



NDB performance tuning

Yves Trudeau, Sun

Innovation Everywhere

Bottlenecks, where?



- Storage nodes CPU?
- Storage nodes disks?
- MySQL?
- Network?



NDBD CPU too high

- Missing index

2009-04-16 15:42:51 [MgmSrvr] INFO -- Node 2: Trans. Count = 4, Commit Count = 0, Read Count = 4, Simple Read Count = 4, Write Count = 0, AttrInfo Count = 4, Concurrent Operations = 0, Abort Count = 0 **Scans = 4** Range scans = 0
2009-04-16 15:42:51 [MgmSrvr] INFO -- Node 2: Operations=5

- Partitioning inefficient

- Realtime options

Using CPU efficiently



- How many NDBD per server
- Assigning threads to CPU
- Managing interrupts load



Disk is full!!!!



- What's on disk?
- Reducing the disk footprint



Temporary tables

- If persistence is not critical
- Good for high write load
- Not so volatile




Controlling query latency



Variation of SchedulerSpinTimer

Histogram of queries by time



About MySQL



- How many SQL nodes?
- How many connections per nodes?
- Number of NDB connections?
- Engine push condition?

Why custom partitioning?



- Basically... to scale queries
- Use case



Optimizing joins



- Nested loop is slow with NDB
- “IN” clause behave better
- NDB API is flexible
- BKA is comint

e1000 nics...



- What's the problem?
- Symptoms
- How to solve

