Database Futures Workshop

https://indico.cern.ch/conferenceDisplay.py?confId=130874

Tony Cass

(Reminder of) Aims

The aim of the workshop is to discuss possible future needs in the database area. In particular, we would like to understand likely future database applications from the different user communities (experiments, accelerator controls, engineering support, IT, AIS, …) and how these might map to different database technologies (relational/nosql) and implementations (oracle/mysql and then pick freely from http://en.wikipedia.org/wiki/NoSQL).

- Tracks
  - Requirements
  - Implementations
  - Technologies

Database Futures Workshop

Rapid Summary

With some reflection after comments at the workshop

Tony Cass

- Oracle
- Other SQL
- NoSQL
- Abstractions

Oracle

- Many mission-critical applications
  - Relatively well understood and relatively stable
- Administrative & Engineering
  - “Trivial” data volumes, but many users
- Experiments
  - Large data volumes, but growth expected to be ~linear with physics data volume
    - In some cases, hardware capability growth outstrips system requirement growth.
- Accelerator
  - Very large data volume, pressure to store more (e.g. Beam Loss Monitor data)
  - O(10PB) by 2020
  - O(EB) for CLIC, however—and LHC initial estimates were off by factor of 50!
- Live long and prosper!

Other SQL

- Other SQL is mostly MySQL; some SQLite.
- Systems run fine now (c.f. low support requirements for ALICE DAQ environment)…
- … but if someone were to offer support, there would be takers.
  - Not clear the support costs can be constrained, however:
    - Unlikely that a single “canned” MySQL environment would be acceptable,
    - and why stop at MySQL? GS/AIS are rejecting interesting SQLServer apps because they are not Oracle based.

NoSQL - I

- The key issue seems to be the difficulty of providing efficient read performance for essentially random queries—databases have been optimised for inserts and “production queries”.
  - Even an issue in BE where people have been well used to tight constraints on read access as this continues the practice of the LEP era.
- The time required to develop an application delivering reasonable performance seems to be less with NoSQL systems than with Oracle (the target of all comparisons).

NoSQL - II

- But, for me at least, the requirements were still unclear.
  - How many “random read” applications are there? Will each need a dedicated NoSQL database?
    - How much easier is “create a new Cassandra table” than “add another index in Oracle”?
  - Is ACID too much of a straight-jacket?
    - No NoSQL application presented aggregates information from multiple sites…
    - … but PhEDEX, which does, is only lazily consistent, apparently.
  - Can a single NoSQL database support all requirements?
    - Unlikely: both ATLAS (2) and CMS (3) have chosen different platforms for different applications based on suitability.
  - Is ease-of-setup at the cost of future maintenance woes?
- Is today’s MySQL support model adequate for the NoSQL world?

From Simon Metson’s talk

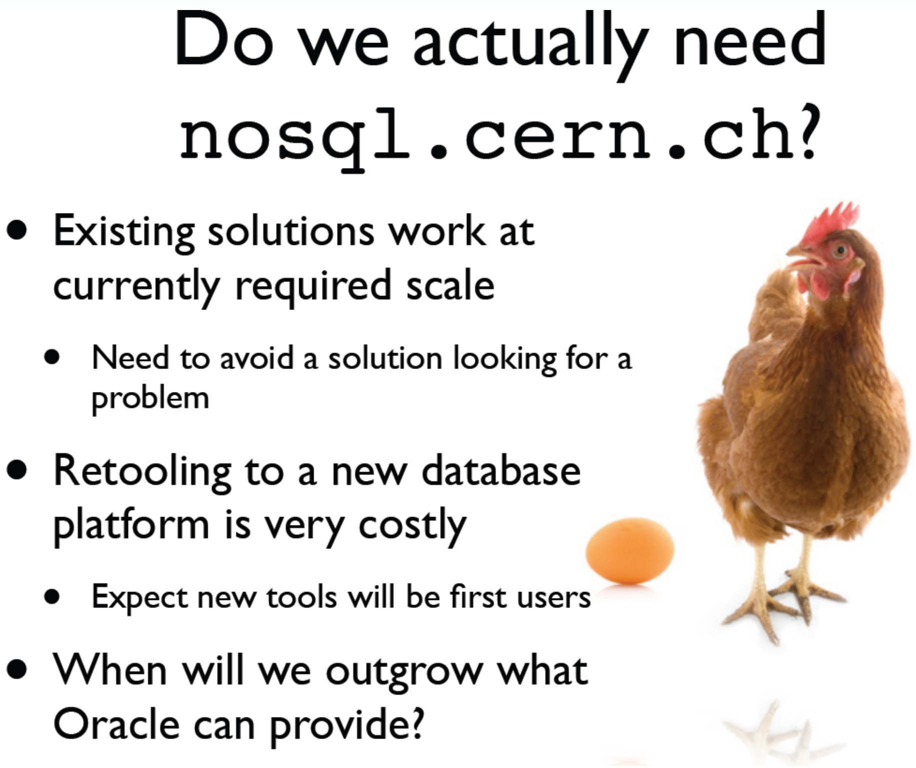

Do we actually need `nosql.cern.ch`?

- Existing solutions work at currently required scale
  - Need to avoid a solution looking for a problem
- Retooling to a new database platform is very costly
  - Expect new tools will be first users
- When will we outgrow what Oracle can provide?

- “Game changing NoSQL use needs significant research before being justifiable”
  - Game changing? Physics data in a database
    - i.e. what analysis programs run against…

Challenges

- BE applications and data volumes are on the edge of what it is possible* to handle in Oracle today, those of the LHC experiments are not.
- Lassi Tuura: “Just because you needed all that data to debug that problem doesn’t mean you need to store it for the next 20 years.”
  - Collect maybe, but store with a sliding window.
- Intelligent data summation: maybe keeping average & max/min values is **better** than simply keeping all data points: brief excursions would be more clearly identified.
- What weight should application developer productivity have in total cost of ownership?
  - More likely: how to transition from ease of development to ease of maintenance and operation?

* for CERN (within current budgets — kW or kCHF)