

Trafodion

Operational SQL-on-Hadoop

SophiaConf 2015
Pierre Baudelle, HP EMEA TSC
July 6th, 2015



Hadoop workload profiles

Operational

- Real-time analytics

• **Operational SQL**
= OLTP +
interactions

Interactive

- Parameterized reports
- Drilldown visualization
- Exploration

Non-interactive

- Data preparation
- Incremental batch processing
- Dashboards, scorecards

Batch

- Operational batch processing
- Enterprise reports
- Data mining

Sub-second

Response Time


Hours




Characteristics of operational DBMS applications

Generalized characteristics and requirements:

- Low latency response times
- ACID (data consistency guaranteed) transactions
- Large number of users
- High concurrency
- High availability
- Scalable data volumes
- Multi-structured data
- Rapidly evolving data requirements (i.e. flexible schemas)



Expose Hadoop limitations



Transaction
Support




Data
Integrity



Real-time
Performance



Operational Query
Optimization



Workload
Management

Trafodion - Introduction

Trafodion is a joint HP Labs and HP-IT research project to develop operational SQL on Hadoop database capabilities

Complete: Full-function SQL

Reuse existing SQL skills and improve developer productivity

Protected: Distributed ACID transactions

Guarantees data consistency across multiple rows, tables, SQL statements

Efficient: Optimized for low-latency read and write transactions

Supports real-time transaction processing applications

Interoperable: Standard ODBC/JDBC access

Works with existing tools and applications

Open: Hadoop and Linux distribution neutral

Easy to add to your existing infrastructure and no vendor lock-in

Leveraging 20+ years of operational database investment!



Hadoop

The diagram consists of two blue rectangular boxes stacked vertically. The top box contains the word 'Hadoop' in white text. Below it is a white plus sign. The bottom box contains the words 'Operational SQL' in white text.

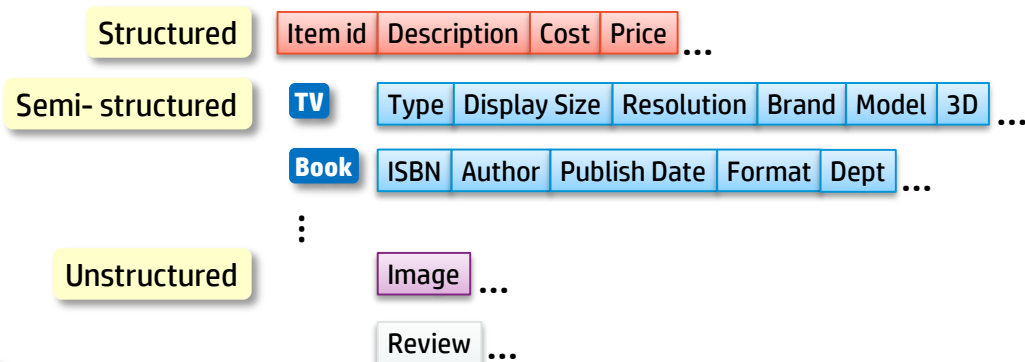
+

Operational SQL

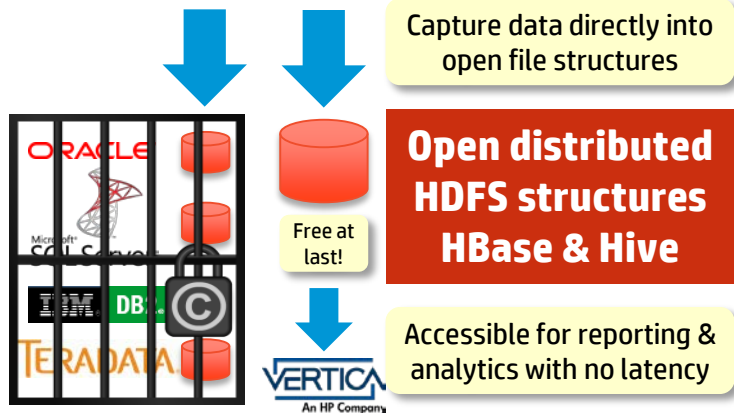


Operational SQL on Hadoop – Use cases

- Integration of structured, semi-structured, and unstructured support
- Using Hadoop for all RDBMS needs – Not Only Big Data
- Operational transactional workloads



```
SELECT all TVs WHERE Price > 2000 and  
Type = 'Plasma' and Display Size > '50'  
and customer sentiment is very positive
```



Trafodion innovation built upon Hadoop stack

Leverage Hadoop and HBase for core modules

Maintain API compatibility

Differentiation

ANSI SQL via ODBC/JDBC

Relational schema abstraction

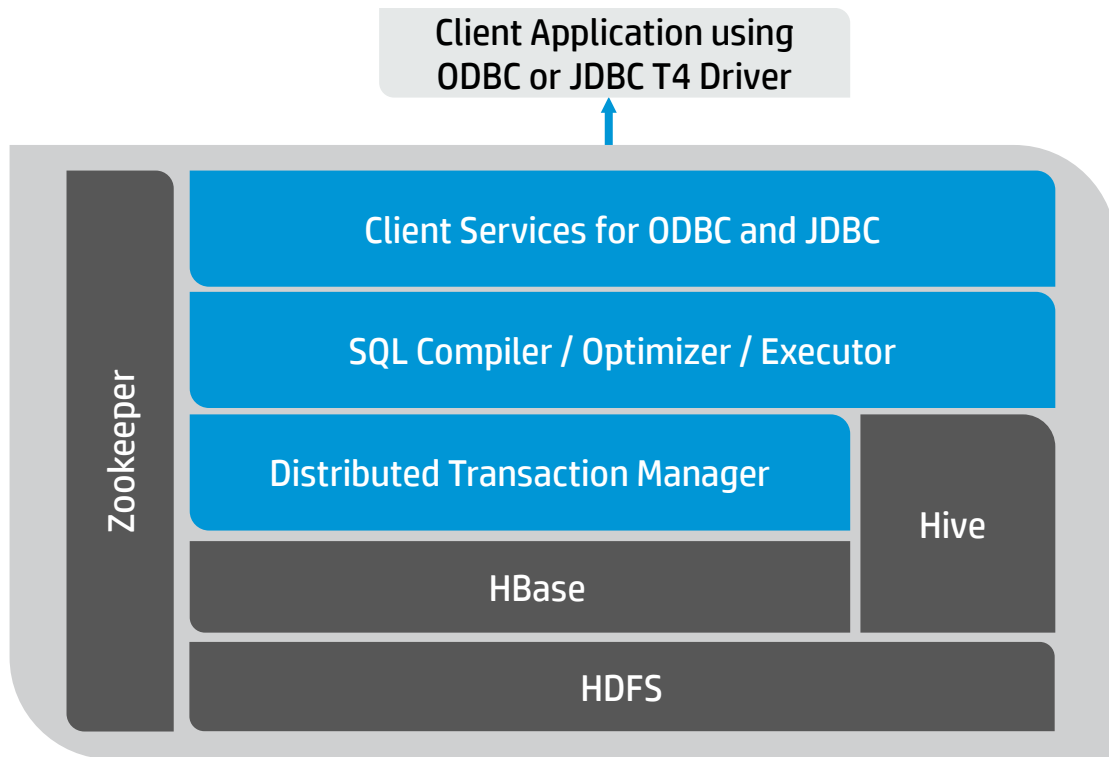
ACID transactions

Low latency read/writes

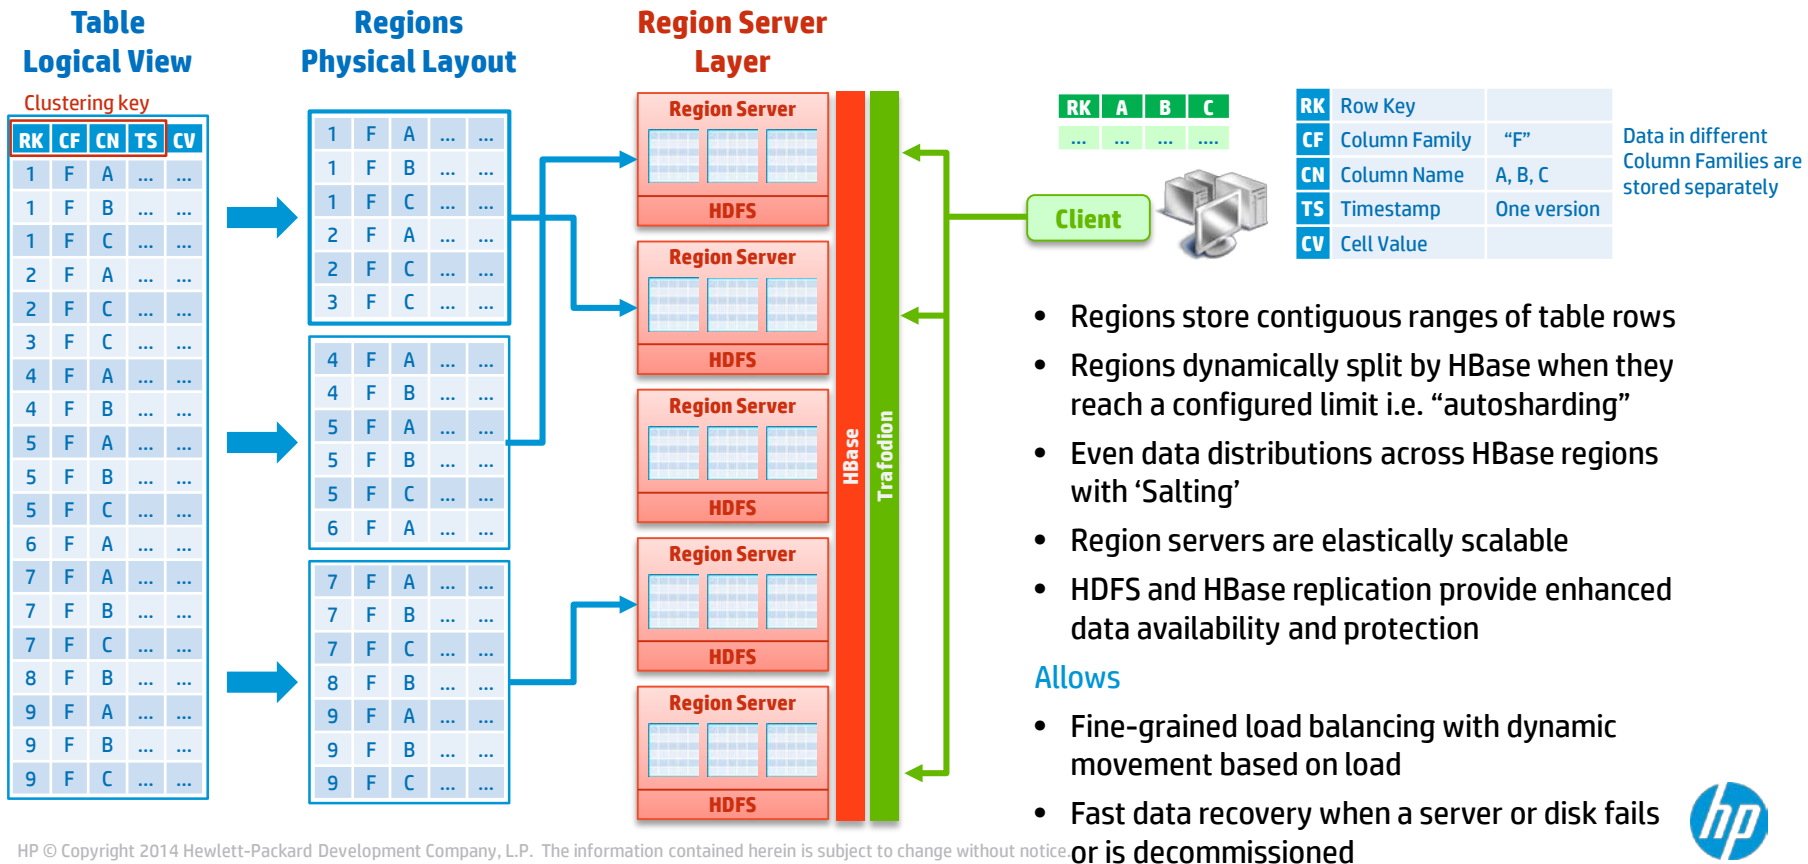
Parallel optimizations

Standard
Hadoop

Trafodion

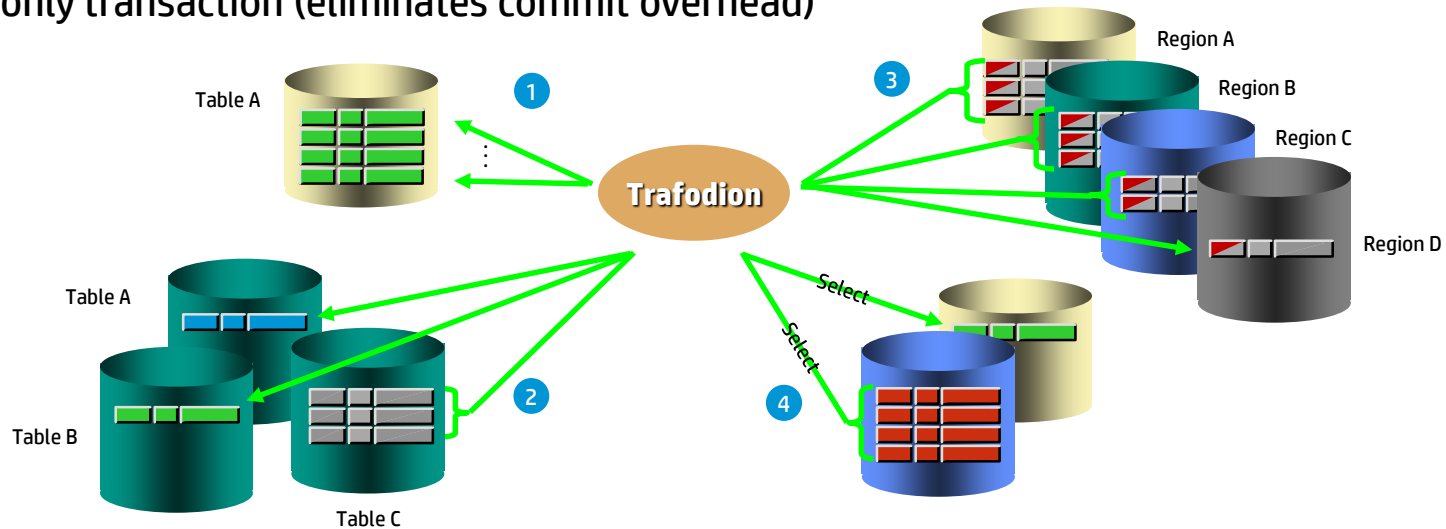


Leveraging HBase for scalability and availability



Distributed transaction protection

- 1 Multiple row inserts, updates, and deletes to a table
- 2 Multiple table and SQL insert, update, and delete statements
- 3 Distributed multiple HBase region insert, update, and delete transaction (2-phase commit)
- 4 Read-only transaction (eliminates commit overhead)



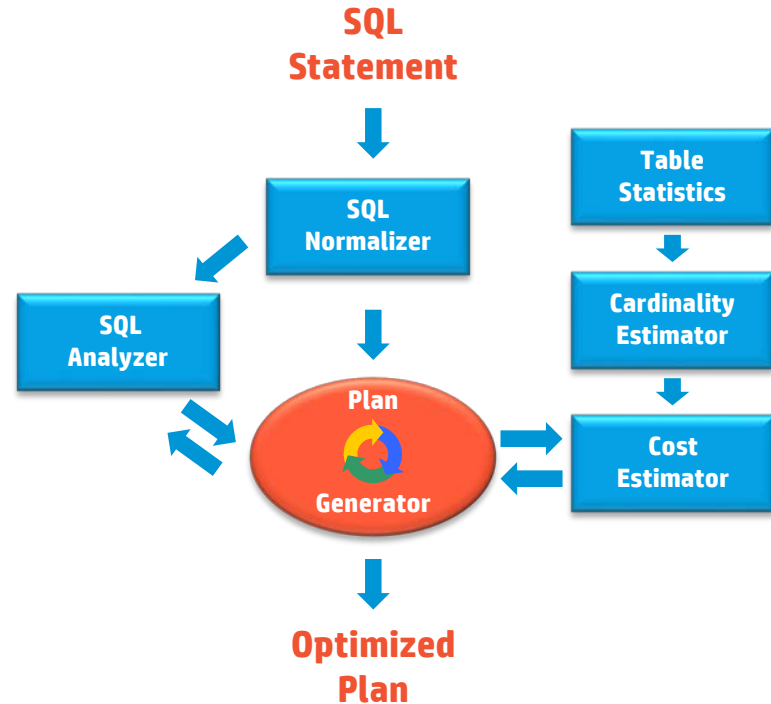
Optimized execution plans based on statistics

Optimizer features

- Top-down, multi-pass optimizations, branch and bound plan pruning considers more potential plans
- Utilizes “equal-height” histogram statistics
- SQL pushdown considerations e.g. predicate evaluation
- Eliminates sorts when feasible, syntactically and semantically
- In-memory vs. overflow considerations
- Optimal degree of parallelism (DOP) considerations including non-parallel plans

Benefits

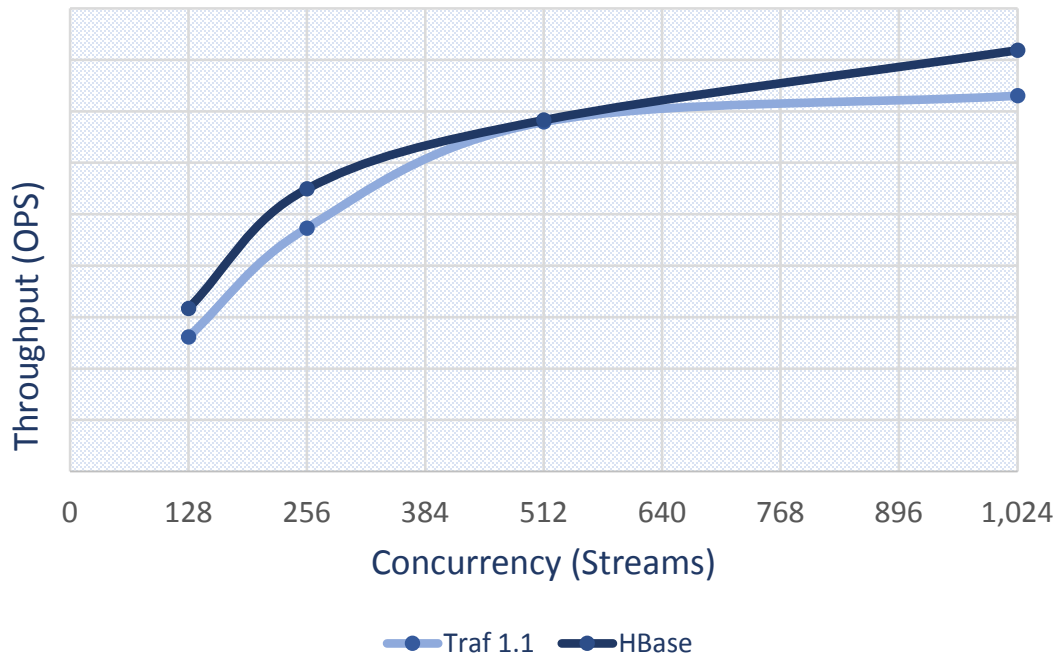
- Facilitates enhanced parallelism and SQL object handling efficiencies
- Optimizations for operational transactions and reporting workloads



Trafodion performance objective

YCSB operation speeds that approach HBase (within 20%)

YCSB Singleton5050 (Workload A)



Meets current objective!

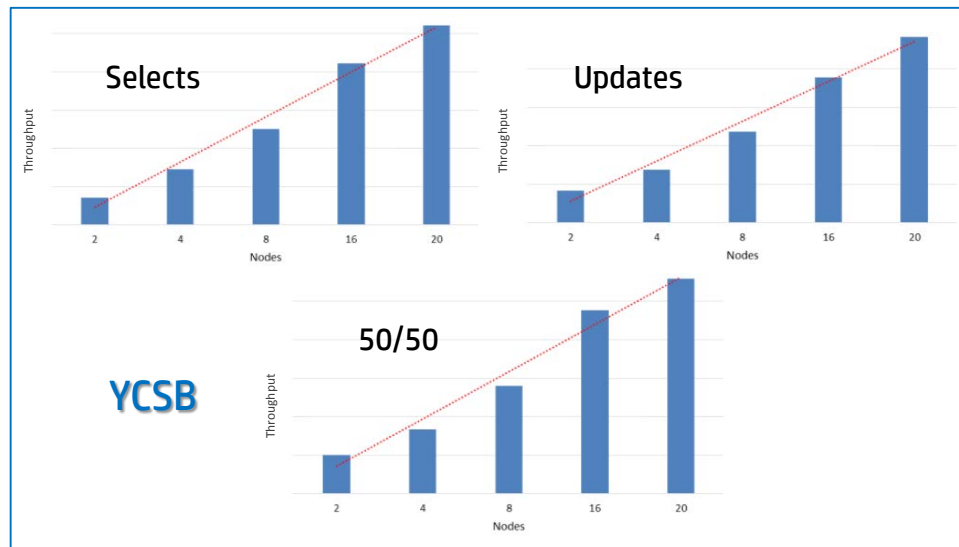
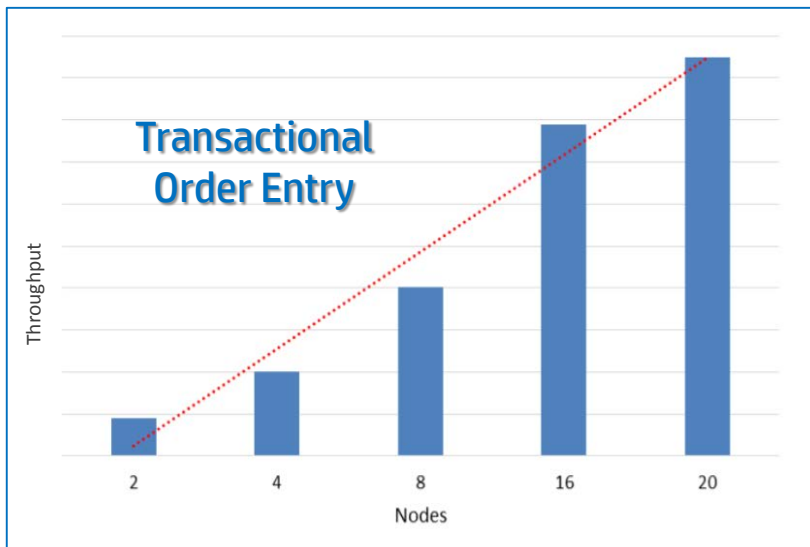
With max variance at 10.8%



Trafodion performance objective



YCSB and Order Entry scale linearly through 20 nodes (240 cores)!



Meets objective!



Profiling Trafodion use cases

HBase Expansion

Expand workloads onto HBase and maximize return out of HBase/Hadoop investment

Assumptions

- Hadoop and HBase already exist, and interested to expand it to process more operational workloads

Value Proposition

- Maximize Return out of current HBase/Hadoop Investment through processing more workloads

Key Technical Requirements

- Use a powerful SQL engine to query / update HBase
- Ensure ACID distributed transaction protection is a key
- Need to federate data between multiple sources very efficiently in a single query
- Provide performance, scalability and HA to meet SLA

RDBMS Migration

Operational workload consolidation and migration onto Big Data platform

Assumptions

- Don't want to have a lock-in to RDBMS vendor proprietary technology (neutral or in favor of open source model) need to reduce overall license cost
- Interested in or already have a Hadoop centric Big Data strategy in place

Value Proposition

- Avoid vendor lock-in and reduce overall cost of data technologies
- Leveraging existing application investment while meeting elastic scalability with schema flexibility

Key Technical Requirements

- Need to resolve scalability issues from RDBMS
- Strong elastic scaling requirements
- Interested in establishing "Big Data" platform on Hadoop eco-system
- Need to leverage existing investment in application with SQL interface capability
- Clearly definition of "operational performance" requirements
- Need to have ACID distributed transaction protection with expected performance and HA to meet SLA requirements on operational workloads)

Operational and Analytics Together

Scalable, flexible and high performing Big Data platform that combines both operational and analytic workloads

Assumptions

- Already committed to Hadoop, and plan to establish a Hadoop centric "big data ecosystem" (e.g., data lake, etc.)

Value Proposition

- Reduce overall time-to-value to process data from operation to analysis

Key Technical Requirements

- Integrate structured, semi-structured, and unstructured data into an Enterprise Data Lake to host operational, historical, external (Big Data), and Master data
- Gain insights across this Enterprise data not limited to just Big Data
- Eliminate massive data movement across platform between operational and analysis, dramatically reduce time-to-value
- Reduce movement and replication of data between databases
- Perform operational / transactional processing to ensure data quality, integrity and transformation on Hadoop



Vehicle telemetry operational data store

Real time capture, monitoring, and analysis at scale with high concurrency

Challenge


- Thousands of devices transmitting today, 2x in couple of years
- Real time onboarding 100's of million of events every day
- Query the data upon arrival along with historical data in real time

Scalability Requirement

- Sub-second query response times
- Sustain performance as concurrent users increase to >100

Proposed Solution

- Trafodion on standard x86 and Linux cluster
- Data load, query and extract in parallel
- User can query both current and historical data



Benefits

- Data driven business management
 - ✓ Central tracking and querying
 - ✓ Alarm reconciliation
 - ✓ Fleet health management
- Scalable, reliable, and real time
- Open source with no vendor lock-in

Re-hosting of telco operational DBMS systems

Single data management platform for rating, CRM, funds, and billing systems

Challenges: Scale, Mixed Workload

Load

- Voice, SMS and data files generated by 100's of million users across the country
- Gigabytes of data to be loaded in minutes

Rating

- Process raw arriving data and load into Trafodion
- Rate freshly loaded data to track usage within minutes

CRM

- High transaction rates
- Comprehensive queries against historical data with high concurrency

Funds

- High transaction updates
- High concurrency

Billing

- Complete 100's of million user accounts processing at the end of month within hours

Proposed Solution




Benefits

- Single integrated and scalable data management platform
- Open source with no vendor lock-in

Trafodion summary

Delivers a full featured and optimized operational SQL on Hadoop DBMS solution with full transactional data protection

- Leveraged use of SQL learnings and expertise versus map/reduce programming
- Foundation for next generation real-time transaction processing applications
- Guaranteed transactional consistency across multiple SQL statements, tables, and rows
- Retention of Hadoop benefits i.e. reduced cost, scalability, elasticity, etc.



~~HP~~ Trafodion V1.1.0

(Open Source since June 2014)



Forrester - Mike Gaultieri (October 22nd, 2013)

'The Future of Hadoop is real time and transactional'

Doug Cutting (October 30th, 2013)

'We're in the middle of a revolution in data processing'

'... it is inevitable that we will see just about every kind of workload be moved to this platform – even OnLine Transaction Processing' (OLTP)

Project Trafodion 1.1 entered Apache Incubator in May 2015

For more information on Trafodion contact Rohit Jain at rohit.jain@esgyn.com



SCS Cluster / Big Data Workgroup

More than 25 Academic and Industrial partners in PACA

- Combine Big Data competencies from SCS members
- Align with the SCS Cluster's strategy
- Deliver documents on Big Data architecture
- Test and validate Big Data approaches via Proof of Concept
- Support the emergence of collaborative projects (French, European) with SCS Cluster



<http://www.pole-scs.org/p%C3%B4le-scs/smart-specialisation-areas/qt-big-data>

Thank You

