

**CMS requirements on IT
provided NoSQL data stores**

Current CMS usage

- Hadoop (just HDFS)
 - CouchDB
 - MongoDB
 - KyotoCabinet
- } see next talk

General NoSQL observations



General NoSQL observations



Fault-tolerance

by @jrecursive





When all you have is a hammer
everything looks like a nail

General NoSQL observations

- Generally simple to set up and manage
- Most designed for “web scale” applications, targeting clustered deployments
- NoSQL, while catchy, is a poor name
 - Especially when so many of them have a SQL-like interface...

General NoSQL observations

- Specialise and drop traditionally "required" features
- Less forgiving than traditional SQL databases
- Optimised for specific access patterns/behaviour, outside of those very suboptimal

General NoSQL observations

- Choice of tool must be done carefully
 - Eventually consistent systems are ***eventually consistent...***
- Very use case dependent
- Lot of choice on the "market"

General NoSQL observations

- Good for startups with limited resources and exposure to risk
- Good for large companies who build data centres with lots of loosely related data and large DevOps teams

Scale

- Designed, and simpler to run, at scales higher than traditional databases
- Majority of NoSQL stores promise automagical horizontal scaling
- For many use cases that isn't actually critical, scale point is beyond where the use case sits
- Experiments can (and already do!) successfully run these tools at low scale

Looking back

- CERN has, to date, provided an excellent Oracle service
- It has not provided a MySQL cluster, groups have managed MySQL themselves
- Create a parallel for a NoSQL solution

What would nosql.cern.ch look like?

- Service, complimenting the Oracle RAC's upon which VO's can build applications on (a la PhyDB)
- Needs to provide a reliable service that experiments cannot provide themselves due to:
 - complexity
 - scale
- It should be a sizeable, expandable cluster, partitioned by VO

What would nosql.cern.ch look like?

- Two classes of NoSQL stores available:
 - Complimentary
 - Game changing
- In my view only game changing stores are really interesting

Community

- "Rolling our own" would be a disaster
- CMS would not be interested in a CERN specific NoSQL project
 - Smaller support base
- Engage with existing community (majority of NoSQL tools are open source)

Community

- Consider community/ecosystem when choosing technology
- Look for systems running at, or beyond, our scales
 - let someone else find those scaling issues
- **CONTRIBUTE**, don't fork/start new project

Technology choice

- Differences between NoSQL databases much larger than between traditional databases, at least to the end user
- Decision between key:value stores and more complex data processing systems
- IMHO technology choice less important than commitment of support

Software?

- MongoDB/CouchDB
 - Already in use, would use a suitable centrally provided system were one available
 - Yet to hit scale issues, and possibly never will
- Cassandra
 - Not used in CMS
 - Apparently not quite at interesting scale, O(150TB)
- Only game changer, in my view, is Hadoop

Why Hadoop?

- Proven from terabyte to multi-petabyte scales
- Vendor and student awareness
- Large, active open source community
 - Already some use in HEP

CMS Requirements

- Pick a single existing project and contribute to it
 - Fragmentation of effort could be a concern, due to number of tools
- CMS *cannot* manage a usefully sized Hadoop cluster, and would need IT to be involved were we to start using it significantly

Let imagine it's Hadoop

- This would be a major project
 - Major changes to purchasing and provisioning of servers, and operations
 - Applications need re-architecting for Map:Reduce paradigm
- Next step would be to identify some demonstrator use cases

Possible demonstrators

- Identify use cases where relational databases are problematic:
 - Log analysis - ideal for Dashboard “historical view” etc.
 - Analysis of application performance
 - More complex data catalogues
- Worth noting that specific classes of analysis/processing become possible in new regime


Do we actually need nosql.cern.ch?

- There are sites “out there” that already have Hadoop clusters
- Could we run demonstrator projects there and move experience back to CERN?
- Minimum cluster size to be interesting $O(100)$ TB



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?



Conclusions

- Would be happy to move to an IT provided CouchDB/MongoDB instance were a suitable service offered
 - Complimentary service, not be a game changer
- Game changing NoSQL use needs significant research before being justifiable
- Agree on demonstrators, possibly run offsite

Conclusions

- NoSQL is no panacea
- Diversity of NoSQL space makes a centrally provided solution difficult
- Solutions beyond the scope of a single organisation