IDUG

2026

Sydney | March 16 - 18

AU Db2 TECH CONFERENCE

Application Modernization – Considered those Db2 for z/OS Capabilities?

Anthony Ciabattoni, IBM

Session Code: C07

Platform:

Application Development | Cross Platform

Agenda

- Pattern and data challenges in IBM Z application modernization
- Db2 capabilities and solution approaches for modern applications storing data in Db2 for z/OS
- Best practices in using Db2 capabilities successfully
- Questions

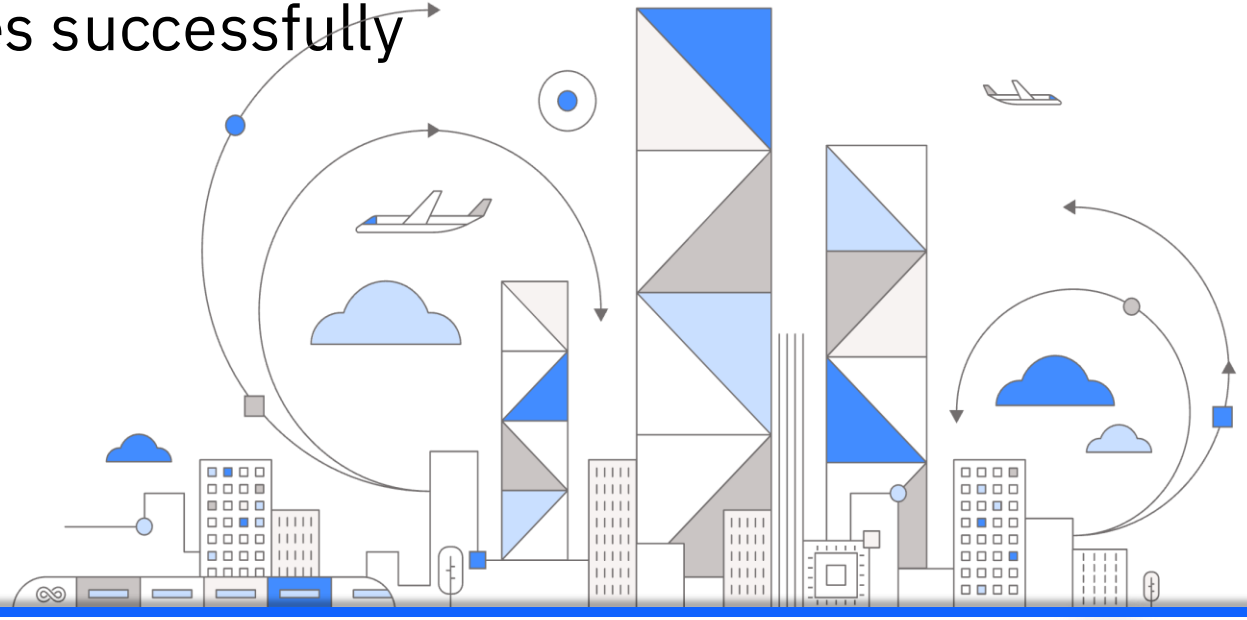

Digital transformation imposes significant demands on existing applications and data

Digital use cases drive data and functionality needs that are hosted on IBM zSystems

Account queries on Mobile using open APIs

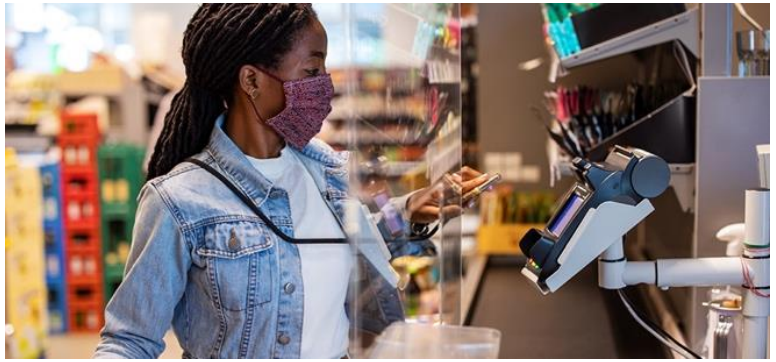

Modernize loyalty program to hybrid cloud

Raise credit card dispute

- New demands, volume, integration are driving the necessity for application and data access transformation

Operational application transformation approaches

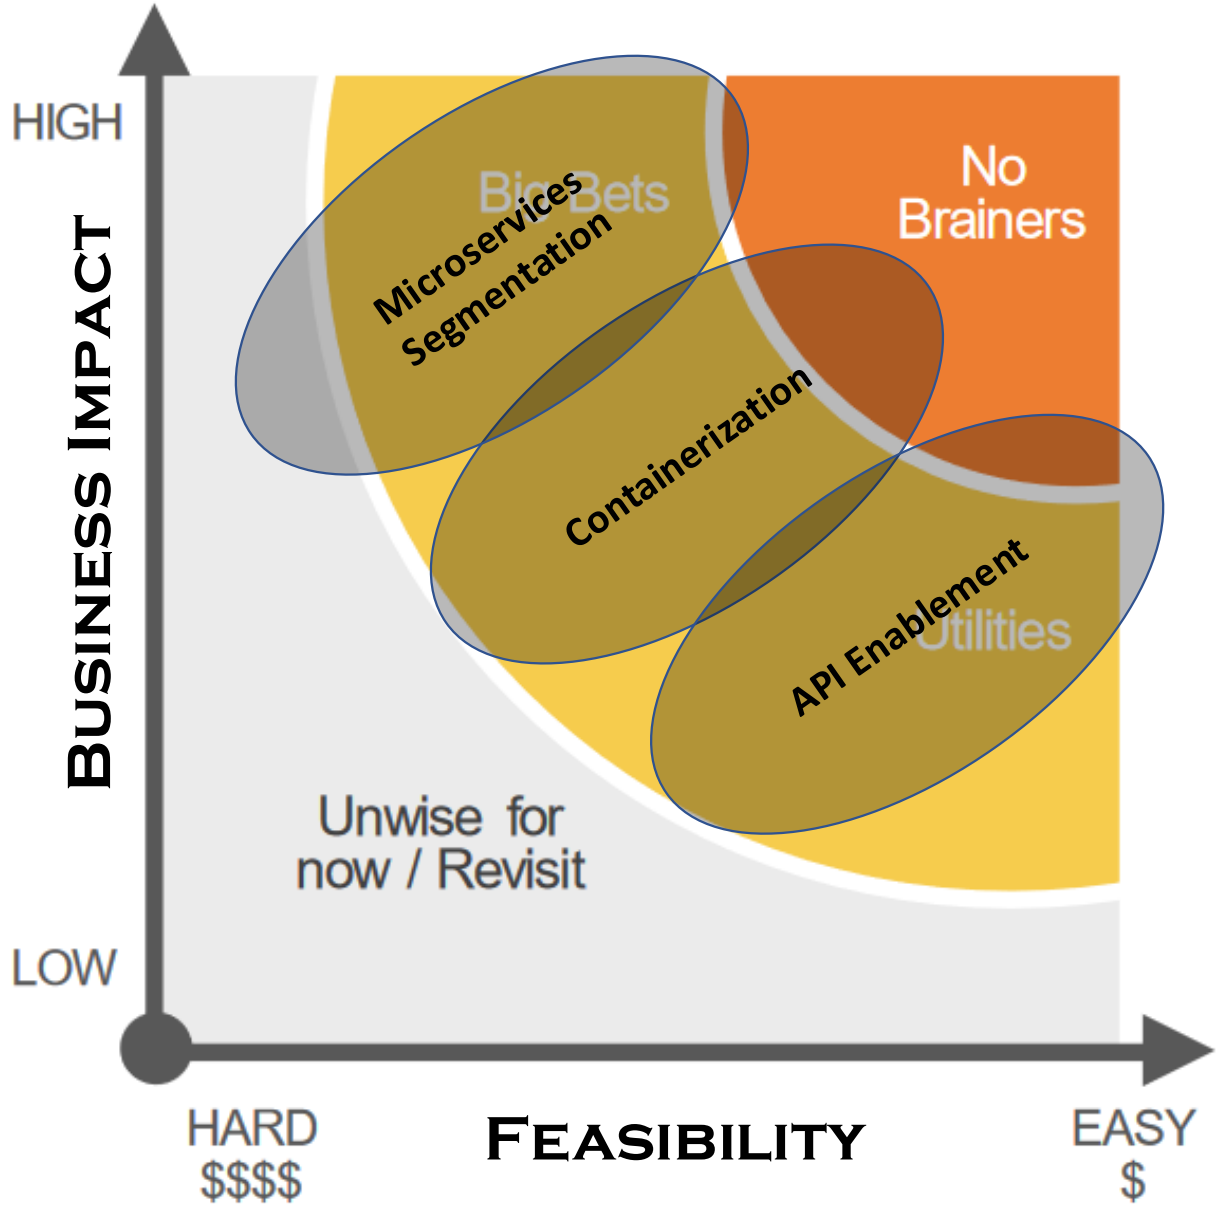

There are three major approaches to Core System Transformations. Each client's environment, end state requirements, risk tolerance, financial justification, etc will define the business impact of each of these approaches – clearly hybrid approaches and variants are viable

- ***API Enablement***: Simplest approach is for these core systems to provide a set of consumable APIs for the Systems of Engagement. The reverse is also viable where the Core System consumes APIs provided by external systems
- ***Containerization***: The systematic movement of application logic to containers that alternately can run anywhere. Preserves data and batch processes
- ***Microservices Segmentation***: Total restructuring of the Core System into a cloud native systems. This is extremely risky and is best achieved by starting with Containerization and developing a data model that can be partitioned. Eventual consistency of data is a fundamental change whose impact needs to be clearly understood

Protecting/leveraging data will drive transformation approach

Data Access vs Data Replication Considerations

- Establishing and maintaining data replication pipeline is expensive and time consuming
- Maintaining a data replication pipeline typically creates data quality and data latency challenges for consuming applications, adds reconciliation processes
- Both Microservices and containerization enable
  - Cloud native
  - Agile processes
- Accessing data in place accelerates the transformation, also improves opportunity for success
- Data access preserves the existing data management and recovery processes

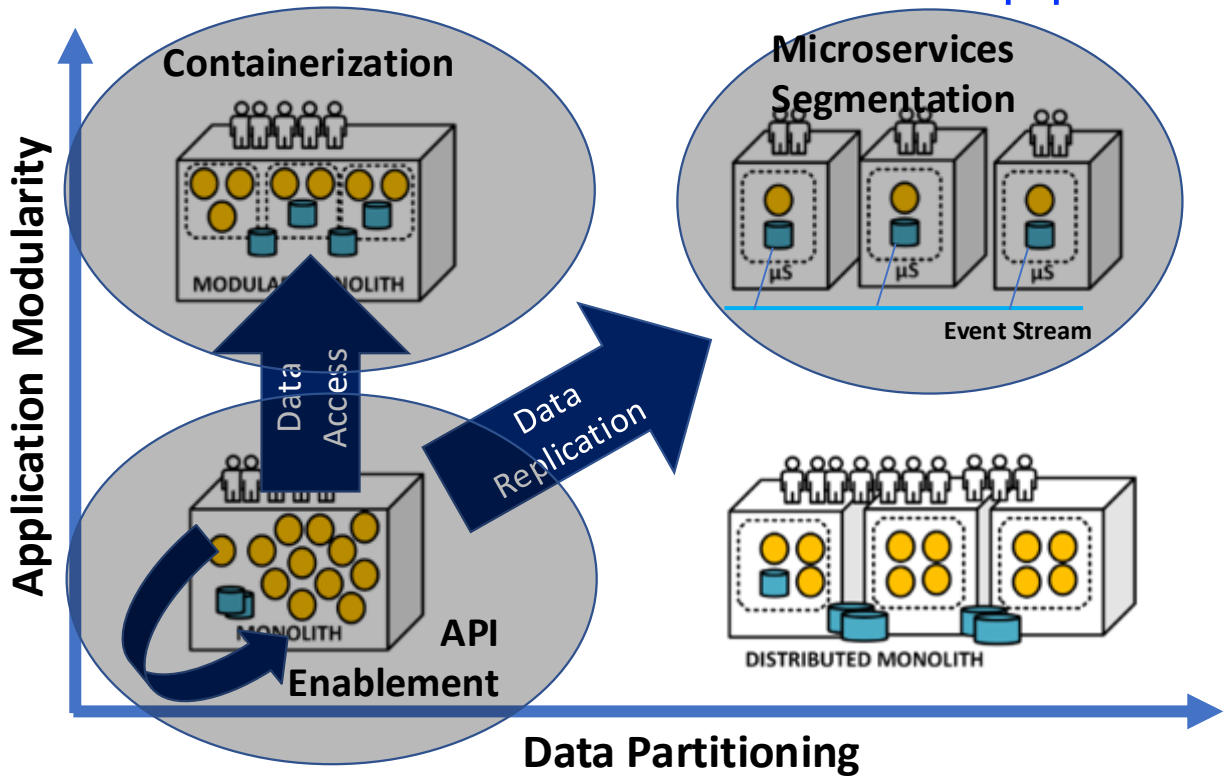

Data Consistency versus Eventual Consistency Considerations

- Data consistency reduces the amount of compensating application logic and/or business processes that need to be developed for systems that are out of sync (one clear system of record)
- Data Consistency addresses the double spend problem
- Eventual consistency enables systems to be developed independent of each other however can result in data quality issues

Incremental IBM Z data consumption modernization approach

Consider Data Fabric for IBM Z data

Communicate widely adopted incremental IBM Z data modernization approach

1. Modernize access to existing IBM Z data

2. Virtualize and govern data access across data sources and platforms

3. Satisfy complex information needs
   - Infuse AI models, historical data) within service processing

Modernize access to existing IBM Z data sources

- Leverage broad API support via SQL and REST to simplify application development
- Enable developers, DevOps engineers, and DBAs to be more productive and agile in developing and deploying applications that access Db2 for z/OS
- Readily provide transformed data via SQL using the Db2 Analytics Accelerator and Data Virtualization Manager for z/OS
- Db2 13 SQL Data Insight delivers Machine Learning via SQL
- Reduce risk and increase data quality by reducing data movement

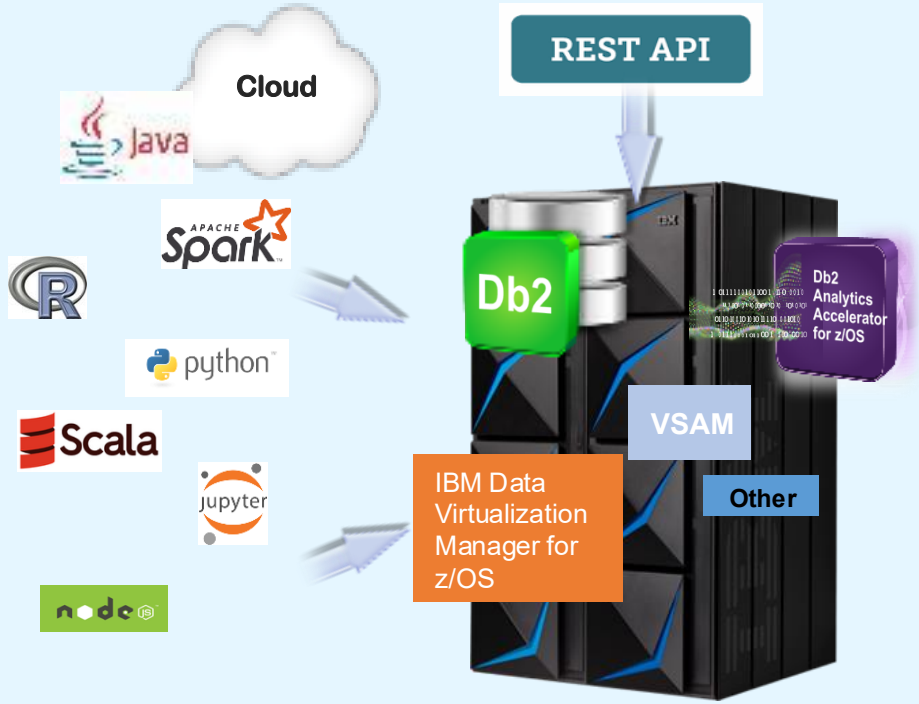

Virtualize and govern data access across data sources & platforms

=> Create Foundation for a Data Fabric

- Virtualize data assets across platform to simplify data access and governance
- Use IBM Knowledge Catalog as a central point for data governance across platform including IBM Z data
- Use IBM Manta Data Lineage
- Compose and aggregate data from multiple systems of record without data movement
- Alternatively, synchronize Z data using IBM Data Gate for watsonx for data availability in Iceberg format for AI consumption

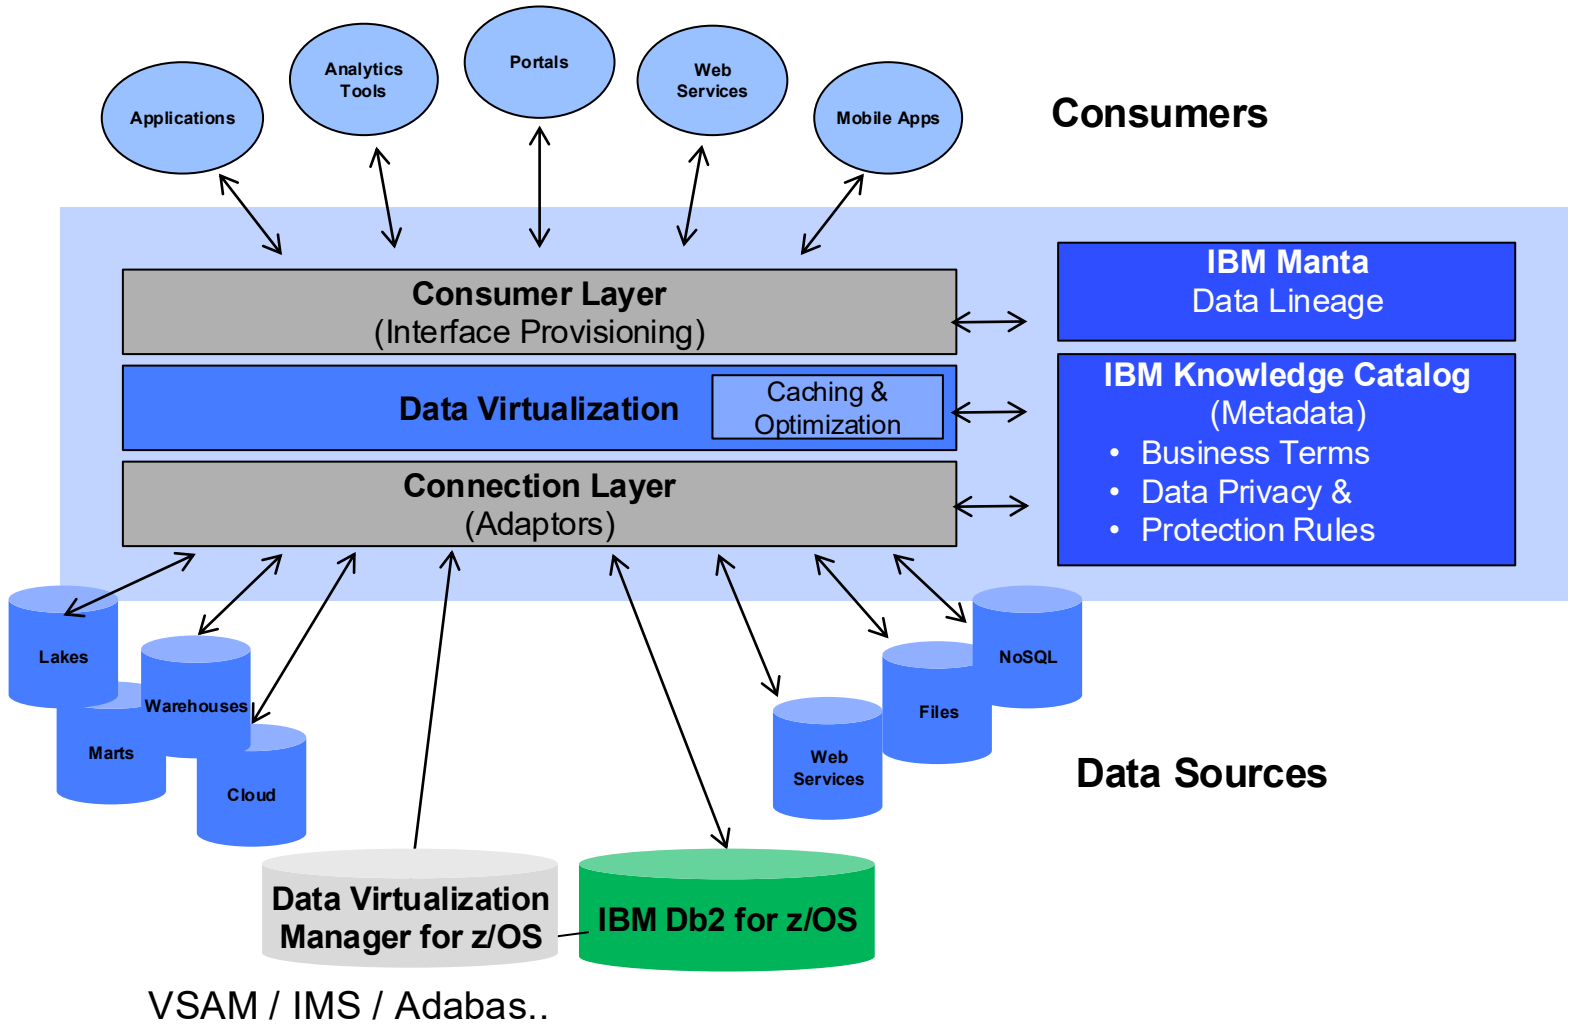

Satisfy complex information needs (e.g. infuse AI models, historical data)

=> Exploit Data Fabric within service processing

With a centralized view of data, including IBM Z data, within IBM Knowledge Catalog you can use Watson Studio to readily build, test and evaluate traditional AI models on your platform of choice

AI models, Predictive Model Markup Language (PMML)/ Open Neural Network Exchange (**ONNX**) can be deployed to Watson Machine Learning for z/OS exploiting IBM Z Telum processors to address more complex information needs within business services

Special purpose LLMs integration with business-critical processes exploiting IBM Z Spyre AI card

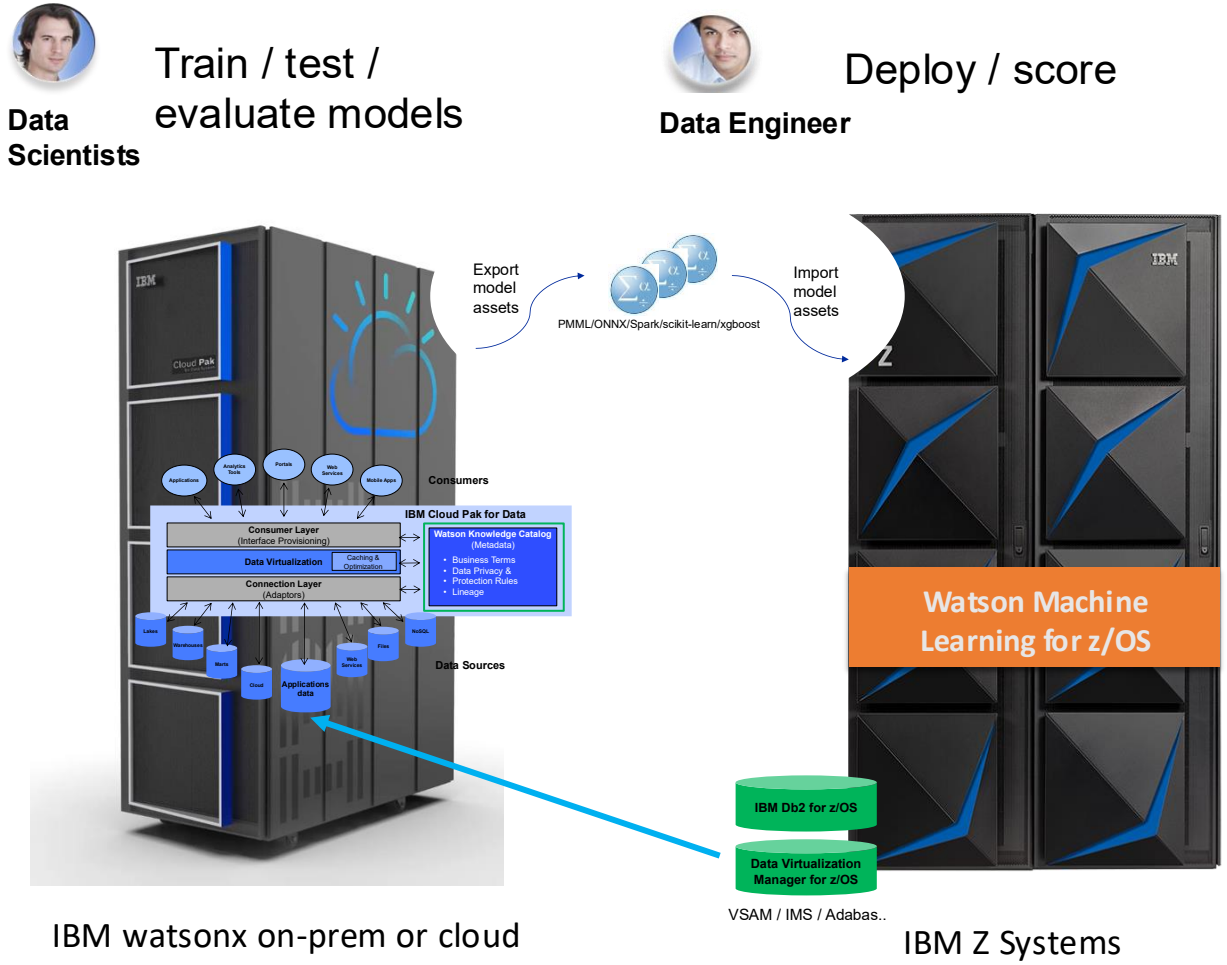

SQL Data Insights

An industry-leading relational database with embedded AI capabilities

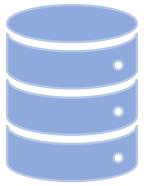

Infuse AI directly into your database on existing data to discover hidden information

Minimizes complexity of deploying AI into your applications

Single model used for a range of inferencing task over multiple fields

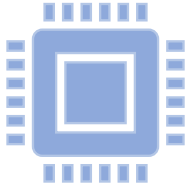

Exploits IBM z acceleration

Db2 13 SQL Data Insight
Extract greater value from Db2 for z/OS data

- No deep data science skills required
- Simply select data, enable training and Db2 for z/OS builds a data relationship model
- Build Neural Network powered relationship maps using unsupervised training over (unlabeled) structured data
- Apply relationship maps and built-in AI-related functions within any SQL statement
- Readily interpret underlying reasons for insight

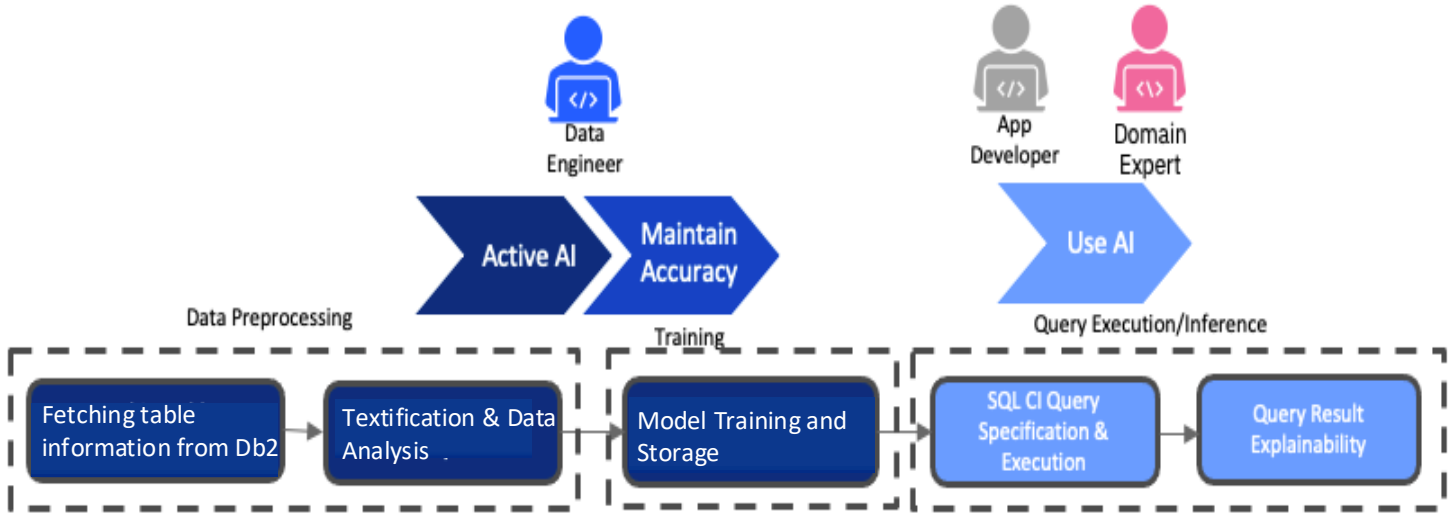

Applicable to a broad range of enterprise critical domains: Finance, Insurance, Retail, Security, HR, IT Management, Data Integration (Entity Resolution; Data Cleaning)

Db2 13 SQL Data Insight …


Tampa Bay Times

FLORIDA'S BEST NEWSPAPER tampabay.com ★★★★ Wednesday, February 2, 2022 | $2

BRADY'S FINAL BOW

Todd Bowles

HEAD COACH

"When you replace a player of that magnitude, first of all, you don't replace him"

Buccaneers replace Tom Brady with former No. 1 overall pick Baker Mayfield

Chris Cwik
Staff writer
Wed, Mar 15, 2023 · 2 min read

806

The Tampa Bay Buccaneers have their answer at quarterback now that Tom Brady has retired "for good." The team signed former No. 1 overall pick Baker Mayfield to a one-year, $8.5 million deal Wednesday, according to Tom Pelissero of NFL Network.

If Mayfield can replicate that, the Buccaneers could get a decent season out of the former No. 1 overall pick and Mayfield could extend his career. In the event Mayfield struggles, the Buccaneers can easily move on after one season.

Either way, Mayfield doesn't seem long for Tampa. Everything about the signing indicates the Bucs will use Mayfield as a stopgap solution until the team can find a more permanent option under center.

Db2 13 SQL Data Insight …

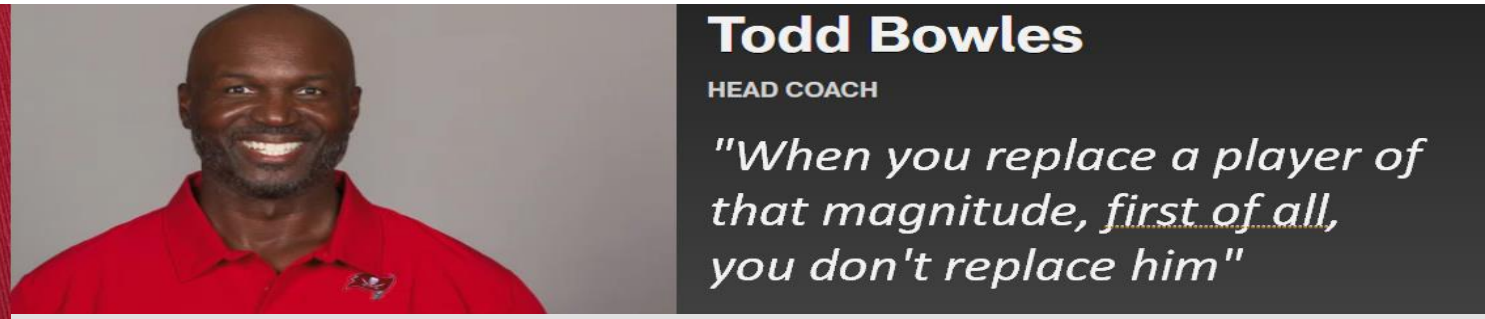

How to find a better option to replace Tom Brady?

Db2 13 SQL Data Insight …

View "NFL_PLAYERS": 7K Rows, 72 attributes (69 to compare similarities) downloadable free from Kaggle.com

*SELECT * FROM NFL_PLAYERS WHERE NAME='Tom Brady'*

ID	NAME	POSITION	TEAM	SEASON	SEASON_TYPE	COMPLETIONS	ATTEMPTS	PASSING_YARDS	PASSING_TDS	INTERCEPTIONS	SACKS	SACK_YARDS	SACK_FUMBLES	SACK_FUMBLES_LOST	PASSING_AIR_YARDS	PASSING_YARDS_AFTER_CATCH	PASSING_FIRST_DOWNS
1625	Tom Brady	QB	NE	2012	REG	401	637	4827	35	8	27	182	1	0	5508	2218	255
1625	Tom Brady	QB	NE	2013	REG	380	628	4343	25	11	40	256	6	2	5304	2014	225
1625	Tom Brady	QB	NE	2014	REG	373	582	4109	33	9	21	134	5	3	4691	1862	220
1625	Tom Brady	QB	NE	2015	REG	402	624	4770	36	7	38	225	5	2	4915	2414	228
1625	Tom Brady	QB	NE	2016	REG	291	432	3554	28	2	15	87	4	0	3353	1744	163
1625	Tom Brady	QB	NE	2017	REG	385	581	4577	32	8	35	201	6	2	5308	1910	230
1625	Tom Brady	QB	NE	2018	REG	375	570	4355	29	11	21	147	4	2	4487	2121	205
1625	Tom Brady	QB	NE	2019	REG	373	613	4057	24	8	27	185	3	1	4613	1863	193
1625	Tom Brady	QB	TB	2020	REG	401	610	4633	40	12	21	143	1	0	5529	1810	233
1625	Tom Brady	QB	TB	2021	REG	485	719	5316	43	12	22	144	3	2	5821	2534	269
1625	Tom Brady	QB	TB	2022	REG	490	733	4694	25	9	22	160	3	2	4944	2315	237

GAMES	TOUCHES	COMP_PERCENTAGE	PASS_TD_PERCENTAGE	INT_PERCENTAGEF	RUSH_TD_PERCENTAGE	YEARS_PLAYED	KEY
16	660	0.63	0.055	0.013	0.174	13	Tom Brady NE 2012
16	660	0.605	0.04	0.018	0	14	Tom Brady NE 2013
16	618	0.641	0.057	0.015	0	15	Tom Brady NE 2014
16	659	0.644	0.058	0.011	0.088	16	Tom Brady NE 2015
12	460	0.674	0.065	0.005	0	17	Tom Brady NE 2016
16	606	0.663	0.055	0.014	0	18	Tom Brady NE 2017
16	594	0.658	0.051	0.019	0.087	19	Tom Brady NE 2018
16	639	0.608	0.039	0.013	0.115	20	Tom Brady NE 2019
16	640	0.657	0.066	0.02	0.1	(null)	Tom Brady TB 2020
17	747	0.675	0.06	0.017	0.071	(null)	Tom Brady TB 2021
17	762	0.668	0.034	0.012	0.034	(null)	Tom Brady TB 2022

Db2 13 SQL Data Insight …

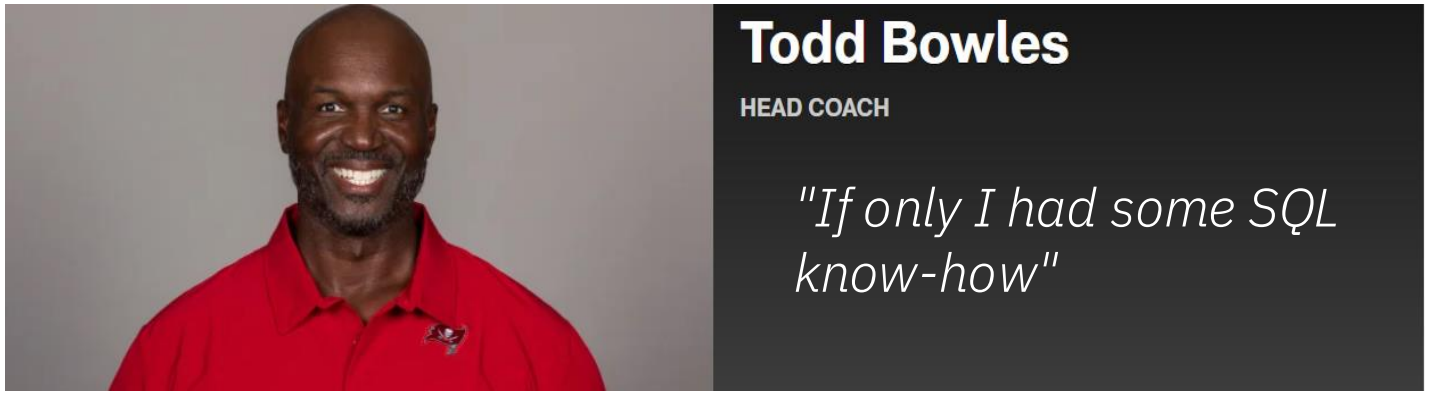


```
Select Name, Team, Position,
AI_SIMILARITY (key,'Tom Brady TB 2022')
AS SIMILARITY_SCORE
FROM NFL_PLAYERS
Where Season = 2022
and Name < > 'Tom Brady'
Order by 4 desc
```

	NAME	TEAM	POSITION	SIMILARITY_SCORE
▸	Jacoby Brissett	CLE	QB	0.599070847034454
2	Geno Smith	SEA	QB	0.564258575439453
3	Carson Wentz	WAS	QB	0.562188148498535
4	Trevor Lawrence	JAX	QB	0.560037314891815
5	Davis Mills	HOU	QB	0.552221238613129
6	Ryan Tannehill	TEN	QB	0.551405906677246
7	Aaron Rodgers	GB	QB	0.5484419465065
8	Taylor Heinicke	WAS	QB	0.539485216140747
9	Matthew Stafford	LA	QB	0.537122488021851
10	Andy Dalton	NO	QB	0.533765852451324
11	Kirk Cousins	MIN	QB	0.533228158950806
12	Lamar Jackson	BAL	QB	0.518120288848877
13	Marcus Mariota	ATL	QB	0.499959647655487
14	Matt Ryan	IND	QB	0.497775435447693
15	Joe Burrow	CIN	QB	0.49701988697052
16	Jimmy Garoppolo	SF	QB	0.494774341583252
17	Jared Goff	DET	QB	0.492530763149261
18	Daniel Jones	NYG	QB	0.474714756011963
19	Tua Tagovailoa	MIA	QB	0.473720908164978
20	Kyler Murray	ARI	QB	0.471944689750671
21	Kenny Pickett	PIT	QB	0.437743186950684
22	Mac Jones	NE	QB	0.436876475811005
23	Russell Wilson	DEN	QB	0.431953370571136
24	Josh Allen	BUF	QB	0.405514121055603
25	Justin Fields	CHI	QB	0.397185266017914
26	Dak Prescott	DAL	QB	0.392007827758789
27	Stefon Diggs	BUF	WR	0.387721300125122
28	Baker Mayfield	LA	QB	0.377517938613892
29	Jarrett Stidham	LV	QB	0.376355230808258
30	Mitchell Trubisky	PIT	QB	0.374222874641418

Db2 Customers Loving Db2 13 SQL Data Insight

"SQL DI gives us the capability to embed AI into business applications in real time."

"With the Db2 engine's ability to train AI models without the need for a Data Scientist businesses can refine and retrain these models as frequently as needed"

"This feature is like Db2 Analytics Accelerator on IBM Z – a real game changer."

Director Global Database Services, Citi

"We saw an 94% accuracy in prediction results. These very promising results have motivated us to integrate this technology into our underwriting processes in the near future."

Citi

- Uses Db2 SQL Data Insights to catch rewards program fraud
- Able to use it to easily learn about new fraud patterns

Swiss Mobiliar

- Uses SQL DI to predict the probibility an insurance marketing offer will be accepted by a potential client
- Second use case: Predict the total value of a household inventory based on number of rooms, age, location, own vs. lease, and number of persons **(implemented this use case in 4 hours in production!)**

Db2 Native RESTful Service Support …

Alternative distributed connectivity to Db2,

- No need for Db2 drivers, Db2 Connect gateway or Db2 Connect license
- Service call direct or via z/OS Connect

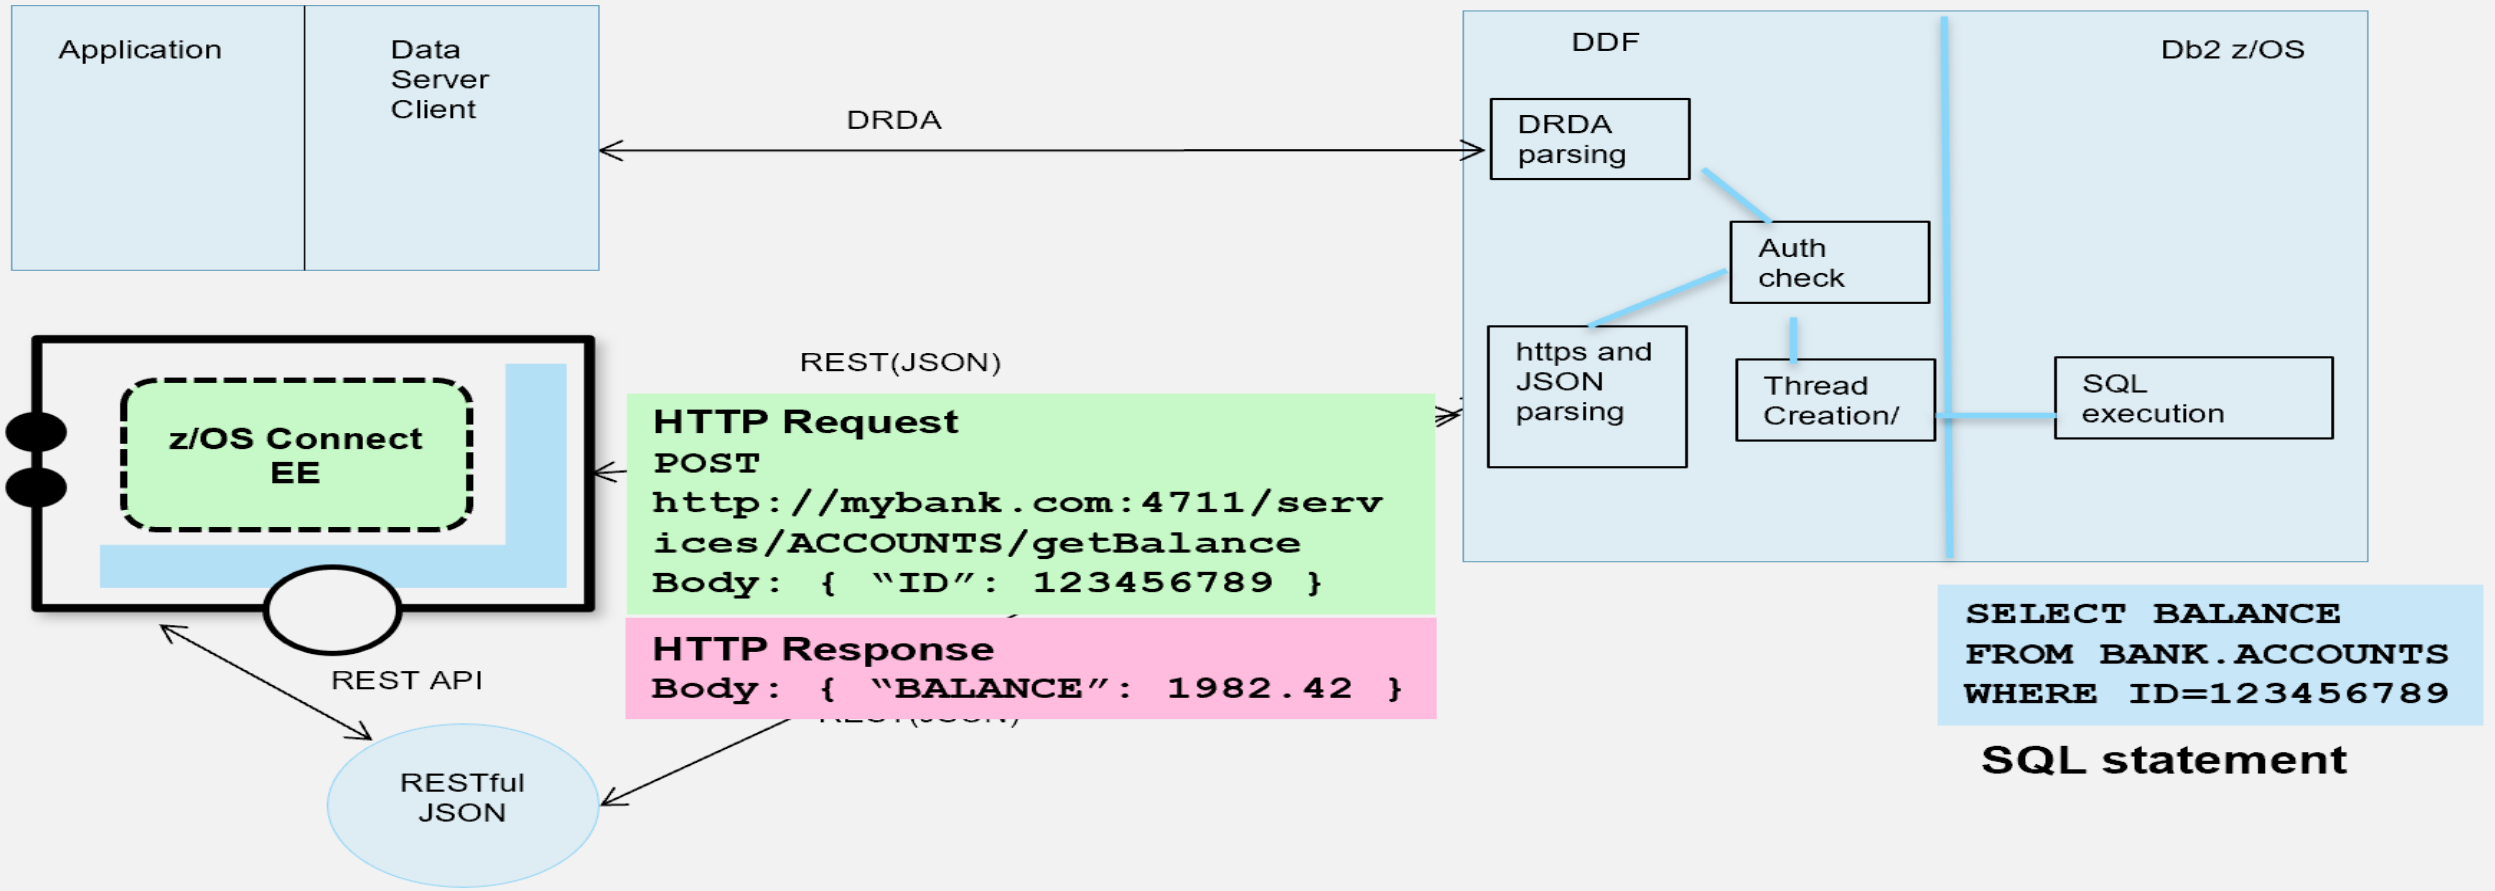

Db2 Temporal - *SYSTEM_TIME*

System-period Temporal Table (STT) contains SYSTEM_TIME period

- Tracks when the data is stored in the table for auditing purposes
- SYSTEM_TIME begin column can be defined as:

  TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS ROW BEGIN

  TIMESTAMP(12) WITH TIME ZONE NOT NULL GENERATED ALWAYS AS ROW BEGIN

- SYSTEM_TIME end column can be defined as:

  TIMESTAMP(12) NOT NULL GENERATED ALWAYS AS ROW END

  TIMESTAMP(12) WITH TIME ZONE NOT NULL GENERATED ALWAYS AS ROW END

- Two table approach: System-period Temporal Table must have a history table
- Inclusive-exclusive model:
  - Beginning value is included in the period, and the end value is NOT included in the period
- Query with FOR SYSTEM_TIME period specification (explicit, base table reference)

Db2 Temporal - *BUSINESS_TIME*

Application-period Temporal Table (ATT) contains BUSINESS_TIME period

- An application-period temporal table (often referred to as a **business timetable**) is a database table that manages data based on when it is considered valid or "in effect" in the real world
- BUSINESS_TIME begin column can be defined as:

  TIMESTAMP(6) NOT NULL ***DATE*** *NOT NULL*
- BUSINESS_TIME end column can be defined as:

  TIMESTAMP(6) NOT NULL ***DATE*** *NOT NULL*

- Inclusive-exclusive model [...): beginning value is included in the period, and the end value is NOT included in the period
- Inclusive-inclusive model: beginning and end value included in period (business time only)
- Implicit check constraint generated by DB2: end_column > begin_column
- New index clause BUSINESS_TIME WITHOUT OVERLAPS: enforce index keys being unique over a period of time
- Temporal UPDATE/DELETE with FOR PORTION OF BUSINESS_TIME period clause (explicit, base

Application modernization is an opportunity to create a solution for challenges of inactive data in Db2 tables

- A large percentage of data in a table can be inactive
  - Results in (very) large tables
- Inactive data tends to be accessed less and less over time
  - One way to address: move inactive data to separate archive
    - Typically, not popular with business because of data accessibility from business functions
- Data maintenance (Image Copy, Recovery, REORG, RUNSTATs) requirements for active, frequently updated data is different than for inactive, static data
  - Inactive data can severely slow down recovery times during unplanned outages
- Large percentage of inactive, seldom accessed data impacts performance accessing active data
  - Lower BP hit ratio, larger index tree

Db2 Archive Transparency

Separate inactive data while keeping access from existing application

Archive table contains rows deleted from archive-enabled table

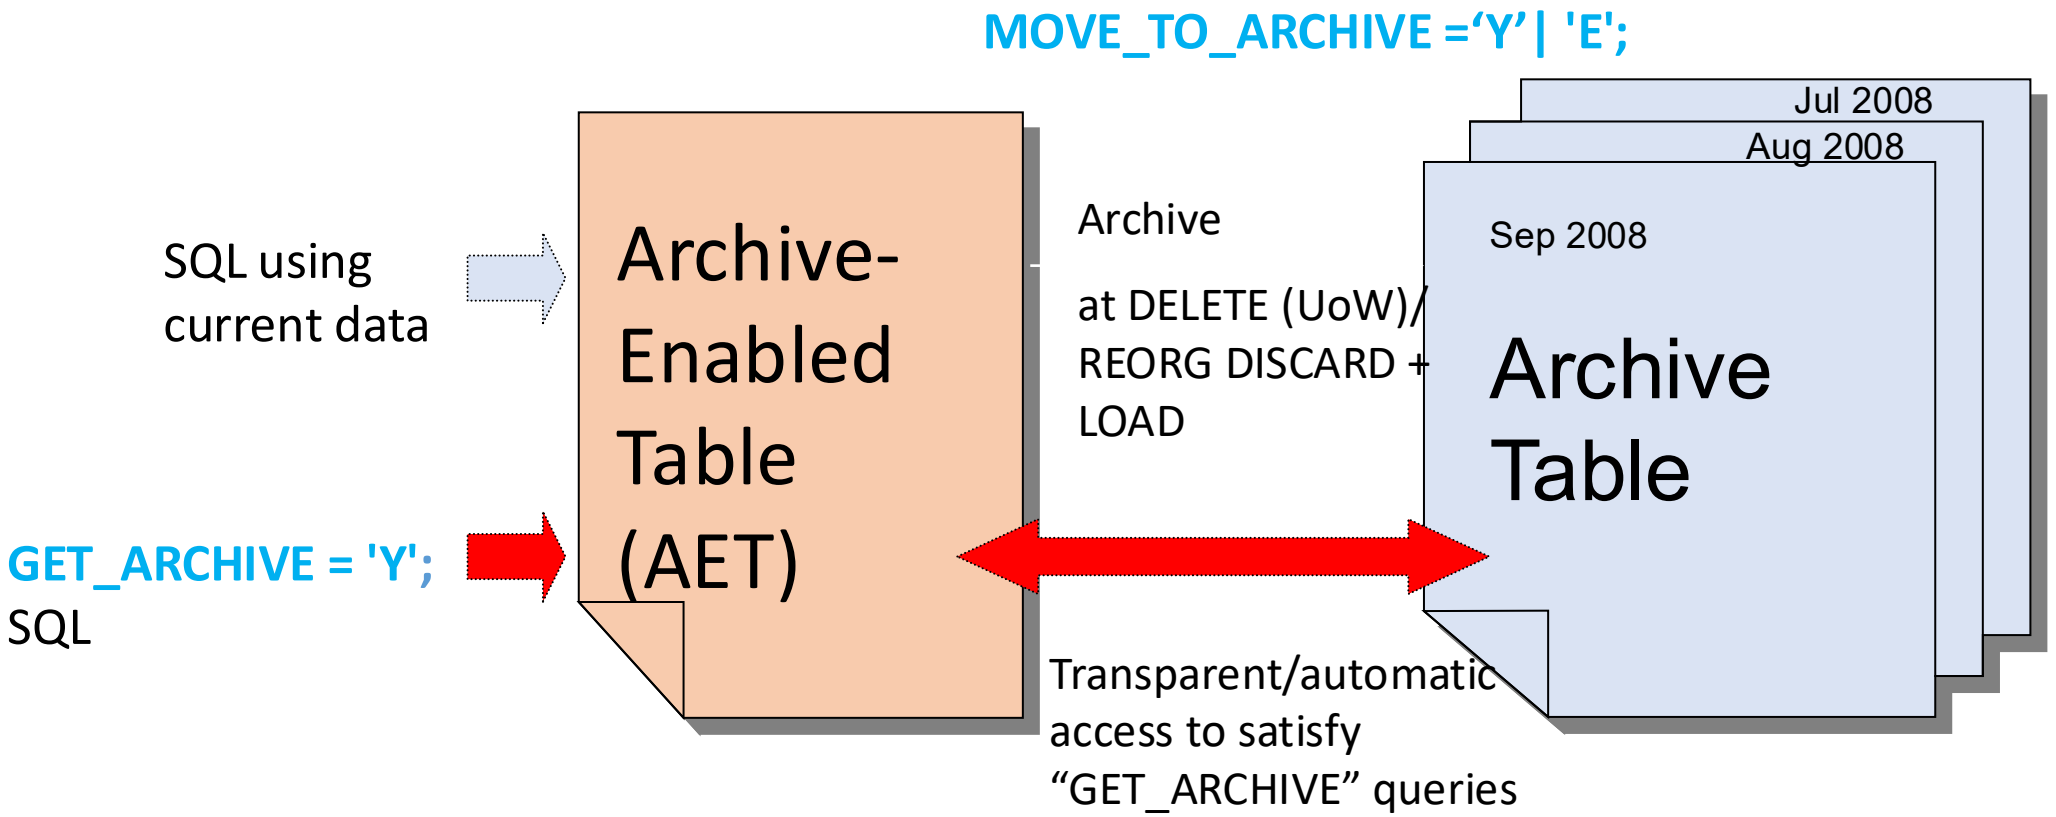

ZPARM MOVE_TO_ARCHIVE_DEFAULT in macro DSN6SPRM

Gap (open requirement) Archive via REORG/DISCARD in single unit of work

Summary

- Application modernization is a challenge and an opportunity
- Incremental modernization delivers value to the business fast with manageable risk
- Learn about pattern in your Db2 Data using Machine Learning delivered with Db2
- Efficient load of data from distributed platforms into Db2 with DRDA Fast Load
- Manage historical data needs with Db2 bi-temporal functionality
  - Also, best performing approach for historical data in Data Warehouse and Data Lakehouse
- Manage inactive data with Db2 archive transparency

IDUG

2026

Sydney | March 16 - 18

AU Db2 TECH CONFERENCE

Application Modernization – Considered those Db2 for z/OS Capabilities?

Anthony Ciabattoni, ACiabattoni@ibm.com

Session Code: C07

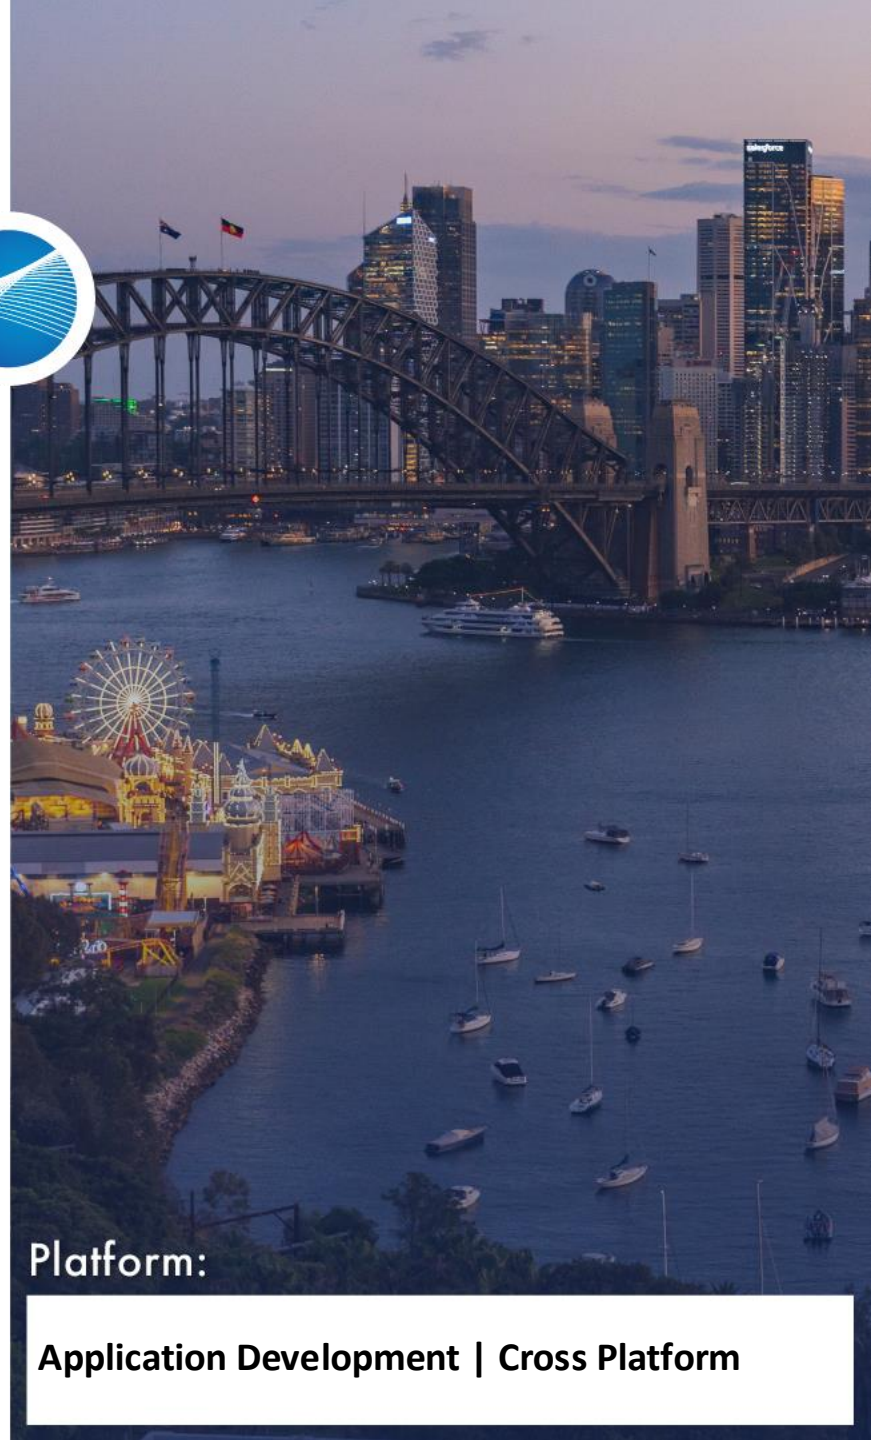

Platform:

Application Development | Cross Platform

Questions


Thank You


IDUG
2026 Australia Db2 Tech Conference