



Brainstorming Session on Database Technologies

Agenda

- 4 main challenges
 - Data size and rate
 - Hardware evolution
 - High availability
 - Machine learning

Data size and rate

Time series workload

- Workload assumptions
 - Billions of individual data points
 - High write and read throughput
 - 250k rec/s from LHC sensors
 - Mostly an insert/append workload
 - Large deletes
 - Multiple use-cases

High ingest

- Accelerator logging 2.0
 - Data coming from different sources
 - Possibly **250k records/s**
 - Data Latency $< \sim 30\text{sec}$
 - No data losses acceptable
 - Temporary data store in the ingestion layer
 - Provide data transformation features
 - Enhancing by adding some context
 - Filtering

High frequency query challenge

- ALICE file catalog
 - Currently
 - Single instance of MySQL
 - Few TBs of data, ~8.5B rows
 - Read access rate – 11500 Hz
 - Run 3
 - 10X increase in data size and insert rate
 - Robust solution is needed

Hardware evolution

Memory evolution

- NVMe
 - Flash over PCI express
 - High bandwidth at low latency
- 3D XPoint
 - Upcoming NV technology
 - Faster and more stable than traditional PCM
 - Density 4-times one of DRAM
- An option for huge In-Memory databases, SPARK analysis
 - Imagine an event index in memory

Computing evolution

- GPU
 - Training of deep learning models
 - Neural network
 - Significant speed-up comparing to CPUs
 - Challenges
 - Algorithms complexity & low level programming
 - Data parallelization
 - Heterogeneous systems

High availability

High availability for Big Data

- New architecture of Hadoop Service
 - At works now
- Ensuring disaster recovery
- Increasing backup capabilities
- Providing sufficient and flexible computing resource
 - Resource expansion for peak usage
- Optimizing cost of operations

Cloud

- One of the options
 - Oracle cloud is being currently tested
- Key factors
 - Feasibility of using cloud with complex environments
 - Experiment's and technical network, SSO...
 - Performance/latency
 - Effort needed for data migration
 - Cost
 - Models differ between companies and even regions

Machine learning

Physics analysis in DB type of workload

- Ongoing projects in experiments
 - CMS Big Data project
 - Speed up HEP analysis 100 times
 - Process an input sample of 1PB within 5 hours
 - SWAN Project
 - Cloud HEP data analysis

Optimizing data placement

- Ongoing investigations in CMS
- Intelligent dataset placement across sites
 - 70 sites
 - ~20 Petabytes of data produces/year
 - For HL-LHC
 - Event size x3
 - Event number x5
- Reduction of transfer and storage costs
- A base for a study on optimal job scheduling

4 main challenges

- Data size and rate
- Hardware evolution
- High availability
- Machine learning