zEnterprise and PureData technologies to accelerate relevant data-intensive and complex DB2 for z/OS queries. Together, the DB2 Analytics Accelerator, DB2 for z/OS and zEnterprise form an integrated hybrid environment that can run transaction processing, complex analytical and reporting workloads concurrently and efficiently.

Blog **66** New technote "How to assess queries for eligibility" available queryassessment | 96 Visits **11** Incremental Update -**ONUTILITYACTION** parameter update | 1 Comment | 94 Visits ۵ View All Forums Important Bookma 👫 DB2 Analytics Accelerat and the accelerator DB2 Analytics Accelerat **15** IBM Knowledge Center archived via HPSS ? Accelerator for z/OS 4.1.0

Page

Page

Members

🖹 1 🕘 🕘 🔛 c

 \equiv 🙂 O 🛛 🗮 UteBaumbach | May 7 | Tags: gueryacceleration 🙂 0 🗢 Patric Becker | May 5 | Tags: cdc incremental \equiv Friendly arithmetic in DB2 for z/OS - SQLCODE +802 Latest post by UteBaumbach | Monday 5:57 PM | No replies How does the new 10 bytes RBA/LRSN affect tables

Latest post by UteBaumbach | Monday 5:55 PM | No replies

CP improvements for loading data into IBM DB2 Analytics Accelerator Latest post by UteBaumbach | Monday 5:54 PM | No replies

Overview





Legal Disclaimer

- © IBM Corporation 2018. All Rights Reserved.
- The information contained in this publication is provided for informational purposes only. While efforts were made to verify the completeness and accuracy of the information contained in this publication, it is provided AS IS without warranty of any kind, express or implied. In addition, this information is based on IBM's current product plans and strategy, which are subject to change by IBM without notice. IBM shall not be responsible for any damages arising out of the use of, or otherwise related to, this publication or any other materials. Nothing contained in this publication is intended to, nor shall have the effect of, creating any warranties or representations from IBM or its suppliers or licensors, or altering the terms and conditions of the applicable license agreement governing the use of IBM software.
- References in this presentation to IBM products, programs, or services do not imply that they will be available in all countries in which IBM operates. Product release dates and/or capabilities referenced in this presentation may change at any time at IBM's sole discretion based on market opportunities or other factors, and are not intended to be a commitment to future product or feature availability in any way. Nothing contained in these materials is intended to, nor shall have the effect of, stating or implying that any activities undertaken by you will result in any specific sales, revenue growth or other results.

We want your feedback!



- Please submit your feedback online at

 http://conferences.gse.org.uk/2018/feedback/II
- Paper feedback forms are also available from the Chair person
- This session is II





