

# GSE 2019 – Accelerate your analytics

## IBM Db2 Analytics Accelerator – An exciting 2019



**GSE session number: ID**

**Chris Watson – IBM Analytics on IBM Z Center of Excellence**

[watson@de.ibm.com](mailto:watson@de.ibm.com) (Linked In: Christopher Watson)

## Legal Disclaimer

- © IBM Corporation 2019. 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.

# Agenda

- **IBM Db2 Analytics Accelerator (a.k.a. IDAA)**
- What's New 2019 - V7.1
- What's Coming 2019 - V7.5
- What's New 2019 - Documentation



# IDAA Data Gravity approach to analytics

---

## **Simplified infrastructure with more resiliency**

- One copy of the data not dozens

## **Much more secure**

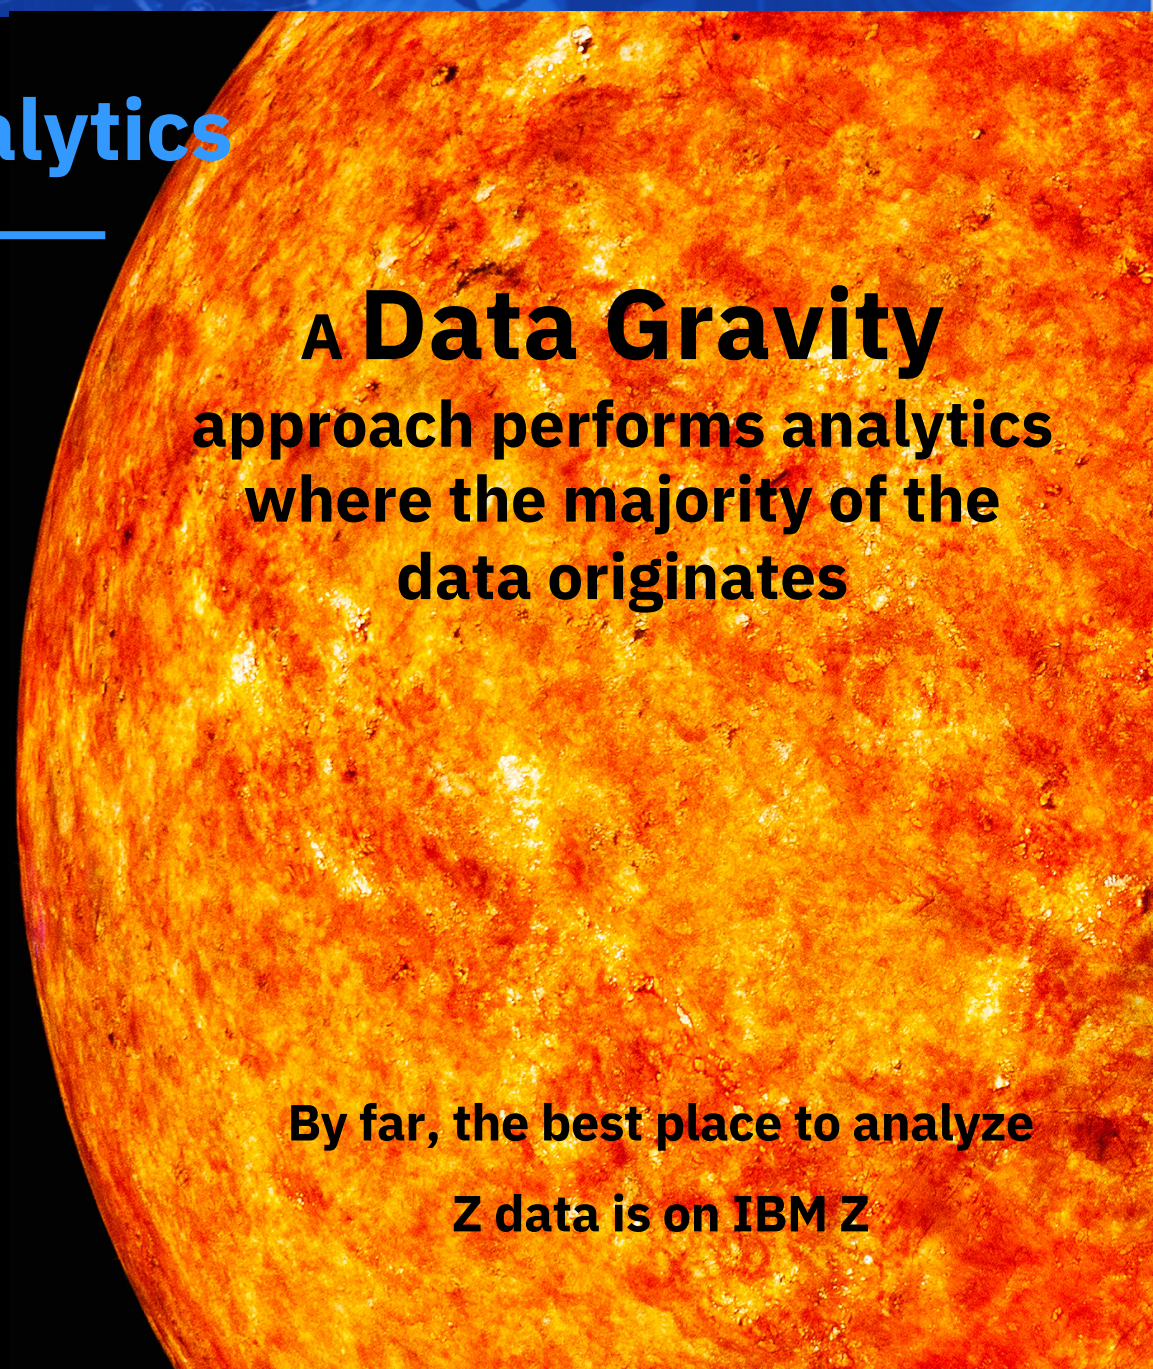
- Z security built in

## **Lower cost**

- Saves money (Infrastructure, SW, and people)

## **Much lower analytics latency**

- Low to no latency with transactional data



**A Data Gravity  
approach performs analytics  
where the majority of the  
data originates**

**By far, the best place to analyze  
Z data is on IBM Z**



# Db2 Analytics Accelerator on IBM Z

## Fully integrated software appliance

- Extends Db2 on Z “in the box”
- Runs full SW stack in a **Secure Service Container (SSC)** LPAR
- “Download & go” – Very easy to deploy + update

## Leverages IBM Z infrastructure

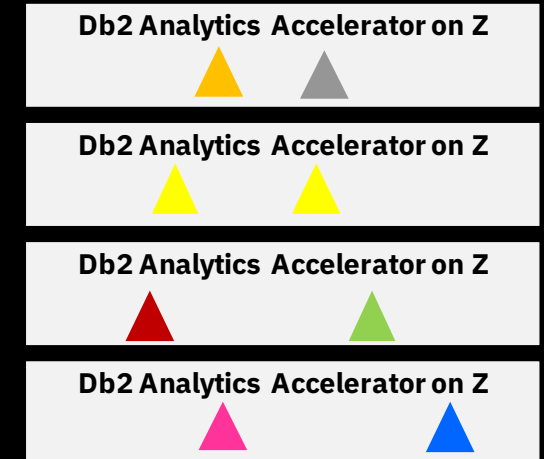


- z13 upwards and LinuxONE
- Leverage unused capacity or dedicated stand-alone system
  - Single-node: 2 – 40 IFLs & 4TB memory    Multi-node 30 – 190 IFLs & 20TB memory
- Leverage existing Z HA/DR solution - GDPS/Metro integration



# Deployment Considerations

## Accelerator Instances 'On IIAS' and 'On Z'

- Single vs. Multiple Accelerators “per box”

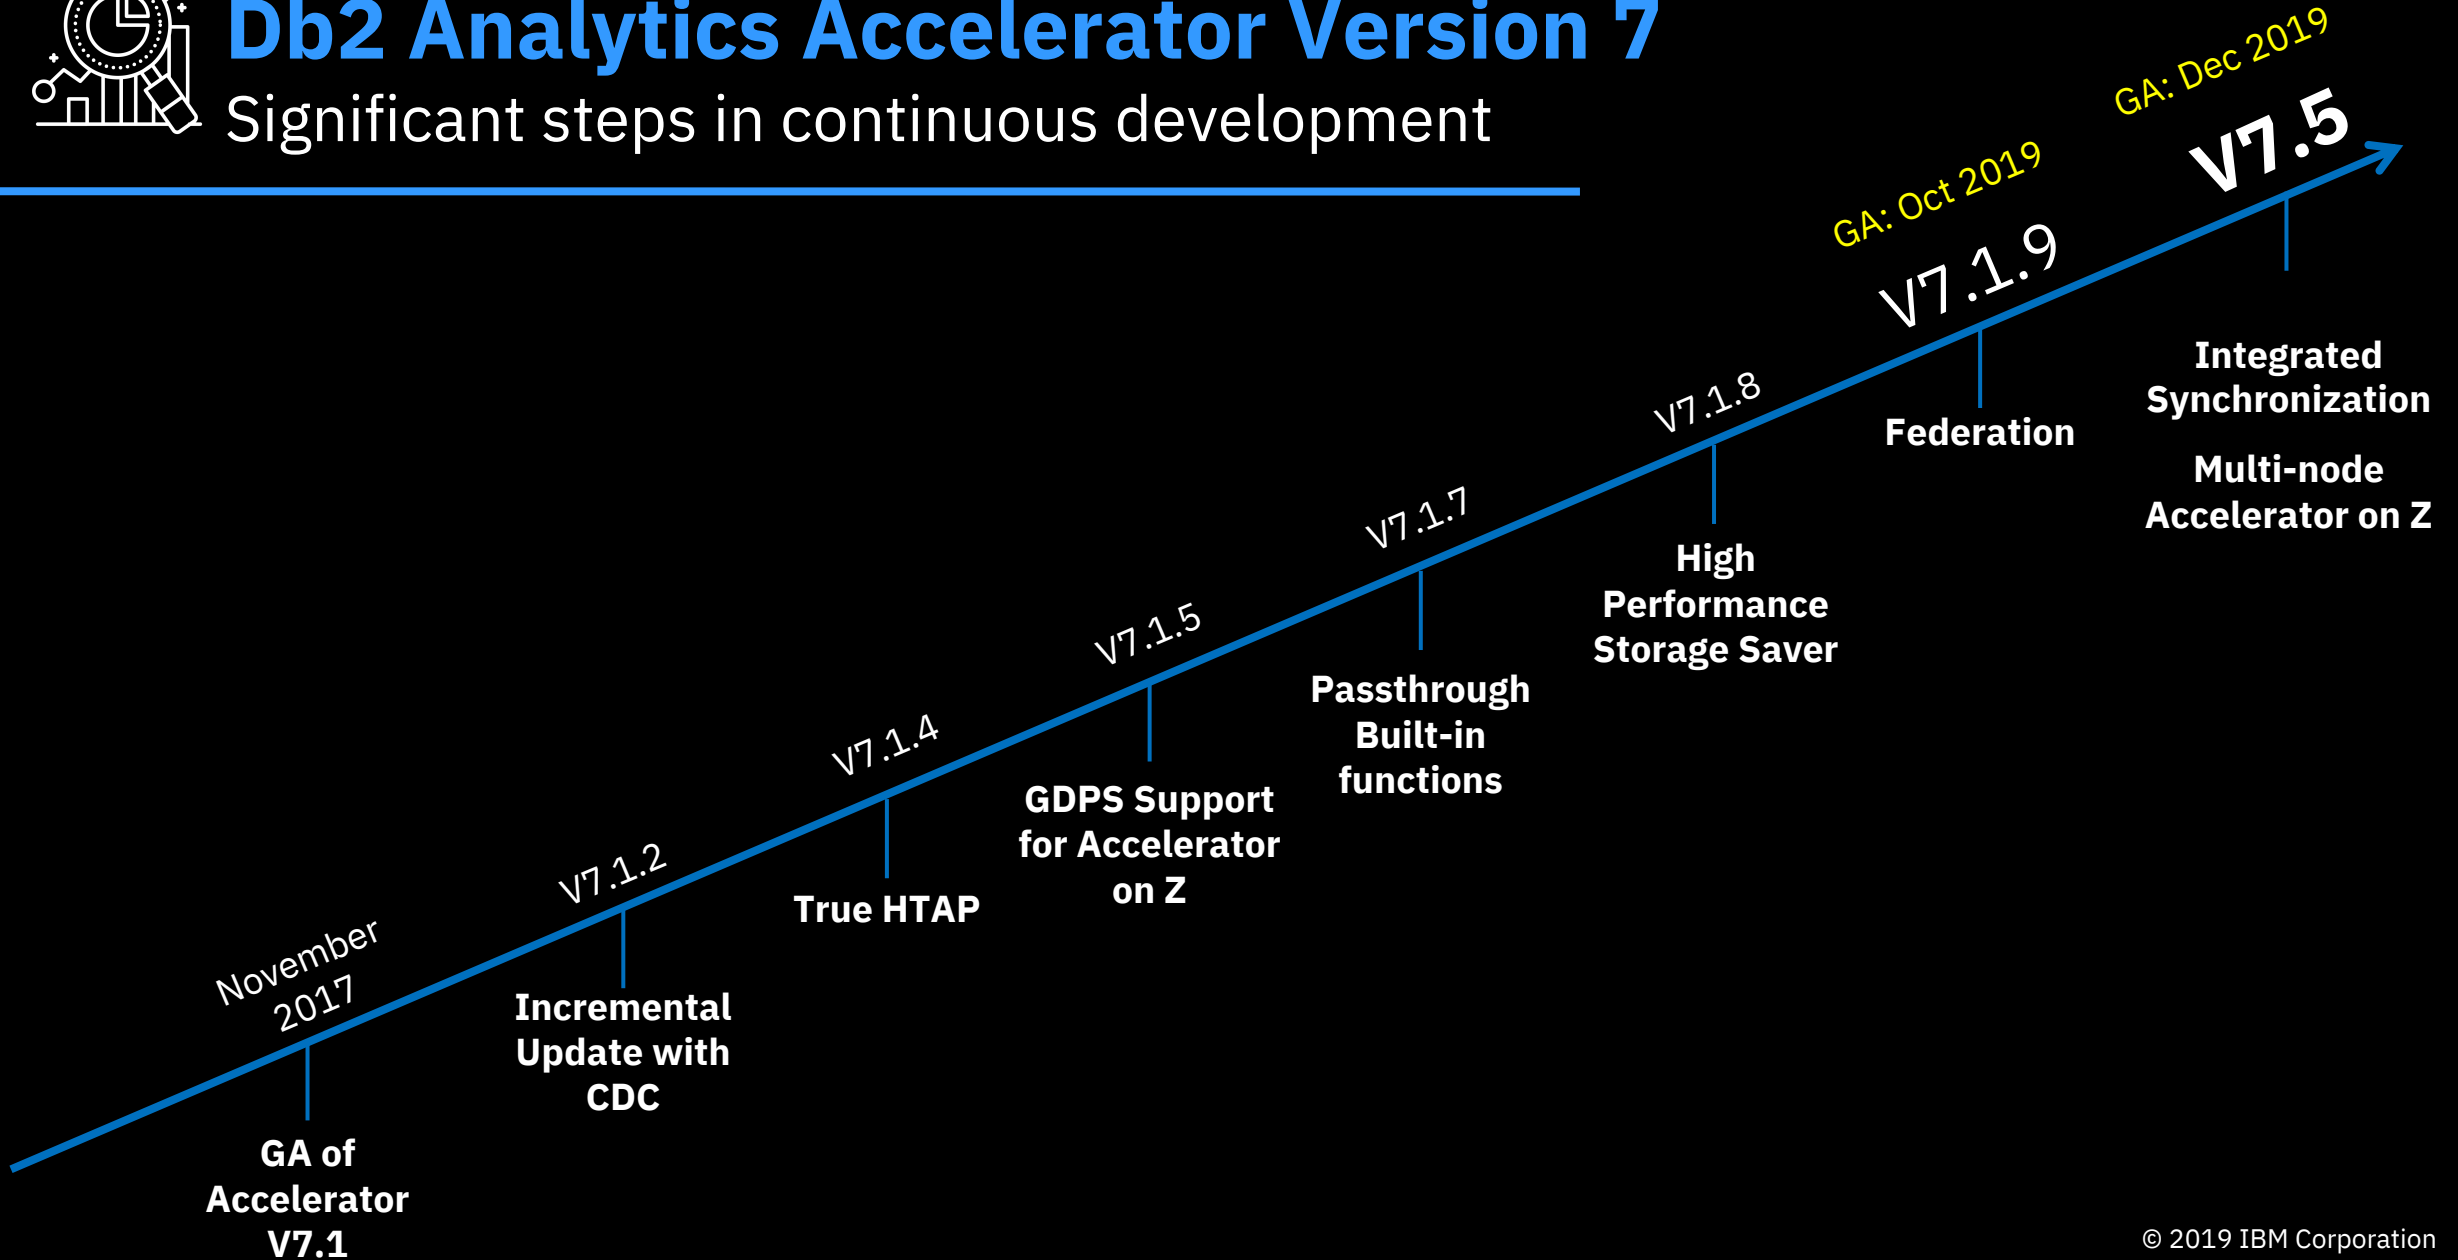


- Consider splitting up large accelerators to smaller instances, for example to ...
  - Isolation of workloads, selective support of HA/DR, address LOB-/workload specific requirements, more flexible management (e.g., maintenance/update/...)
- Consider consolidating multiple accelerators on the same CPC, for example to ...
  - Reduce cost by sharing IFLs, flexible DEV/TEST environments, flexible resource adjustments between accelerators
- Consider creating additional accelerator instances, for example to help with
  - Update or new feature testing, support for new workload development, sandboxes,...



# Db2 Analytics Accelerator Version 7

Significant steps in continuous development





# Agenda

- IBM Db2 Analytics Accelerator (a.k.a. IDAA) - Motivation
- What's New 2019 - Documentation
- **What's New 2019 - V7.1**
- What's Coming 2019 - V7.5

# IBM Db2 Analytics Accelerator Version 7.1

*Continuous delivery – Fixes and new features approx every other month*



Accelerator on IBM Z



Accelerator on IAS (Appliance)

## Enhancements in 7.1.6 (Feb 2019)


- **Encryption of data-in-motion:** Network traffic optionally encrypted
- **Call Home for Accelerator on IAS,** via HMC

## Enhancements in 7.1.7 (April 2019)

- **High Performance Storage Saver (HPSS)** – Technical Preview
- Support for more **pass-thru built-in functions** (e.g. REGEXP\*)
- **Improved compression:** Full table reloads, alter key, or archive actions
- Support for **Partition reloads of replicated tables**

# High Performance Storage Saver

## Storing historical data in Accelerator only



### Support for partitioned-by-range tables, for example:

Time-based partitions, where only the recent partitions are used in a transactional context (frequent data changes, short running queries), but the entire table is used for analytics, regulatory/audit purposes, etc. (data intensive, complex queries)



# Pass-through support for Db2 Warehouse built-in functions

## *Enhancing Db2's native SQL Capabilities with the Accelerator*

- Many built-in functions that are supported by the underlying DBMS (Db2 Warehouse) in the Accelerator are not supported natively by Db2 for z/OS (yet).
- Some functions can now be used in SQL queries routed to the Accelerator with the new Built-In-Function (BIF) Pass-through support
  - Db2 for z/OS is "aware" of the Accelerator, **when parsing** the SQL statement.
  - If a BIF is referenced, which is only available on the Accelerator, the Db2 for z/OS parser validates the signature and allows its invocation within the rewritten SQL
  - Db2 for z/OS still needs to validate parameters, return types, .... Therefore limited to commonly requested BIFs
- Accelerator V7.1.5: First set of supported BIFs (OLAP functions):
  - CUME\_DIST, FIRST\_VALUE, LAG, LAST\_VALUE, LEAD, NTH\_VALUE, NTILE, PERCENT\_RANK, RATIO\_TO\_REPORT
- Added in V7.1.7:
  - Aggregate functions: CUME\_DIST, PERCENT\_RANK
  - Scalar functions: REGEXP\_COUNT, REGEXP\_INSTR, REGEXP\_LIKE, REGEXP\_REPLACE, REGEXP\_SUBSTR
- Db2 12 only, FL504

# IBM Db2 Analytics Accelerator Version 7.1

*Continuous delivery – Fixes and new features approx every other month*



Accelerator on IBM Z



Accelerator on IAS (Appliance)

## Enhancements in 7.1.8 (June 2019)

- High Performance Storage Saver (HPSS) – GA Release
- Improved compression (Initial load, Re-load, Alter key, Archive)
- Mapping of Db2 for z/OS REAL column type to DOUBLE on Accelerator
- Call-back interfaces for external load programs
- Several Stability and User Experience Improvements, e.g.
  - Query performance: automatic column-group statistics
  - Improved table size display
  - Improved fetch performance for queries with large resultsets
- Integrated Synchronisation: Early Support Program

# IBM Db2 Analytics Accelerator Version 7.1

*Continuous delivery – Fixes and new features approx every other month*



Accelerator on IBM Z



Accelerator on IIAS (Appliance)

## Enhancements in 7.1.9 (Oct 2019)

- Federation
  - Enable read access to all data (from multiple Db2 subsystems) in one Accelerator through a single DB2 subsystem
- “True” HTAP (Db2 z/OS 12 enhancements)
  - Support for **static** SQL
  - Shorter wait times with **object level wait**
- Workload balancing between Accelerator V5 and V7 (Db2 z/OS 12)
- Load performance improvements
- Software Update
  - Show software update installation progress
  - Essential-only install option to reduce SW update time
  - SW Update robustness and stability



# IBM Db2 Analytics Accelerator Version 7.1

*Continuous delivery – Fixes and new features approx every other month*



**Accelerator on IBM Z**



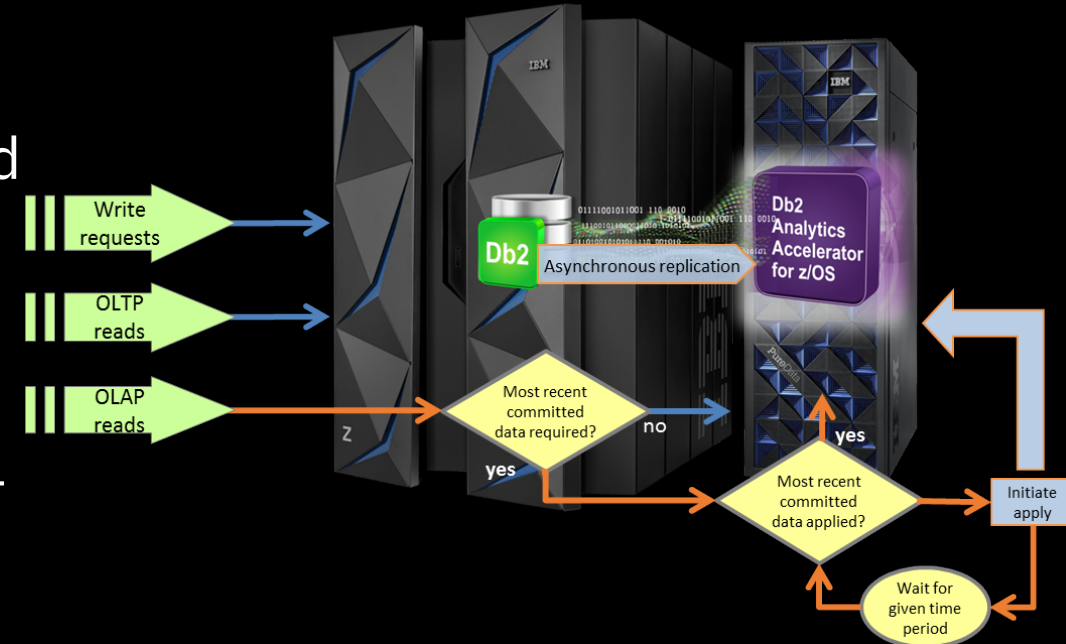
**Accelerator on IAS (Appliance)**

## Enhancements in 7.1.9 (Oct 2019)

- Several Stability and User Experience Improvements
  - Improved system level space reporting
  - Warning for tables with missing statistics
  - Diagnostics (Save Trace improvements)
  - Compression: Autogrooming to compress uncompressed rows
  - Accelerator on Z: Isolation and limitation of various accelerator storage areas to limit impact of running out of disk space
  - Replication using CDC: Reduced memory consumption on Accelerator to enable replication for up to 10 paired Db2 subsystems in parallel
  - Improved query performance when using TIMESTAMPDIFF function

## “True HTAP” support for static queries

- HTAP support in the Accelerator removes the need to manage the replication latency between Db2 and the Accelerator.
- Ability to delay the execution of queries until the latest incremental updates have been applied.
- Guarantees that any change committed on Db2 will be seen by a subsequent query executed on the Accelerator.
- Dynamic SQL support available since V7.1.4
- Static SQL support added with V7.1.9



- Prerequisites:
  - Db2 12 with APAR PH14116
    - Adds a new column ACCELERATION\_WAITFORDATA to SYSACCEL.SYSACCELERATEDPACKAGES

## “True HTAP” support for static queries (contd.)

- New bind option: ACCELERATIONWAITFORDATA
  - For BIND, BIND COPY and REBIND PACKAGE
  - For BIND DEPLOY for SQL PL procedure package
  - QUERYACCELERATION bind option must be used in combination
- New option: ACCELERATION WAITFORDATA
  - For CREATE/ALTER PROCEDURE/FUNCTION for SQL PL procedures and SQL PL scalar functions
- Statement level monitoring counters in IFCID 401
  - E.g. Accumulated wait time, Accumulated number of expired statements for that wait time



# “True” HTAP – Object Level LRSN

- **Queries do not wait for non-relevant changes**

- DB2 sends LRSN for the objects referenced in the query to accelerator

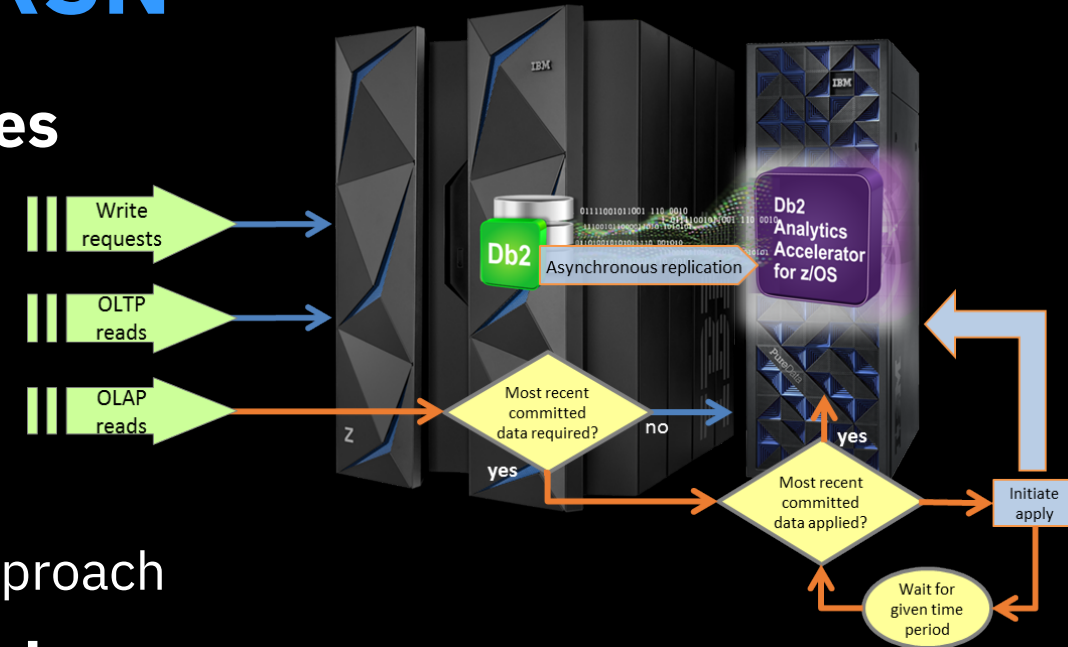
- Example: UPDATE AAA; COMMIT  
UPDATE BBB; COMMIT  
SELECT X FROM AAA => Waits for AAA **only**

- Advantage: Shorter wait times as with previous approach

- **Object granularity is on the table space level**

- Multiple tables in a single table space will have the same LRSN regardless of which table is updated
- Recommended to have a single table per table space (UTS)

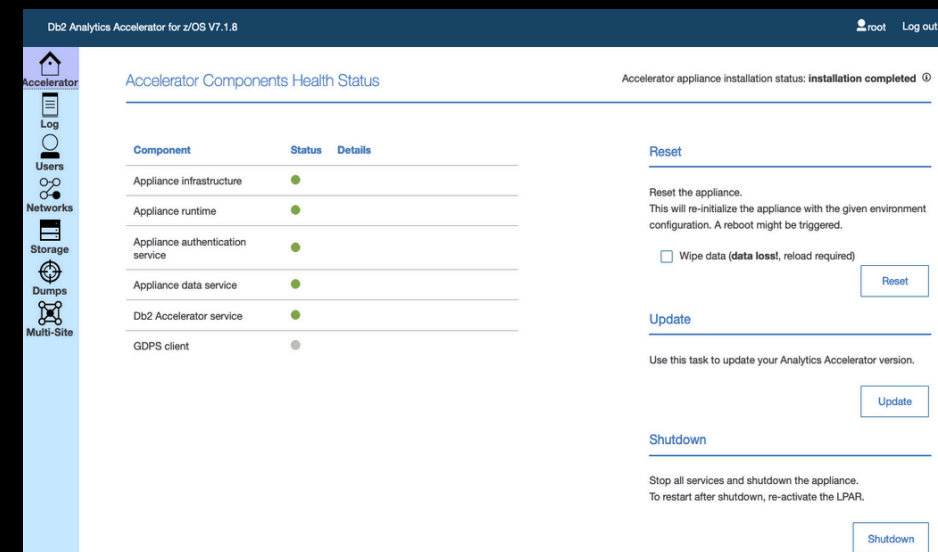
- Prerequisites: Db2 v12 FL 505 with APAR PI98126 and Accelerator v7



# Accelerator “on Z” setup via JSON configuration files

## Up to 7.1.8:

1. Define and activate SSC LPAR via HMC
2. Upload Accelerator image with SSC installer
3. Use the Accelerator Admin UI:
  - Complete network / storage settings in UI panels
  - Apply settings and initialize the Accelerator



## With 7.1.9:

1. Define and activate SSC LPAR via HMC
2. Upload Accelerator image with SSC installer
3. Using the Accelerator Admin UI:
  - Upload a JSON configuration file containing network and storage settings
  - Apply settings and initialize the Accelerator

```

1- {
2   "accelerator_name": "lgor01",
3   "accelerator_description": "SAMPLE-single-node",
4   "accelerator_type": "single-node",
5   "db2_posting_ipwr": "10.1.1.181/24",
6   "network_interface_bindings": {
7     "mgmt_nm": "activation-profile",
8     "db2_nm": "osa-0a-p0"
9   },
10  "runtime_environments": [
11    {
12      "cpc_name": "CPC001",
13      "lpar_name": "LPG001",
14      "network_interfaces": [
15        {
16          "name": "osa-0a-p0",
17          "device": "0.0.0c00",
18          "port": "0"
19        }
20      ]
21    }
22  ],
23  "primary_storage": {
24    "boot_device": {
25      "type": "dasd",
26      "device": "0.0.5e29"
27    },
28    "runtime_devices": {
29      "type": "dasd",
30      "devices": [
31        "0.0.5e25",
32        "0.0.5e26"
33      ]
34    }
35  },
36  "data_devices": {
37    "type": "dasd",
38    "devices": [
39      "0.0.5e14",
40      "0.0.5e80",
41      "0.0.5e8f"
42    ]
43  },
44  "hyperpav": "auto"
45 }
46 }
47 }

```

# Sample Config File: *Simple, Single-Node Accelerator*

```

1 {
2   "accelerator_name": "igor01",
3   "accelerator_description": "SAMPLE-single-node",
4   "accelerator_type": "single-node",
5   "db2_pairing_ipv4": "10.1.1.101/24",
6   "network_interface_bindings": {
7     "mgmt_nw": "activation-profile",
8     "db2_nw": "osa-0A-p0"
9   }
10  "runtime_environments": [
11    {
12      "cpc_name": "CPC001",
13      "lpar_name": "IGOR01",
14      "network_interfaces": [
15        {
16          "name": "osa-0A-p0",
17          "device": "0.0.0a00",
18          "port": "0"
19        }
20      ]
21    }
22  ]
23  "primary_storage": {
24    "boot_device": {
25      "type": "dasd",
26      "device": "0.0.5e29"
27    },
28    "runtime_devices": {
29      "type": "dasd",
30      "devices": [
31        "0.0.5e25",
32        "0.0.5e26"
33      ]
34    },
35    "data_devices": {
36      "type": "dasd",
37      "devices": [
38        "0.0.5e14",
39        [
40          "0.0.5e80",
41          "0.0.5e8f"
42        ]
43      ]
44    },
45    "hyperpav": "auto"
46  }
47 }

```

- Complete example for a basic accelerator:
  - Network interfaces for administration and Db2 acceleration
  - ECKD (DASD) storage for the various required storage pools
- General accelerator properties
- Runtime environments
  - Specify each LPAR the accelerator should run in
  - Network interfaces valid in this LPAR
- Storage definition
  - “primary\_storage” is the set of storage devices used for the initial installation of the accelerator

# Agenda

- IBM Db2 Analytics Accelerator (a.k.a. IDAA) - Motivation
- What's New 2019 - Documentation
- What's New 2019 - V7.1
- **What's Coming 2019 - V7.5**



# Introducing Db2 Analytics Accelerator Version 7.5

Unprecedented enterprise-grade transactional and analytic processing

## Db2 Analytics Accelerator Version 7.5 delivers:

- A wider range of scalability for Db2 Analytics:  
**Multi-node** Accelerator on IBM Z deployments
  - From very small to very large
- **Integrated Synchronization (a.k.a. InSync)**  
A new advanced data synchronization technique

**Announcement:**

October 15, 2019

**GA:**

December 6, 2019





# Very Large Accelerator Deployments

*aka Multi-Node Deployment*

TechPreview in V7.1.9

GA in V7.5

## Background

- IAS accelerator deployments are available in different sizes, ranging from M4002-002GoD (1/3-rack “growth-on-demand” with 3 capped server nodes) to M4002-020 (double-rack with 13/14 server nodes)
- In V7.1 Z accelerator deployments are supported from very small (2 IFLs, 64 GB) to large (40 IFLs, 4 TB)
  - This extends the offering at the low end → significantly smaller deployments than on IAS are possible and valuable
  - It fully covers the majority of actual accelerator deployment sizes
  - But it represents a limitation for growth at the high end, since 35 IFLs / 2.5 TB provide performance/capacity “comparable” to M4002-003

**→ The new “Multi-Node Deployment Option on IBM Z” will close that gap and is designed to scale up to performance/capacity comparable to M4002-020**

# Accelerator on IBM Z – Deployment Options

SINGLE-NODE	MULTI-NODE
1 Accelerator = 1 LPAR	1 Accelerator = 6 LPARs on the same CPC
2 – 40 IFLs, 64 GB – 4 TB memory (32 GB per IFL minimum requirement)	30 – 190 IFLs, 1.25 TB – 20 TB memory (total) (limited by maximum #IFLs on single CPC – z15 → 190)
Accelerator size can be flexibly adjusted within full range without data reload	
1 LPAR for combined head + data	1 LPAR group with absolute capping at total #IFLs
	1 LPAR for head: 30 IFLs (shared), 256 GB memory
	5 LPARs for data: (total #IFLs / 5) IFLs (shared), 256 GB – 4 TB memory
Ideal for “white space” deployments in existing CPCs	Ideal for deployment on a dedicated IBM Z or LinuxONE system (can be shared between multiple accelerators and/or other Linux on Z workloads)
Recommended up to 20 TB uncompressed user data	Recommended up to 100 TB uncompressed user data
Less if special performance requirements exist, e.g., IU latency, query elapsed time,...	
More in special cases like online archives	

# Accelerator on IBM Z – Scaling Targets ,c‘ = IFL


Netezza

N3001-02	N3001-05		N3001-10				N3001-20	
----------	----------	--	----------	--	--	--	----------	--

Accelerator on Z Single-node

← 1 →  
2c – 40c  
0.125 TB – 4 TB

**This chart shows design goals and preliminary data.**  
Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.




Accelerator on Z Multi-node

preliminary



Accelerator on IIAS



# Integrated Synchronization – “InSync”

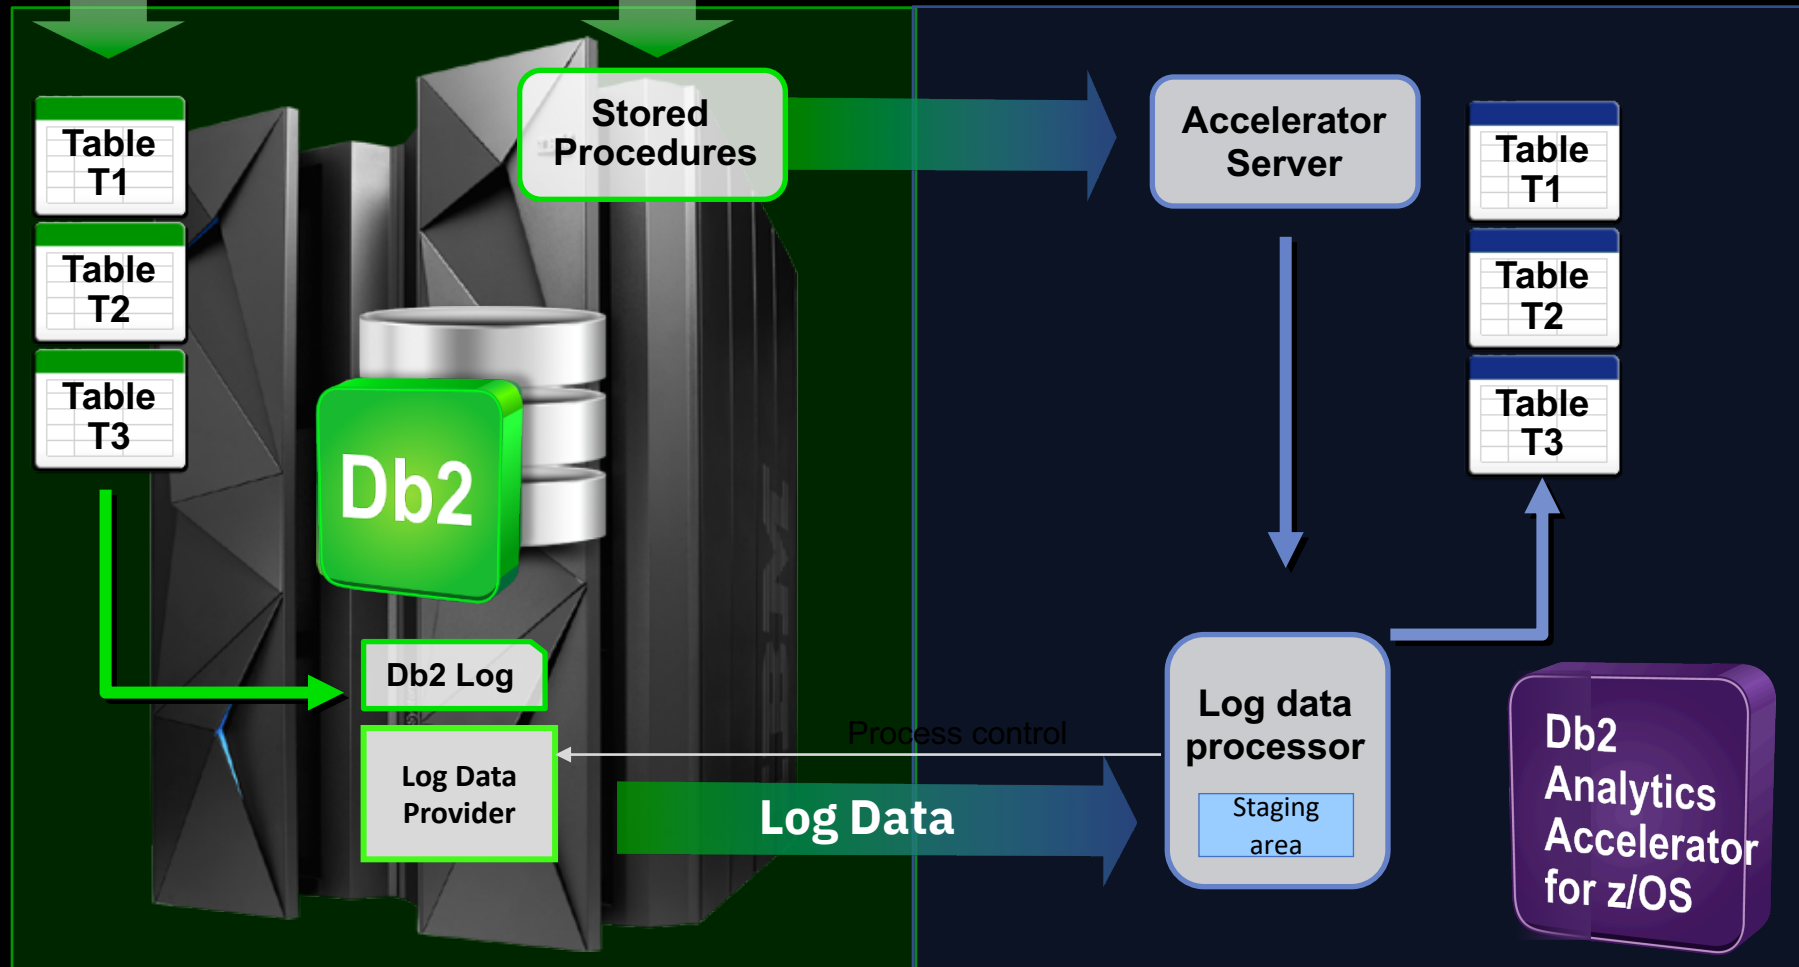
## Deeper integration of Db2 and Accelerator

- Integrated, low-latency data coherence protocol between Db2 for z/OS and the Db2 Warehouse engine inside the Accelerator
  - zIIP enabled
  - Complete application transparency
  - Enterprise-grade HTAP enabler
- Allows better handling of non-logged changes to Db2 tables
  - Future item to be able to replicate selected non-logged utility actions, such as load with dummy partition or reorg discard full partition
- Simplified installation, packaging, administration, upgrade

# Integrated Synchronization - Db2/Z to-Accelerator data synchronization

Applications executing  
I/U/D Statements on replicated tables

Accelerator Users enabling  
tables for replication





# Initial Performance Test Environment

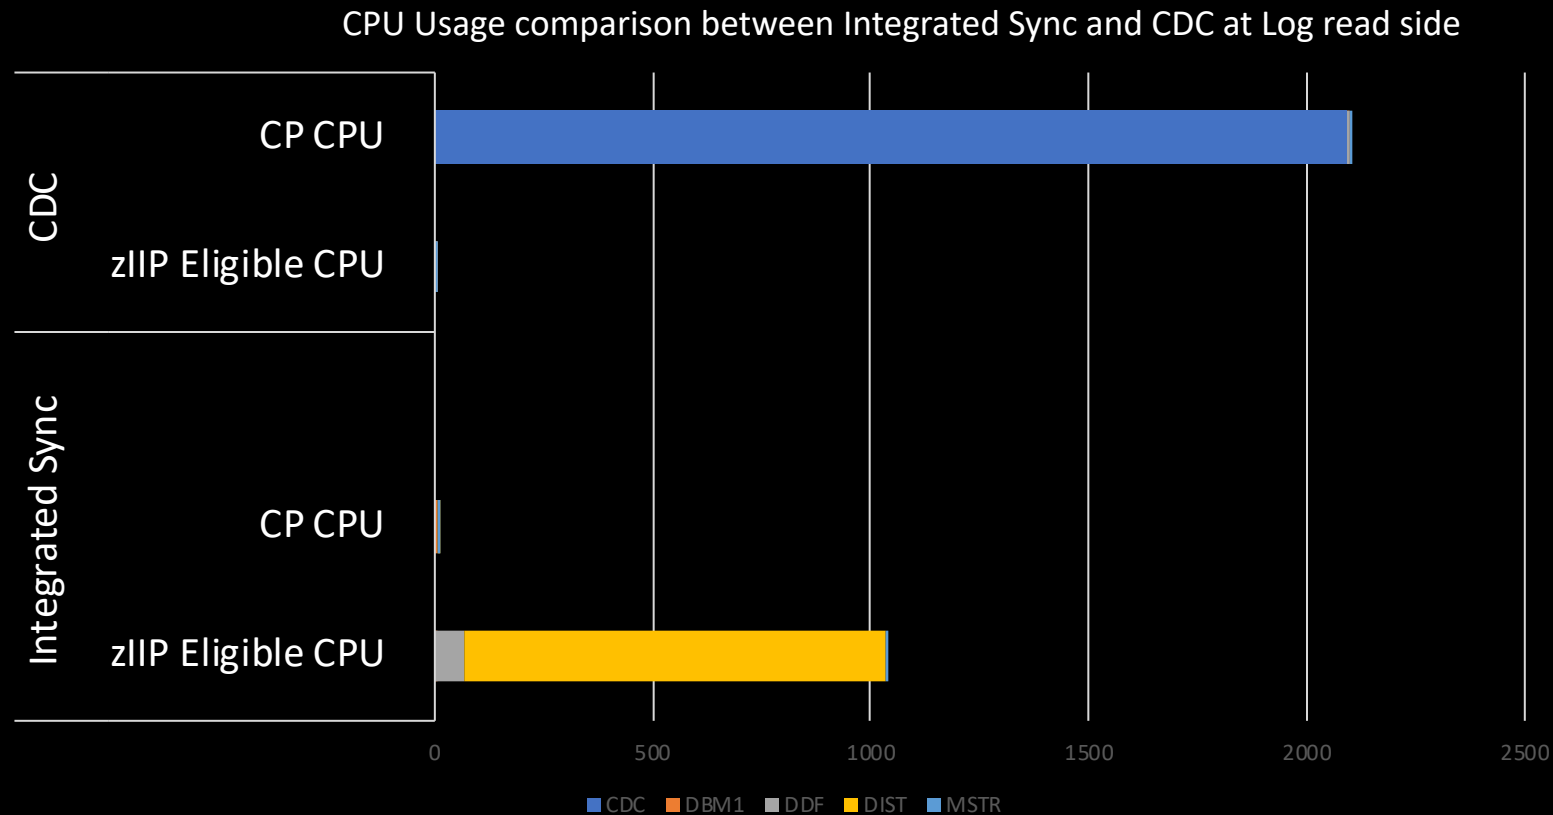
## Hard- and Software Configuration

- 6 z/14 GCPs and 2 ZIIPs
- 96G Memory
- 10Gbit network to Accelerator
  
- z/OS 02.02.00
- Db2 for z/OS v12
- M4002-010 (IIAS 1-rack)

## Test Workload and Test Case

- Total of 500 partitioned tables with concurrent insert/updates.
  - 50 Java clients committing every few seconds
- Max Db2 transaction rate with distributed Java test case
  - 64.4K rows/second inserted
  - 64.4K rows/second updated
  
- Scenario 1: Capture and apply with **concurrent Db2 transactions**
- Scenario 2: Capture and apply the **pre-populated log**

# CP and %zIIP-Eligible Comparison of InSync and CDC (Prepopulated logs)



**InSync uses about 49% of CPU time vs. CDC for the same work: CDC 2114 seconds vs InSync 1076 seconds**  
**Combined InSync work across all address-spaces is 98% zIIP offloadable**

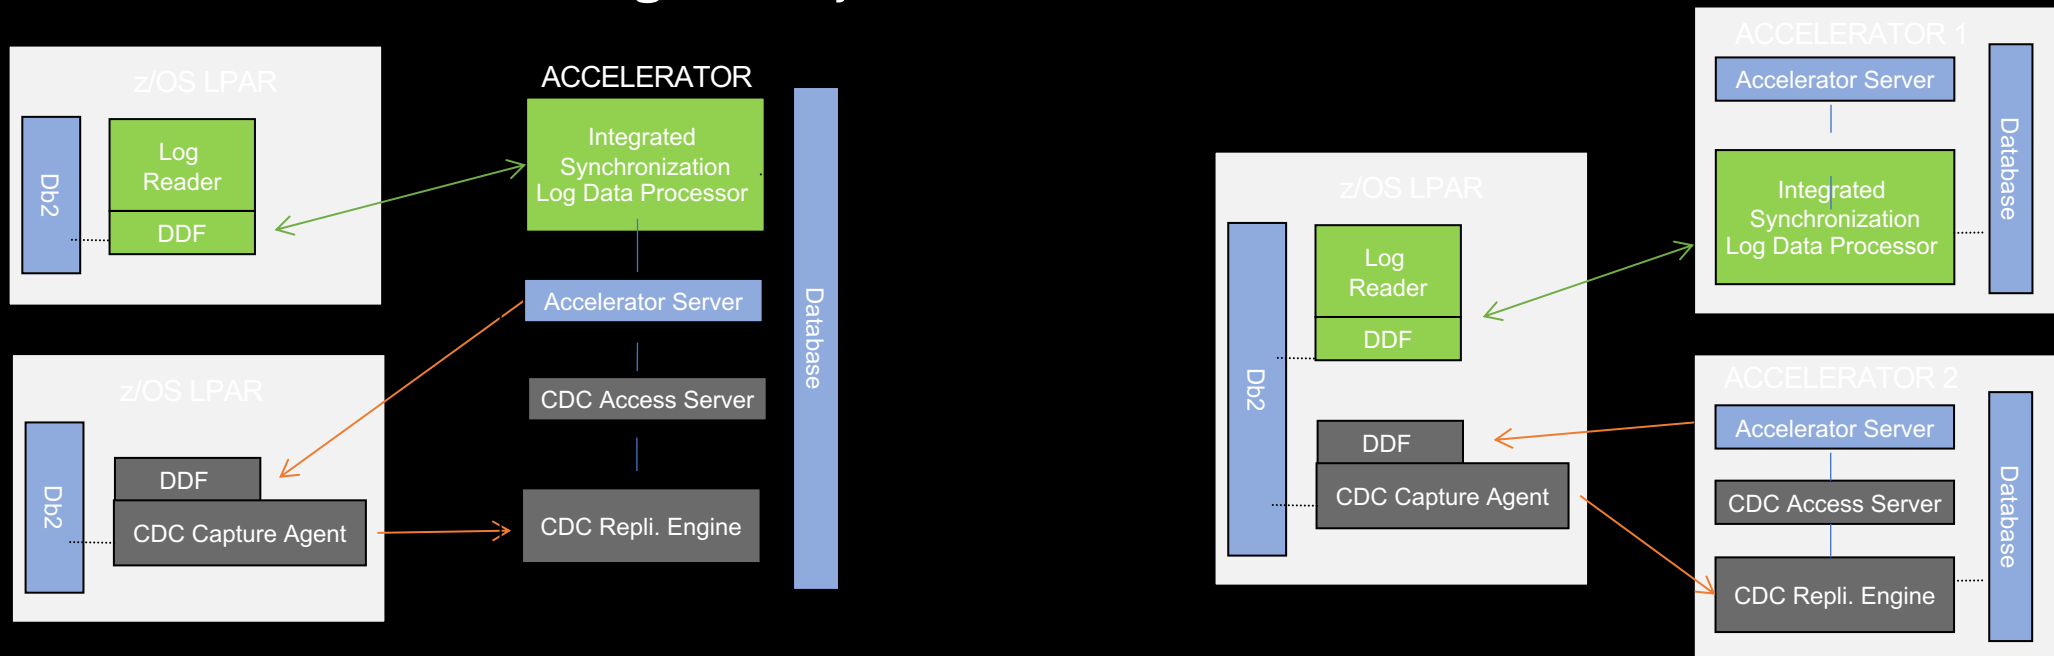
# Performance Summary\*

- CPU Usage
  - **Up to 49% CPU time reduction** for log reading vs. CDC
  - **InSync log reading is 99% zIIP eligible** (counted in DIST address space)
  - All this **despite massive filtering** of “out-of-scope” log records for test!
- Latency and Throughput
  - **Significant reduction in replication latency** and **big improvement in replication throughput**
    - CDC: Latency increased during ~1 hour test timeframe up to 20 min.
    - InSync: Maximum latency during ~1 hour test timeframe of 10 sec
- HTAP
  - With InSync, **HTAP query scalability is significantly improved.**
    - No more impact on replication throughput and latency
  - First lab measurements **show significant reduction in query wait times**
    - Average wait time dropped from 51 seconds to 6 seconds across a total of 89 queries
    - Average latency of the workload dropped from 50 seconds to 3
  - Expect WAITFORDATA to be applicable to a much broader set of use-cases with InSync

\* Results of first internal tests with pre-populated logs scenario

# Mixed CDC and Integrated Synchronization environments

- For each Db2 connection to an Accelerator (“pairing”) you can decide whether to enable Incremental Update based on CDC or based on Integrated Synchronization (using Configuration Console)
  - One Accelerator supports both Incremental Update methods
  - From the same Db2 subsystem data can be replicated to one Accelerator with CDC and to another one with Integrated Synchronization



# Prerequisites and setup

- Prereqs:
  - Accelerator V7.1.8 or V7.1.9 + Join Early Support Program
  - Db2 12 for z/OS with APAR PH06628 PTF UI63356 installed (Db2 needs to be recycled)
  - Db2 running in function level V12R1M500
  - DDF configured with secure port access and encryption through AT-TLS
  - For data sharing: The ability to connect to a specific member, e.g. through a member-specific secure port defined with a location alias
- Setup:
  - Install Accelerator and Db2 prereqs
  - Configure DDF or data sharing group to meet the prerequisites
  - Enable Incremental Update for a Db2 subsystem using Integrated Synchronization using the Accelerator Configuration Console




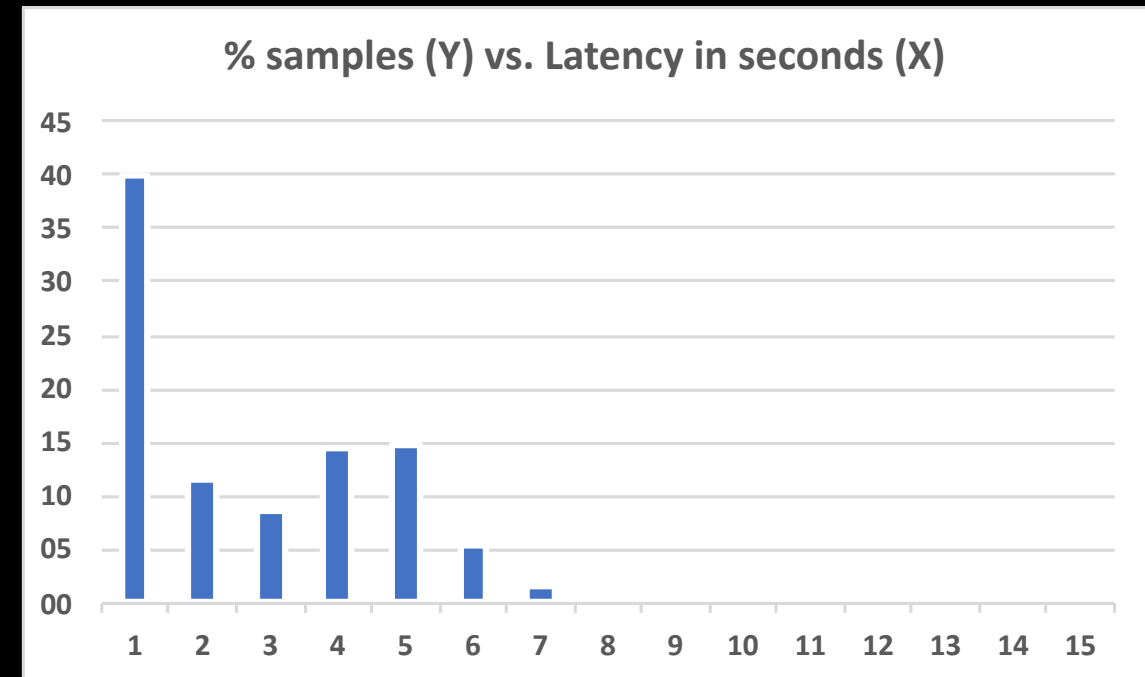
# Integrated Synchronization Technical Preview

- To validate our goals for this technology, we have started an **Early Support Program**. This will allow us to work closely with our customers and gather feedback about our objectives.
- To participate in the program, a customer must accept specific terms and conditions, and provide resources:
  - Install prerequisites
  - Allow installation of ++APAR & early code and libraries
  - Separate non-production environment for Db2 Z and the Accelerator (Do not test this in production)
  - Allow remote support (via Web-Ex) together with IBM Support
  - The customer will actively test and provide feedback about the new functionality. This feedback will be integrated into further development efforts and into the Accelerator in general.
  - Allow IBM to share testing results in presentations at conferences (Either anonymized or presented by the customer)
  - Regular status calls with customers, including education provided by IBM
- Contact: [dwhz@de.ibm.com](mailto:dwhz@de.ibm.com)

# Integrated Synchronization

## First results from Early Support Program

- Successful installation and configuration
  - Resolved initial challenges with security certificates for AT-TLS
- Successful replication of Inserts, Updates, and Deletes
  - Multiple tables with millions of rows each
  - Very low latency 
- Successful transaction rollback tests



# Agenda

- IBM Db2 Analytics Accelerator (a.k.a. IDAA) – Motivation
- What's New 2019 - V7.1
- What's Coming 2019 - V7.5
- **What's New 2019 - Documentation**

## Solution Assurance Guides for Accelerator V7

- Essential help for planning and initial installation
- **Updated:** Solution Assurance guide for **Accelerator on IIAS**
  - Now includes chapter for Call-Home and Encryption
- **New:** Solution Assurance guide for **Accelerator on Z**
  - Platform specific guide
  - Additional instructions as IBM engineer not essential on site
- Copies available via email

# Accelerator V7 High Availability Redpaper (May 2019)

- How IBM Db2 Analytics Accelerator V7 can easily integrate into or complement existing IBM Z environments for HA and DR
- Both deployment options are covered (IIAS and Z)
- For Accelerator on Z the GDPS Metro integration is covered
- Download link:
  - <http://www.redbooks.ibm.com/abstracts/redp5536.html?Open>

## Table of contents

Chapter 1. Business continuity and disaster recovery

Chapter 2. High availability and disaster recovery in IBM Z enterprise environments concepts

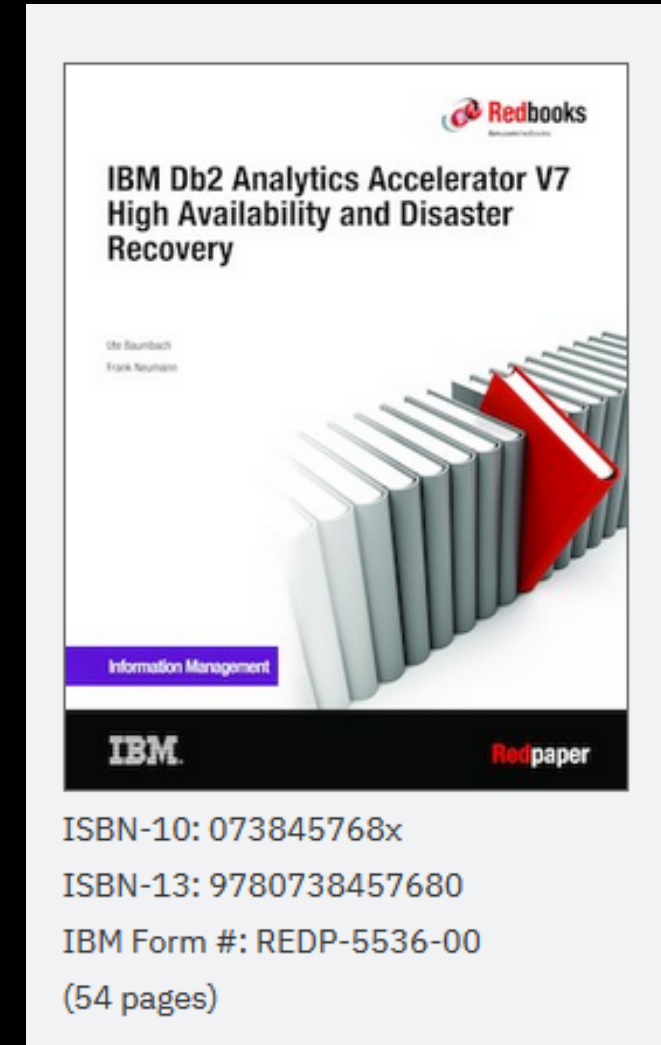
Chapter 3. Introducing high availability concepts for IBM Db2 Analytics Accelerator

Chapter 4. Combining versions and deployment platforms

Chapter 5. Planning for HA and DR with IBM Db2 Analytics Accelerator on IIAS

Chapter 6. Planning for HA and DR with IBM Db2 Analytics Accelerator on Z

Chapter 7. Summary





# Links

- Release Notes V7.1.7
  - <http://www-01.ibm.com/support/docview.wss?uid=ibm10876158>
- Release Notes V7.1.8
  - <https://www-01.ibm.com/support/docview.wss?uid=ibm10885773>
- Release Notes V7.1.9
  - <https://www.ibm.com/support/pages/node/1075299>
- V7.1 Prerequisites and maintenance
  - <https://www-01.ibm.com/support/docview.wss?uid=swg27050440>
- Knowledge Center
  - [https://www.ibm.com/support/knowledgecenter/SS4LQ8\\_7.1.0/com.ibm.datatools.agt.doc/idaa\\_kc\\_welcome.html](https://www.ibm.com/support/knowledgecenter/SS4LQ8_7.1.0/com.ibm.datatools.agt.doc/idaa_kc_welcome.html)



## Learn more

IBM Demos Site

Visit the Db2 Analytics Accelerator on IBM Demos:  
<http://ibm.biz/Acceleratordemos>

- What's available?
  - Product videos - for a quick glance at product features and capabilities
  - A guided demo - provides a guided walk-through of the Accelerator's basic capabilities
  - A hands-on lab - offers a full product experience with a guided tutorial

**New Integrated  
Synchronization  
Video  
Now Available!**

### IBM Db2 Analytics Accelerator for z/OS

#### IBM Db2 Analytics Accelerator for z/OS

IBM® Db2® Analytics Accelerator is a high-performance component tightly integrated with Db2 for z/OS®. It delivers high-speed processing for complex Db2 queries to support business-critical reporting and analytic workloads. The Accelerator transforms the mainframe into a hybrid transaction and analytic processing (HTAP) environment. It drives out cost and complexity and enables analytics on transactional data as it is generated. Leverage your business-critical data where it originates to integrate real-time insight with real-time operational decisions.

#### Products:

[IBM Db2 Analytics Accelerator for z/OS](#)

#### Videos

Watch technical experts walk you through common use cases, highlighting product features key capabilities.

#### Introduction to Db2 Analytics for z/OS

Video 0hr 7min

Introductory video explaining the features and benefits of running the Db2 Analytics Accelerator with Db2 for z/OS.

#### Connecting the Db2 Analytics Accelerator to a Db2 subsystem

Video 0hr 3min

This video outlines the easy steps to connect the Accelerator to a Db2 subsystem, define and attach an Accelerator and start/stop query acceleration.

#### True HTAP™ queries on the Db2 Analytics Accelerator

Video 0hr 7min

Detailed overview of how to setup and run True HTAP queries on the Db2 Analytics Accelerator.

# Please submit your session feedback!

- Do it online at <http://conferences.gse.org.uk/2019/feedback/ID>
- This session is ID

1. What is your conference registration number?

💡 This is the three digit number on the bottom of your delegate badge

2. Was the length of this presentation correct?

💡 1 to 4 = "Too Short" 5 = "OK" 6-9 = "Too Long"

1 2 3 4 5 6 7 8 9

3. Did this presentation meet your requirements?

💡 1 to 4 = "No" 5 = "OK" 6-9 = "Yes"

1 2 3 4 5 6 7 8 9

4. Was the session content what you expected?

💡 1 to 4 = "No" 5 = "OK" 6-9 = "Yes"

1 2 3 4 5 6 7 8 9



# GSE 2019 – Accelerate your analytics

## IBM Db2 Analytics Accelerator – An exciting 2019



**GSE session number: ID**

**Chris Watson – IBM Analytics on IBM Z Center of Excellence**

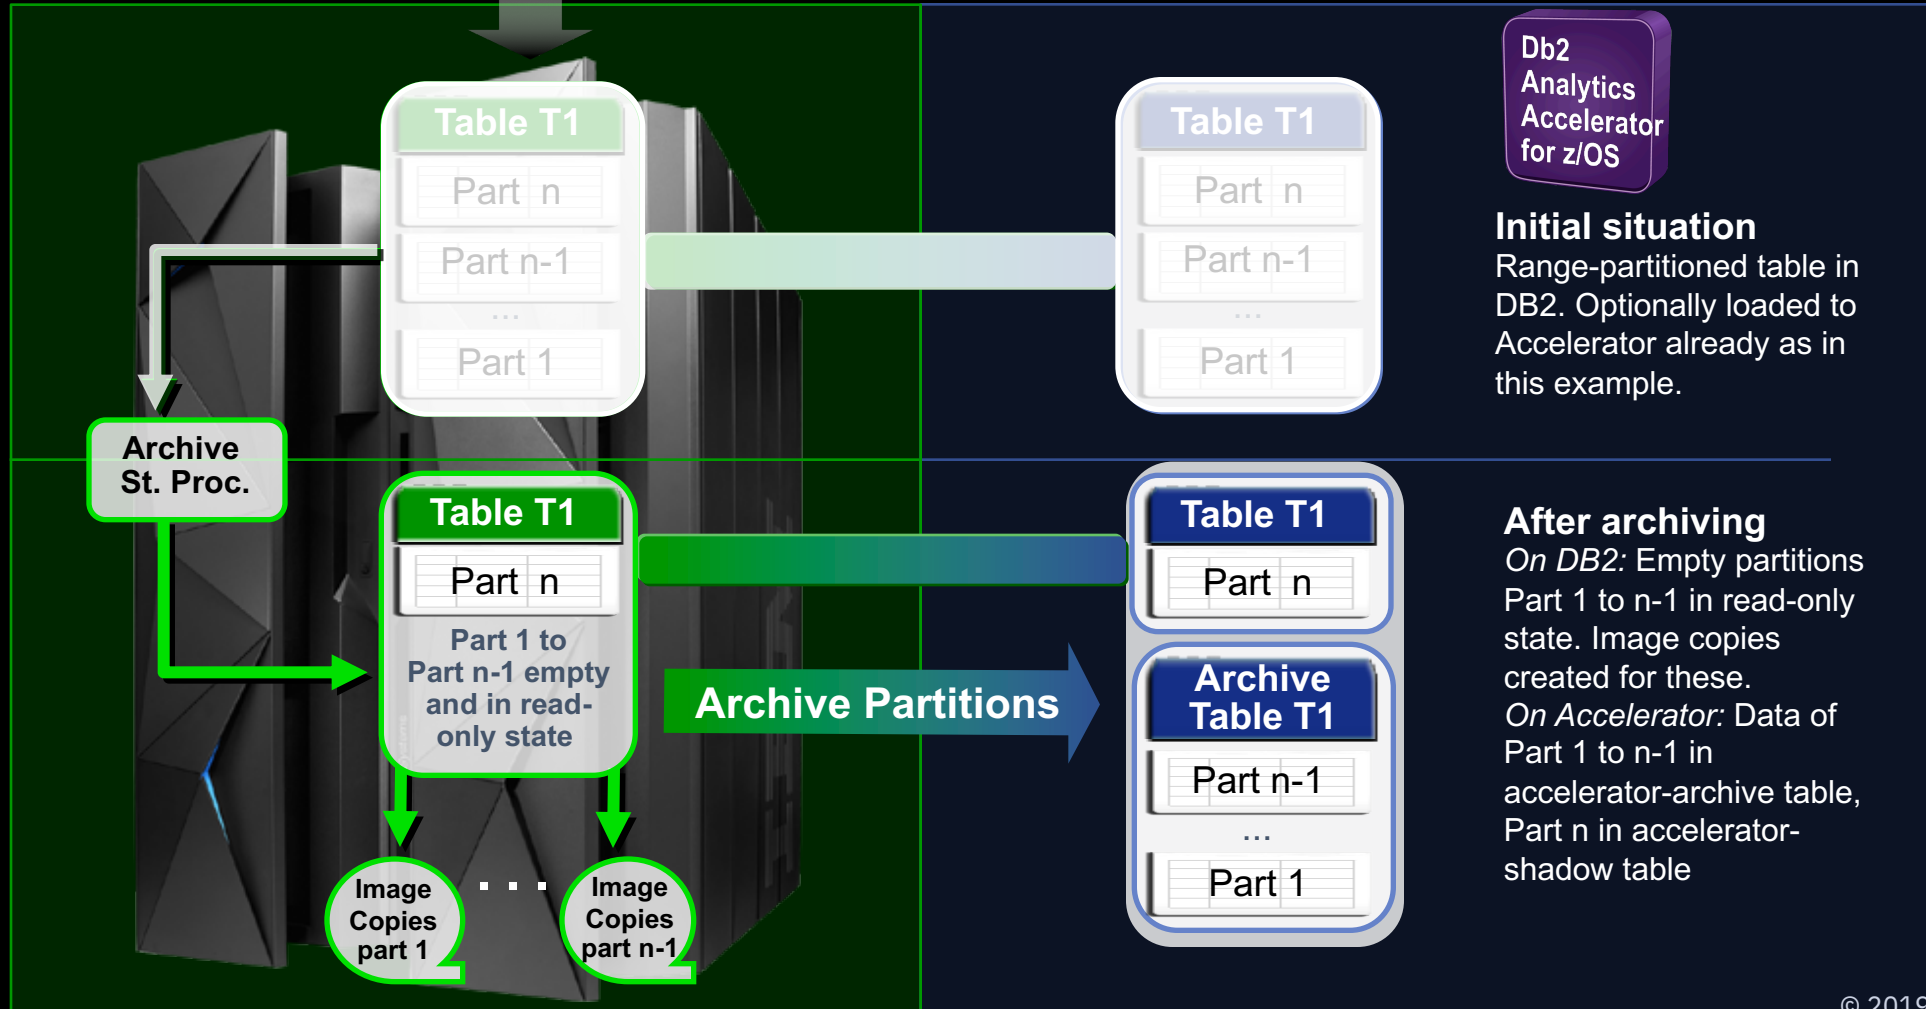
[watson@de.ibm.com](mailto:watson@de.ibm.com) (Linked In: Christopher Watson)

# Backup

# High Performance Storage Saver process



User requests to archive partitions 1 to n-1 using the ACCEL\_ARCHIVE\_TABLES stored procedure





# High Performance Storage Saver (HPSS) – GA with 7.1.8

- Supported Use Cases
  - Archiving of partitions of range-partitioned tables
  - Restoring of partitions
    - Restore uses the Db2 image copy taken during archiving
  - High availability: Archiving to multiple accelerators
    - Archive to second accelerator done using the Db2 image copy taken during first archive
  - Migration: Archiving of partitions to V7 that are already archived on V5
    - Archive to V7 accelerator done from Db2 image copy taken during archive to V5 accelerator
- Not yet available:
  - Incremental Update for active (non-archived) partitions of Accelerator-archive table