

MySQL Cluster 7.2 – what's new

October 2011

Johan Andersson

Severalnines AB

johan@severalnines.com

Cell +46 73 073 60 99

several**nines**

Agenda



- Fast JOINS
- Memcache interface
- Distributed GRANTS

Until now..

The logo for Severalnines, featuring the word "several" in a light grey sans-serif font and "nines" in a bold blue sans-serif font.

- JOINS are painfully slow in 7.1
 - Mostly due to optimizer in MySQL Server
- JOINS are executed in the MySQL Server
 - Table by table record by record.
- Queries must be FORCED INDEXed and/or STRAIGHT JOINed
 - Human head better than MySQL Optimizer
 - No index statistics
- GRANTs are not distributed

Until now..

The logo for 'severalnines' is displayed in a white sans-serif font. The word 'several' is in white, and 'nines' is in a vibrant blue color.

- Nested Loop Join

```
FOR EACH ROW <a> in TABLE T1 (matching conditions on T1)
```

```
    FOR EACH ROW <b> in TABLE T2 (matching condition on T2 given <a>)
```

```
        FOR EACH ROW <c> in TABLE T3 (matching conditions on T3 given <b>)
```

- FOR EACH is implemented using one of the 4 primitive data access methods in ndb
- Fetching each ROW is one network roundtrip between MySQL server and Data nodes.
 - NOTICE: Everything is done 1 row at a time. Zero parallelism!

Fast JOINS (AQL)



- MySQL Optimizer generates a query plan
 - Determines if ALL, NONE, or PARTs of the the Query can be pushed down.
 - Pushable parts of Query plan is sent to data node
- Data Node executes pushed down query
 - Parallezing if possible etc etc

And ?

severalnines

- ~20x Performance gain over MySQL Cluster 7.1 on typical three way joins.

Memcache



- The memcache API allows for memcache requests to be transformed into NDBAPI requests.
- Requires Memcached 1.6 to be installed
 - Memcached loads a plugin

Memcache



- Prime the MySQL Cluster with Ndb Memcache meta data:
- `mysql -uroot -p -h<hostname> < share/memcache-api/ndb_memcache_metadata.sql`

Memcache



- Download and install Memcache 1.6
 - http://memcached.googlecode.com/files/memcached-1.6.0_beta1.tar.gz
 - `./configure && make && sudo make install`
- Start Memcached with NDB Cluster plugin:
 - `bin/memcached -E lib/ndb_engine.so -e "connectstring=mgmd_host:1186;role=db-only" -vv -c 20`
 - `ndb_engine.so` comes in the cluster distro.

Memcache



- telnet localhost 11211
- set key 0 900 4
data
- SELECT * FROM ndbmemcachedb.demo_table;

Distributed GRANTS



- GRANTS are distributed to all MySQL Servers
 - Before: GRANT on each server

- **SPROCs/TRIGGERS/VIEWS are still NOT distributed**
 - Ref manual suggests you can do `ALTER TABLE mysql.proc ENGINE=NDB;`
 - VIEWS is no no.

- Groundwork:
 - Install a particular

Distributed GRANTS



- `mysql options -uroot < share/mysql/ndb_dist_priv.sql`
- Backup the old privilege tables
 - `mysqldump options -uroot mysql user db tables_priv columns_priv procs_priv > /backup/my_backup`
- `mysql> CALL mysql_cluster_move_privileges();`
- `mysql> CALL mysql_cluster_privileges_are_distributed();` → returns true if ok, false if failed.

Distributed GRANTS



- If you do a Initial Restart of MySQL Cluster...
 - Restore the backup
 - CALL mysql_cluster_move_privileges();
 - You could also rename the backup tables..

several**nines**

Q&A





Thank you for your time!

johan@severalnines.com