

HTAP Roadmap and Strategy



IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion. 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. The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract. The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Database Workloads

There are two broad types of database workloads today

- Online transaction processing (OLTP)
 - Example: Workloads driven by point of sales systems, banking systems, billing systems, human resources systems, manufacturing systems, etc.
- Analytic processing systems
 - Example: Traditional Data Warehouse or Data Mart analytics where complex queries process vast amounts of data to answer business questions

The data that resides in analytic processing systems usually comes from the OLTP systems

- Example a retail store will process the sales transactions
 - To record sales, ship products to customers, manage returns, etc.
- That data is then extracted from the OLTP system and stored in a Data Warehouse or Data Mart for analysis
 - Looking for trends in buyer behavior, next best offers, merchandise optimization, etc.

What is Different About The Two Workloads

Transaction Processing

- Typically 100s or 1000s of simultaneous transactions
- Relatively simple in terms of the logic the transaction is processing
 - A transaction may be made up of many steps or statements but overall each step is fairly simple
- No need for parallelism of an individual statement due to the simplicity of statements
- Goal of the DBMS is to maximize throughput (the number of transactions that can be processed in a given unit of work)

Analytic Processing

- Typically smaller numbers of users or queries running simultaneously
- Relatively complex statements
 - Often a mix of simple lookup type queries mixed in with very complex multi table joins aggregating information over vast amounts of data
- High degree of parallelism is often required to divide and conquer
- Goal of the DBMS is to minimize response time of the complex query
- Operational Analytic Processing (aka Operational Datastore (ODS) processing)
 - A subclass of analytic processing characterized by a mix of complex analytic statements and a high volume of concurrent simple statements (eg. lookup queries accessing only a few rows)

What is HTAP?

 When you combine OLTP and Analytics into one system = HTAP (Hybrid Transactional and Analytics Processing system)

There is a broad spectrum of application types when discussing HTAP

- Characteristics often associated with HTAP systems include:
 - Zero (or near-zero) latency
 - · Analytical queries have access to the latest committed transactional data
 - In-memory exploitation
 - Industry definitions of HTAP often include strong in-memory processing, particularly for speeding up analytical queries

Reporting on top of OTLP systems

- This is the most generic and often not considered true HTAP
- Virtually all OLTP systems have some reporting that is run on them but usually these reports are run as nightly batch jobs so that they do not impact the daily transaction performance

• Two more challenging types of HTAP are described on the next slide

What is HTAP - continued?

Point of Decision HTAP

- In this type of application transactions are running on the data and analytics are running on the data at the same time (but <u>not</u> in the same unit of work)
 - Example: Someone withdraws a significant amount of savings from their bank account while at the same time a loan application is being processed for that person to check their past history of transactions to check their credit score
 - The withdrawal transaction is happening in the same database as the credit check analytics but they are not part of the a single unit of work

In-process HTAP (the most challenging HTAP category)

- In this type of application a transaction and an analytics query run within the <u>same</u> unit of work
 - Example: A trade is made on the stock exchange and <u>within the same</u> unit of work a query checks the person's risk position
 - If that risk is too great, the trade will roll back (undo) the purchase of that stock

What is Available in HTAP Today for DB2

 DB2 has features to deliver reporting on OLTP as well as in-process HTAP and point of decision HTAP





 Exploiting Single Table BLU In-Memory for both OLTP and Analytics for a superior HTAP solution





 BLU Columnar Engine subsumes HTAP workloads for operational analytics







 Exploiting Single Table BLU In-Memory for both OLTP and Analytics for a superior HTAP solution



Competitive Comparison

• All vendors use a row store in front of columnar for HTAP

- DB2 Shadow Tables (full row based table and full column based table)
- Oracle In-Memory Database (full row based table plus a row buffer in front of full column based in-memory only table)
- SQL Server (full row based table plus a row buffer in front of columnar index)
- SAP HANA ("not talked about" row buffer in front of full column based table)

BLU for OLTP – Column only store

- Insert/Update/Delete ALL run directly on the columnar data as do the analytic queries without any row buffer in front
- The only vendor to not require a row buffer or row store in front of columnar
 - No need for merging data into row and then again into columnar = fewer CPU cycles and less memory required for OLTP
 - No need to compensate or route complex queries between columnar and row buffers for analytic workloads

The Future of HTAP

- Today HTAP systems are focused on "traditional transactions"
- The future "transaction" already looks very different
 - Millions of transactions per second per node
 - Streaming data in and analytics on all data (not just streaming analytics)
 - Volume of streams and spark with analytics of BLU = BLU Spark
 - STAP = Streaming volume of transactions with analytic processing



	V BAR	