Database Management Systems

Trends and Directions

Namik Hrle
IBM Fellow
CTO Private Cloud and z Analytics

March 2017

IBM

Please note:

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice and 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.
- Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.
- IBM, the IBM logo, ibm.com are trademarks of International Business Machines Corporation, registered in many jurisdictions worldwide. Other product and service names might be trademarks of IBM or other companies. A current list of IBM trademarks is available on the Web at www.ibm.com/ legal/copytrade.shtml

IT transforms businesses like never before

Taxi company owns no vehicles

Accommodations company owns no real estate

Media company creates no content

Retail company caries no inventory

Demography of one

Right person, right place, right time, right offer

Democratization of IT

Cognitive computing

Thinking-like ability

Illuminating dark data

Augmented Intelligence

Implications to DBMS technology

Coexistence with ‘new’ technologies
Spark, Hadoop, Key-Value stores, Graph databases, …

Hybrid Transactional/Analytical Processing
Bringing analytics to transactional data

Hybrid cloud delivery
Fast deployment, continuous delivery, uniform experience …

Higher standards for traditional quality of services
Performance, scalability, continuous availability, security, …

Implications to DBMS technology

Coexistence with ‘new’ technologies
Spark, Hadoop, Key-Value stores, Graph databases, …

Hybrid Transactional/Analytical Processing
Bringing analytics to transactional data

Hybrid cloud delivery
Fast deployment, continuous delivery, uniform experience …

Higher standards for traditional quality of services
Performance, scalability, continuous availability, security, …

Where Does Performance Improvement Come From?

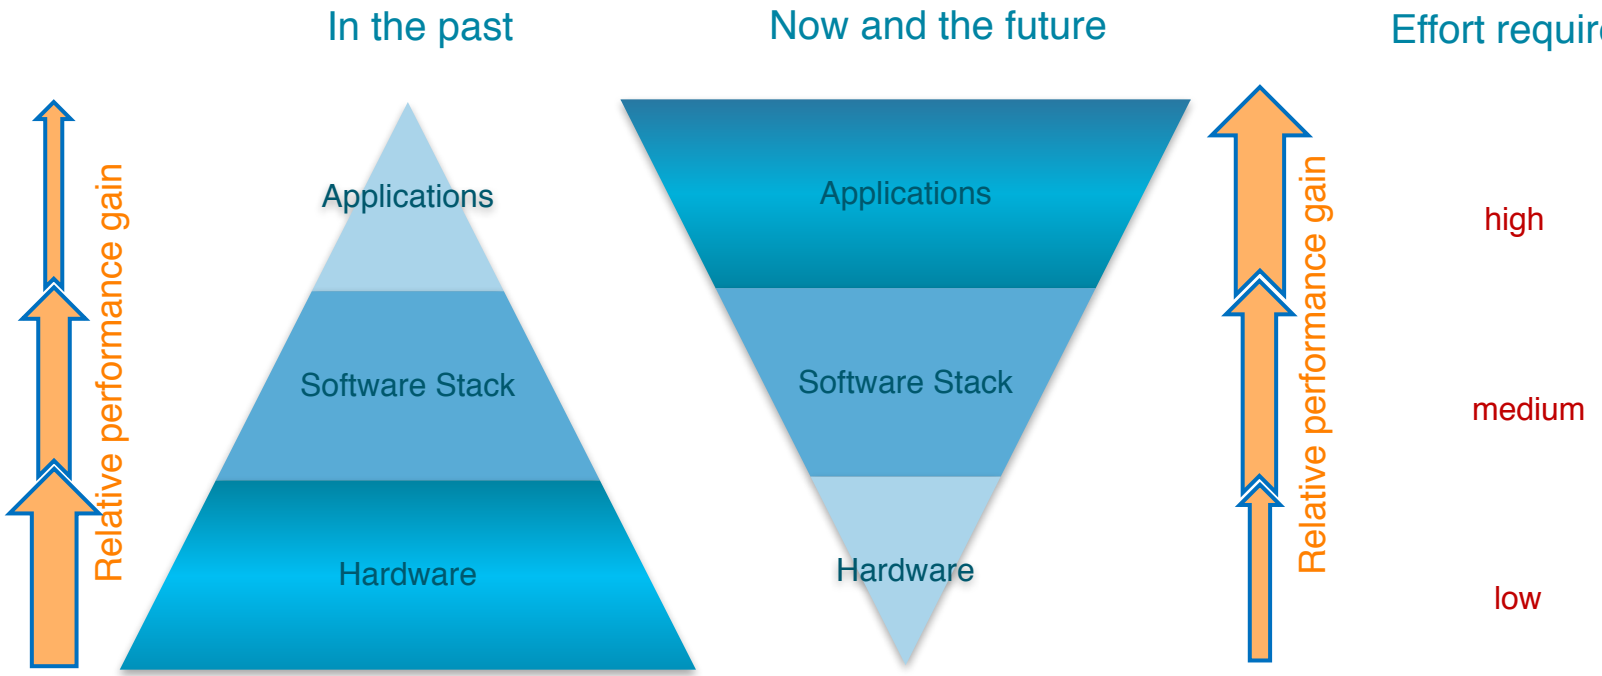

Fading of Moore's Law: Small fraction of performance improvement will come from technology scaling and transparent hardware features.

Getting harder: Bulk of performance improvements will need to come from software innovation and software exploitation of new hardware features.

Fit-for-purpose example

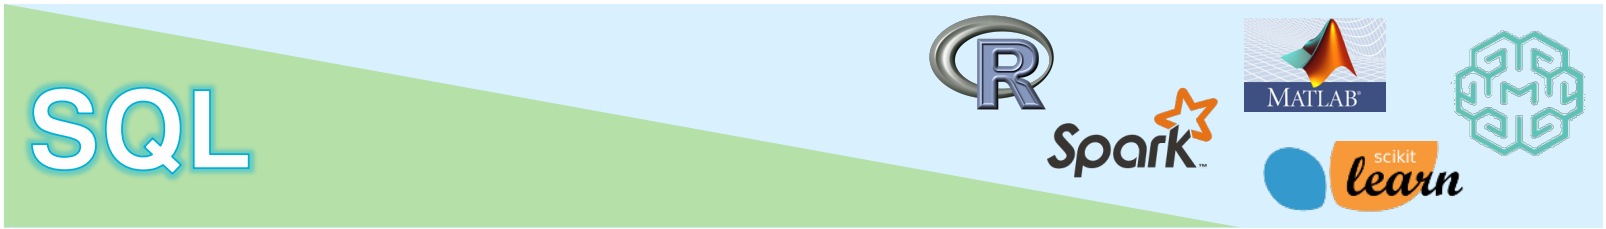

dashDB Spark Integration Examples

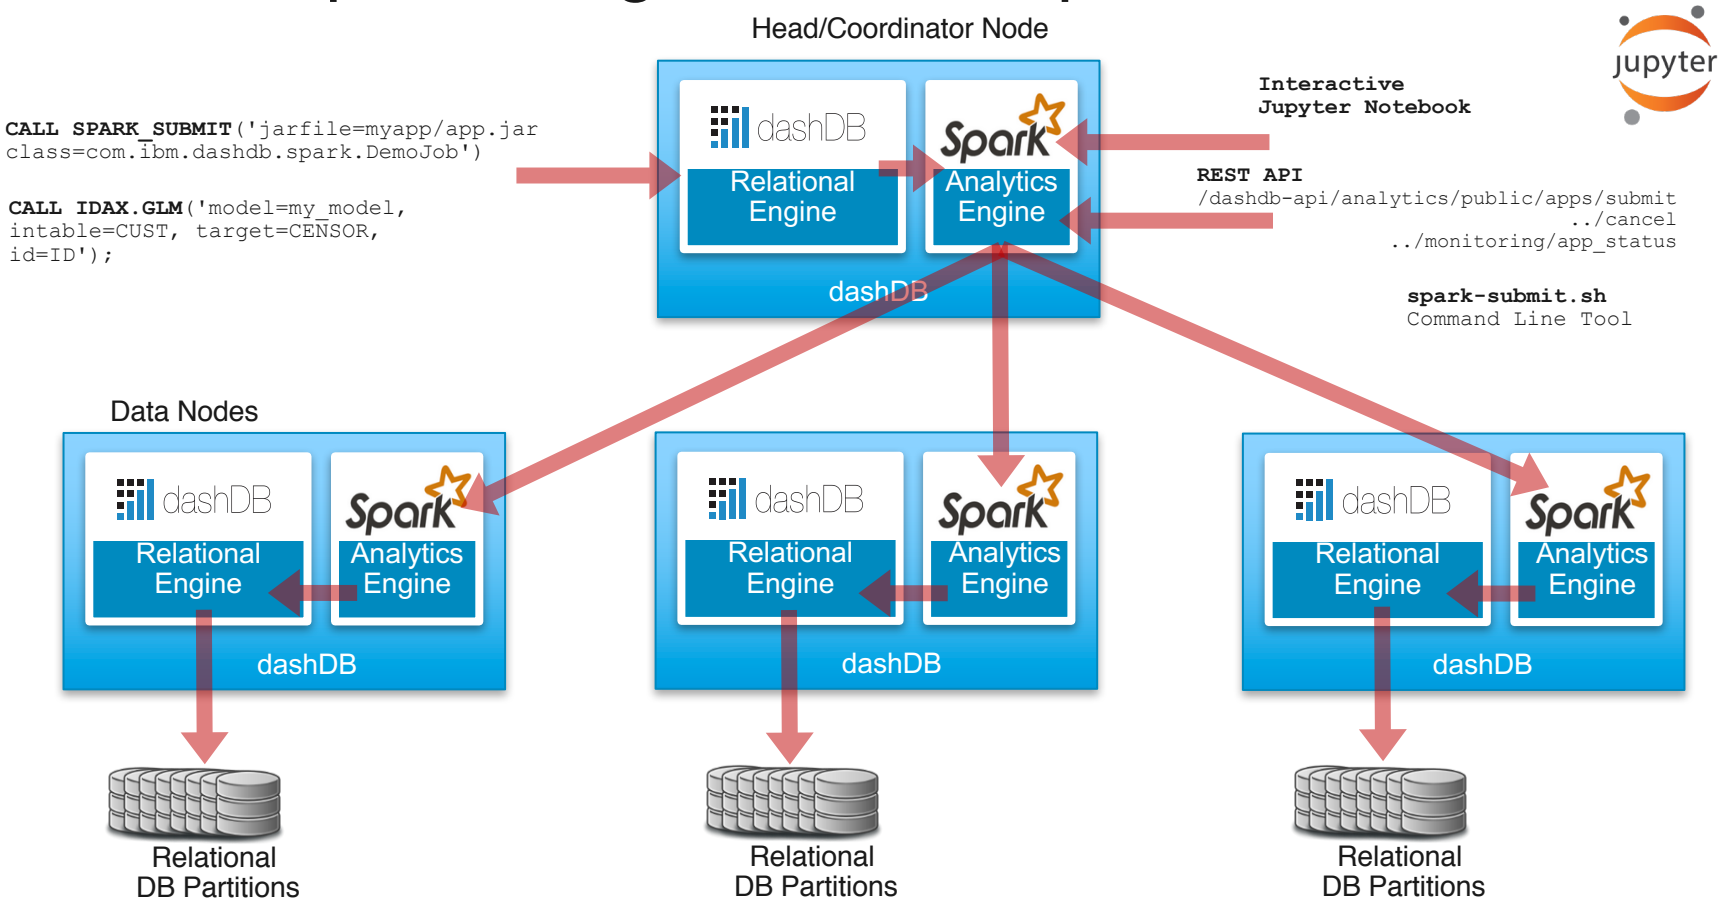

dashDB Spark Integration Examples: Future Options

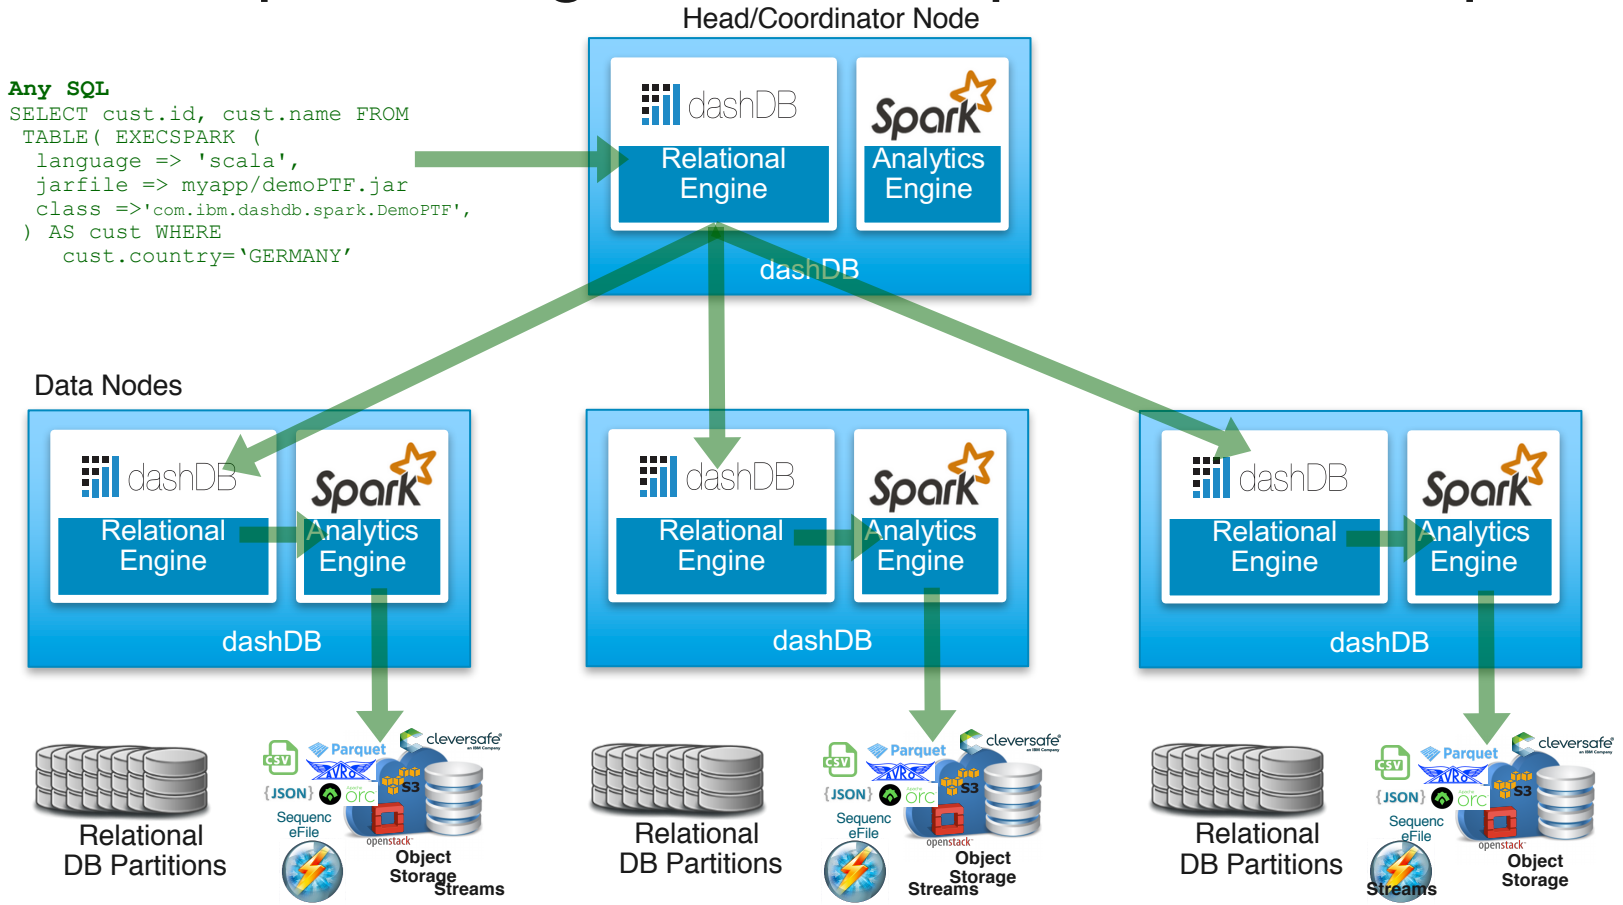

IBM BigSQL

- **Comprehensive, standard SQL**
  - SELECT: joins, unions, aggregates, subqueries . . .
  - UPDATE/DELETE (HBase-managed tables)
  - GRANT/REVOKE, INSERT … INTO
  - SQL procedural logic (SQL PL)
  - Stored procs, user-defined functions
  - IBM data server JDBC and ODBC drivers
- **Optimization and performance**
  - IBM MPP engine (C++) replaces Java MapReduce layer
  - Continuous running daemons (no start up latency)
  - Message passing allow data to flow between nodes without persisting intermediate results
  - In-memory operations with ability to spill to disk (useful for aggregations, sorts that exceed available RAM)
  - Cost-based query optimization with 140+ rewrite rules
- **Various storage formats supported**
  - Text (delimited), Sequence, RCFile, ORC, Avro, Parquet
  - Data persisted in DFS, Hive, HBase
  - No IBM proprietary format required
- **Integration with RDBMSs via LOAD, query federation**

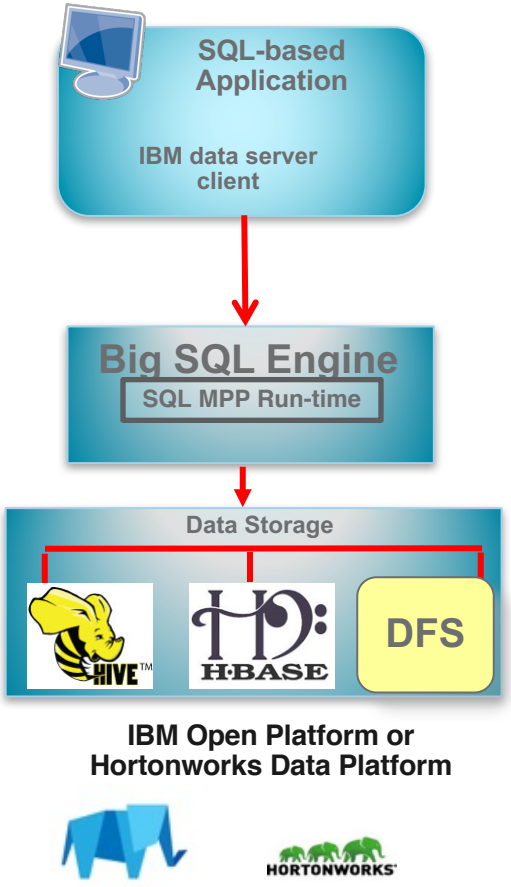

Wildfire – DBMS for new generation BigData Apps

- These applications want even more from DBMSs …
  - Higher ingest and update rates
  - Versioning, time-travel
  - Ingest and update anywhere, anytime (“AP” system)
  - Real-time analytics on real-time data (HTAP)
  - Rich analytics
- … but still want the traditional database goodies
  - Updates
  - Transactions (ACID)
  - Point queries (and not only via primary key)
  - Complex queries (joins, …) that require optimizer technology

Wildfire goals

HTAP: transactions & queries on same data

- Analytics over latest transactional data
- Analytics over 1-sec old snapshot
- Analytics over 10-min old snapshot

Open Format

- All data and indexes in Parquet format on shared storage
- Directly accessible by platforms like Spark

Leapfrog transaction speed, with ACID

- Millions of inserts, updates / sec / node
  - Multi-statement transactions
  - With asynchronous quorum replication (sync option)
- Full primary and secondary indexing
  - Millions of gets / sec / node

Multi-Master and AP

- Disconnected operation
- Snapshot isolation, with versioning and time travel
- Conflict resolution based on timestamp

Challenge: getting all of these simultaneously

Wildfire - Data lifecycle

Grooming: take consistent snapshots resolve conflicts
Post-grooming: make efficient for queries

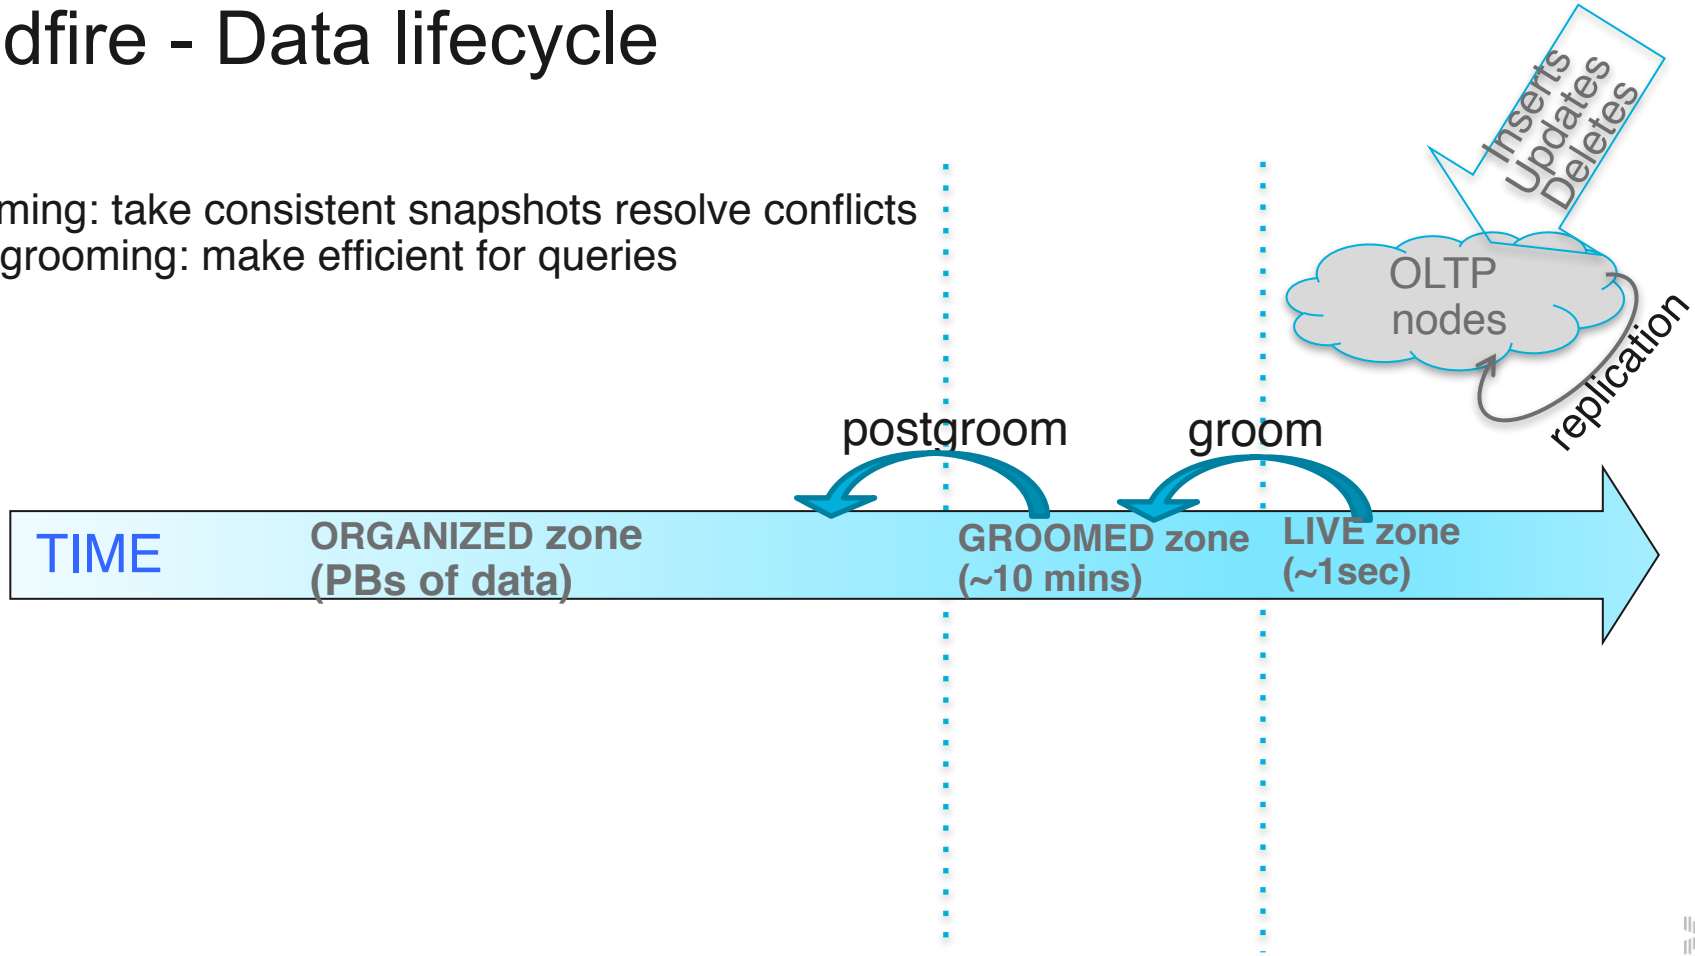

Wildfire - Data lifecycle

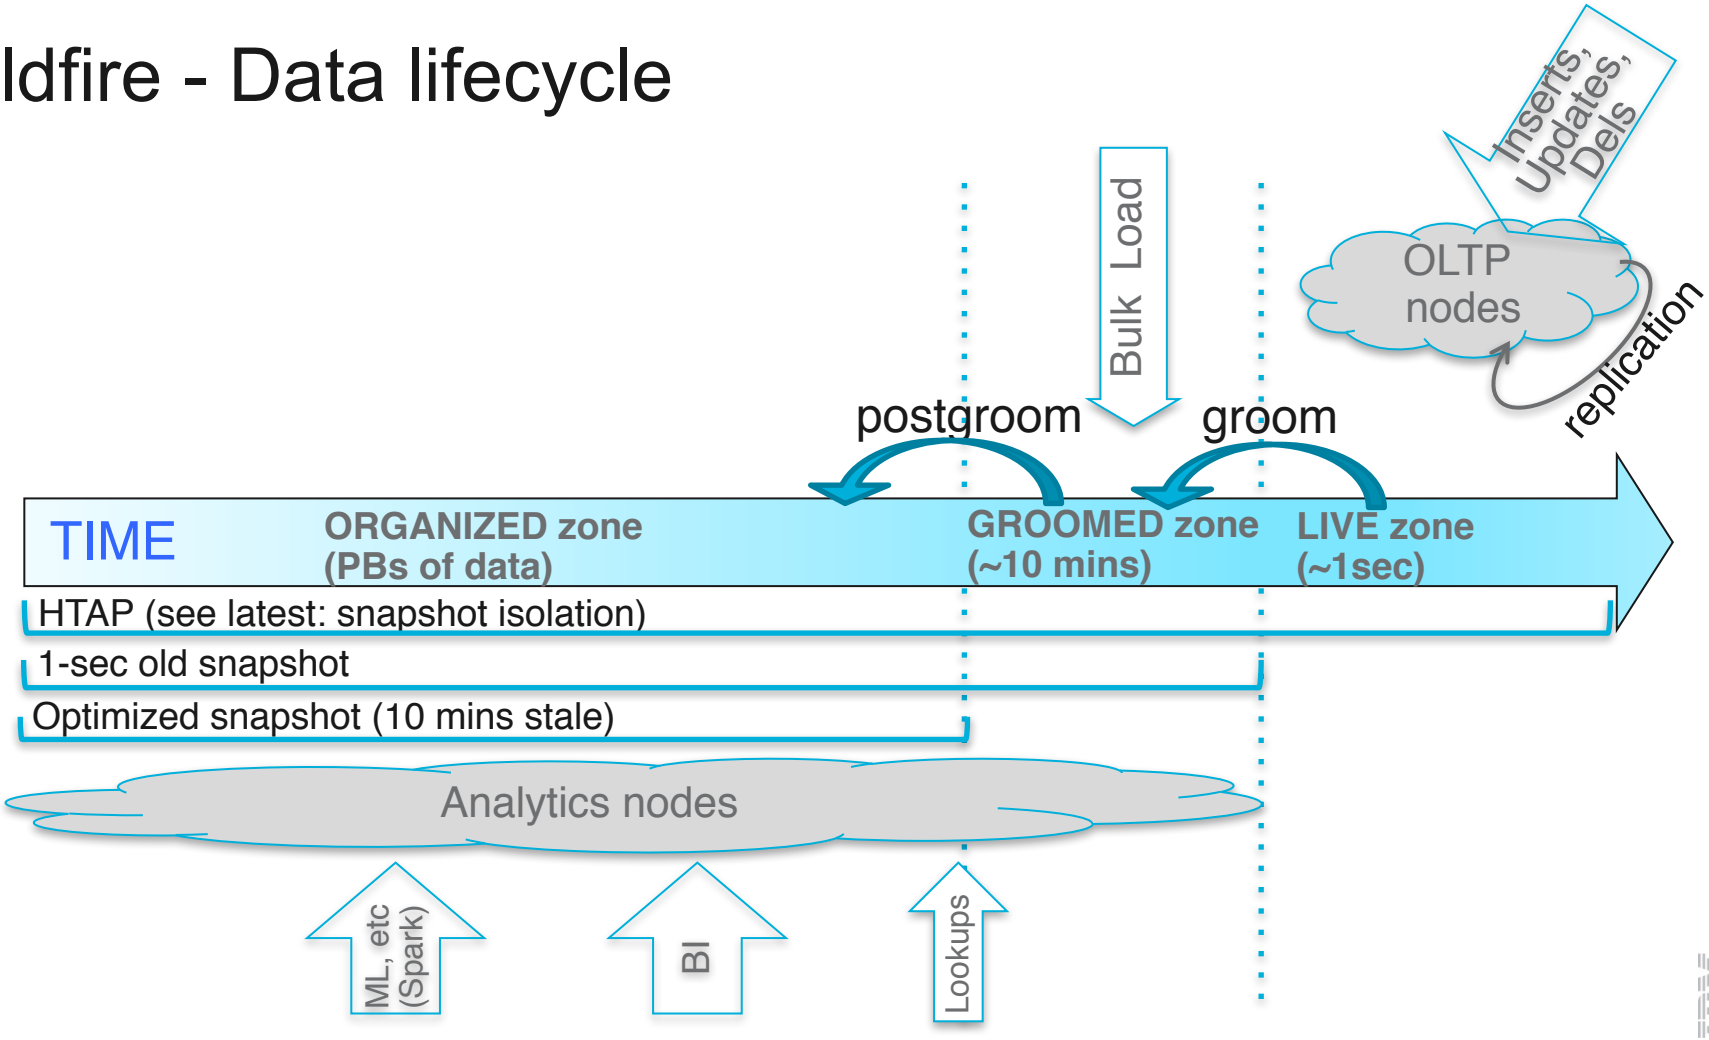

Implications to DBMS technology

Coexistence with ‘new’ technologies
Spark, Hadoop, Key-Value stores, Graph databases, …

Hybrid Transactional/Analytical Processing
Bringing analytics to transactional data

Hybrid cloud delivery
Fast deployment, continuous delivery, uniform experience …

Higher standards for traditional quality of services
Performance, scalability, continuous availability, security, …

Move the analytics, not the data

Real-time analysis is a game changer	**91%**	of data scientists are interested in real-time data for modeling
Transactional data is critical to real-time predictive models.	**85%**	of data scientists value transactional, operational and customer reference data
Time is wasted wrangling data	**94%**	of data scientists reported barriers, including time spent getting at data that may not be fresh.
Moving data bears security concerns	**63%**	of IT managers have security concerns around data transfer.

Base: 100 data science and data analytics leaders at enterprises within the US
Source: A commissioned study conducted by Forrester Consulting on behalf of IBM, May 2016 Forrester Thought Leadership Paper

Hybrid Transactional/Analytical Processing

Benefits
- Eliminating latency between data creation and data consumption
- Uniform access to any data for different types of applications
- Reducing redundancy of data by consolidating all or some of the layers
- Efficient data movement within the system, often not involving network
- Uniform policies and procedures for security, HA, DR, monitoring, tools, …

Challenges
- Mixed workload management capabilities
- Ensuring continuous availability, security and reliability
- Seamless scale-up and scale-out
- Providing universal processing capabilities to deliver best performance for both transactional and analytical workloads

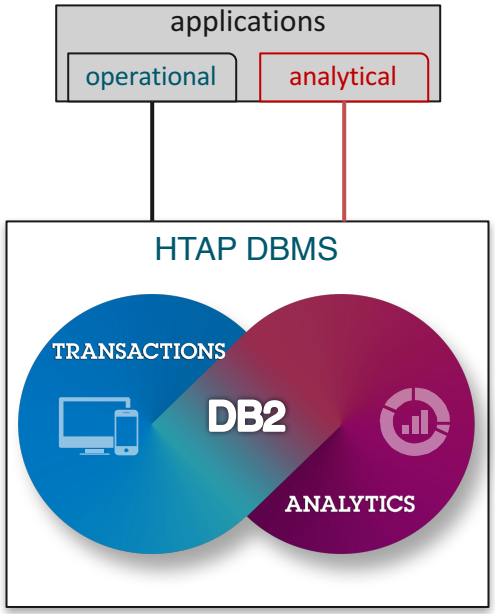

Approaches
- Large RAM enable 'In-memory' databases
- Columnar stores
- Large number of sockets, cores, servers
- Massively parallel processing
- Vector processing
- Hardware acceleration through special purpose processors: FPGA, GPU, ...
- Appliances

Building on proven technology base
- DB2 already provides superior technology to address most of the challenges
- The remaining challenge is being addressed by DB2/IDAA hybrid approach

CREATE TABLE T1 (C1 …, C2 …, C3 …, C4 …, C10…)
INSERT INTO T1 WITH VALUES (… 1000 rows ….)

Row Based RDMS

For every row
- Find a target page for the row
- Fix (aka pin) the page in the bufferpool
- Write a log record for the row insert
- Insert the row on the page
- Unfix (unpin) that page

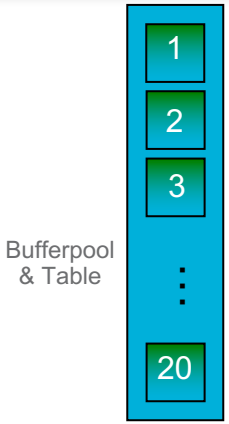

O(#rows)

1000 page fixes
1000 log records written

Non-optimized Columnar RDMS

For every row
  For every column
- Find a target page for the column value
- Fix (aka pin) the page in the bufferpool
- Write a log record for the column insert
- Insert the column on the page
- Unfix (unpin) that page

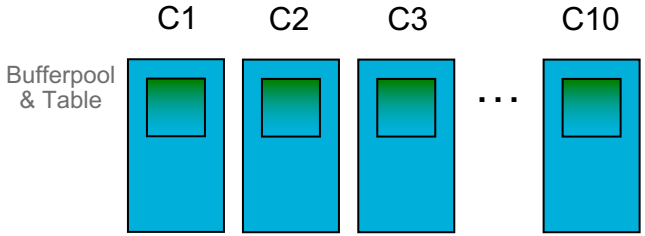

Let's say columnar compression allows data to be stored in half the total #pages

O(#rows*#columns)

10000 page fixes
10000 log records written

DB2 BLU

For a batch of rows
  For every column
- Fill local column buffers with col data

When local buffers full (or commit)
  For every full local column buffer
- Find a target page(s) in bufferpool
- Fix (pin) page(s) in bufferpool
- Write log record for all data buffer
- Copy data from buffer to pages
- Unfix (unpin) page(s)

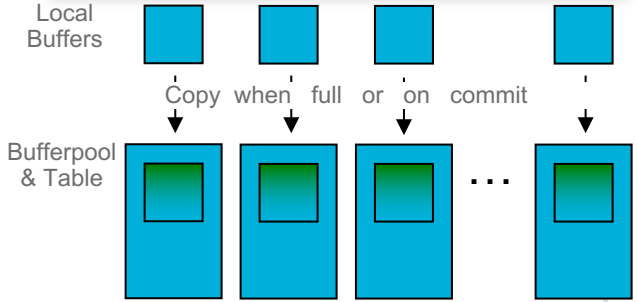

O(#pages)

10 page fixes
10 log records written

The ultimate HTAP platform

Supports transaction processing and analytics workloads concurrently, efficiently and cost-effectively

Delivers industry leading performance for mixed workloads

The unique heterogeneous scale-out platform

Superior availability, reliability and security

conventional transactions

analytics workloads

HTAP workloads

Basic idea

Reading most recent committed data during asynchronous replication

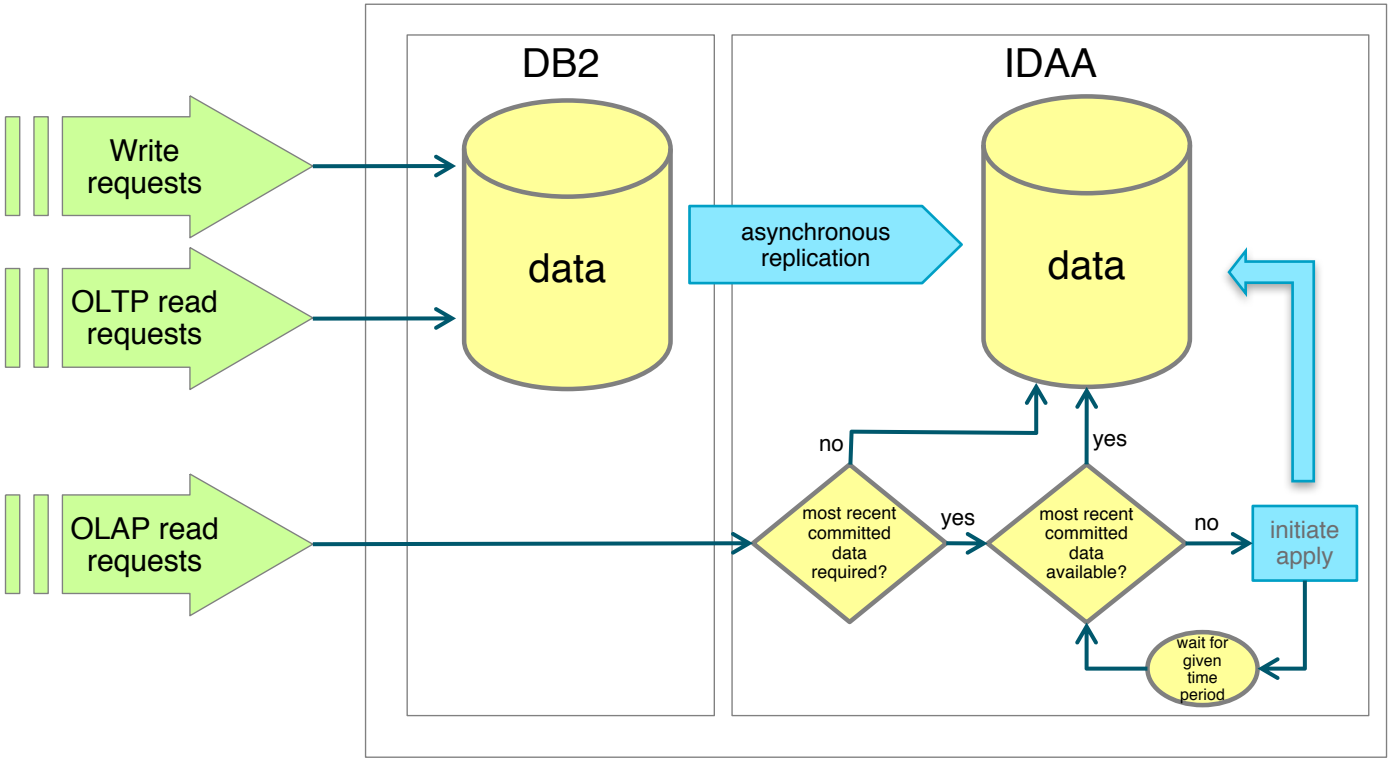

Implications to DBMS technology

Coexistence with ‘new’ technologies
Spark, Hadoop, Key-Value stores, Graph databases, …

Hybrid Transactional/Analytical Processing
Bringing analytics to transactional data

Hybrid cloud delivery
Fast deployment, continuous delivery, uniform experience …

Higher standards for traditional quality of services
Performance, scalability, continuous availability, security, …

Why enterprises choose to run in the cloud

- No infrastructure to install
- Get started right away
- Grow as fast as needed
- Scale, resiliency
- Advice on optimization and value
- Reduce needs on internal IT
- Low entry price
- Monthly operating expenses
- Subscription-based pricing low commitment
- Comprehensive view of software/service usage, cost, levels, entitlement: pay per usage for software/services
- Consolidated software catalog
- Consolidated monitoring, events, logs, access & analytics
- Real-time collaboration and views with service and support
- Automated deployment, update, security (scanning, compliance, access), backup, …
- Simple access to value-add cloud services

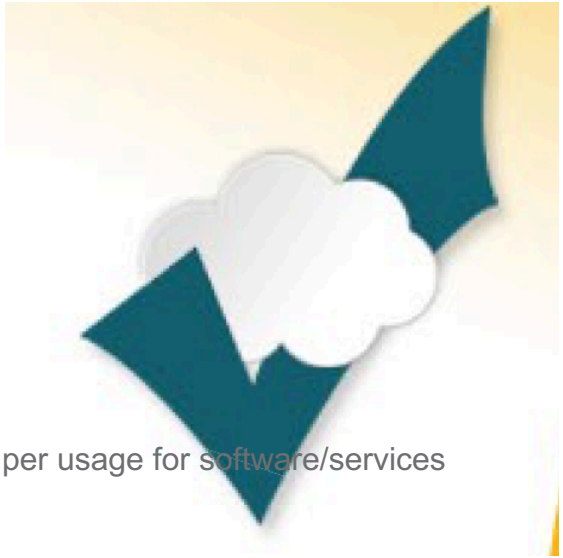

Why enterprises choose to run on premises

Control

Security

Downtime & Maintenance

Data gravity & Latency

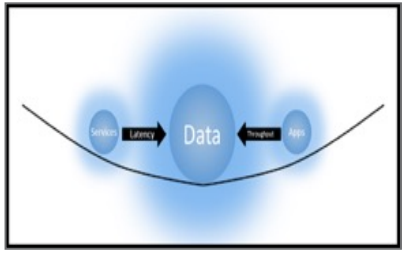

ROI

Regulatory

Accessibility

Visibility

dashDB – a hybrid data warehouse

- Single data architecture that supports all deployment models
- Start anywhere and no obstacles to migrate or expand