



Virtual Technology Summit

Hands-On Learning With Oracle and Community Experts

Where Technology and Community Meet



NoSQL and MySQL

Sheeri Cabral
Data Team Manager
Mozilla
November 18, 2014

Please Stand By. This session will begin promptly at the time indicated on the agenda. Thank You.



Copyright © 2014 Oracle and/or its affiliates. All rights reserved. |

Safe Harbor Statement

The following is intended to outline our general product direction. It is intended for information purposes only, and may not be incorporated into any contract. It is not a commitment to deliver any material, code, or functionality, and should not be relied upon in making purchasing decisions. The development, release, and timing of any features or functionality described for Oracle's products remains at the sole discretion of Oracle.



What is NoSQL?

- "Not Only SQL"
- Relational/SQL requires strict schemas
- Allows for more flexibility



What is NoSQL?

- Allows for simple CRUD
- e.g. get/set for key/value



What is NoSQL?

- Non-relational
 - Document
 - Key-Value store
 - Graph

MySQL and Hadoop Integration



- Custom scripts
- Apache Sqoop
 - Bi-directional
 - Batch

MySQL Applier for Hadoop



- Uni-directional
 - MySQL to Hadoop
 - Real-time, streaming
- Uses binary log

MySQL Applier for Hadoop



- Gets row data
- Translates to proper format
- Appends to HDFS file
 - Using libhdfs



Great 10-minute tutorial
on configuring, installing and using
MySQL Applier for Hadoop

<http://bit.ly/MySQLApplierHadoop>

How New is NoSQL?



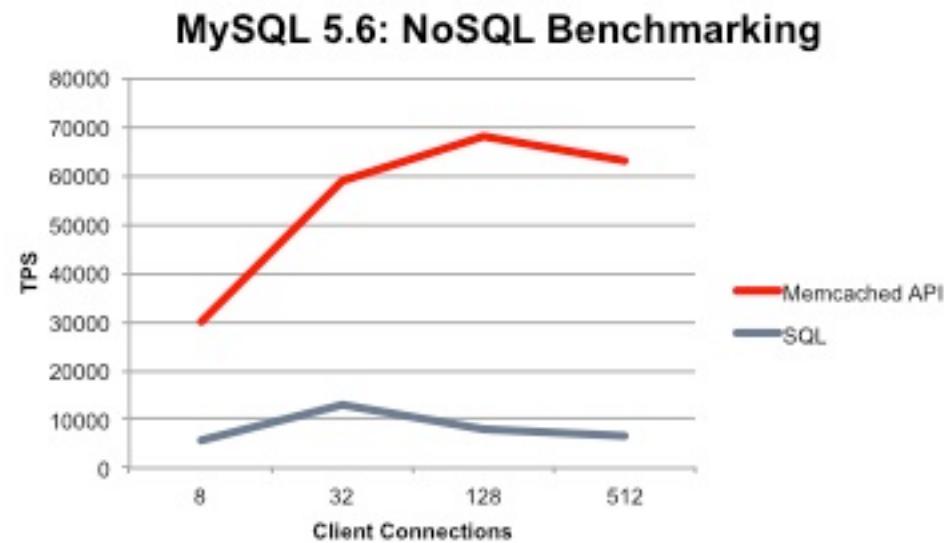
- memcached is a key-value store
 - In-memory only
 - Used for caching
 - Initial release in 2003

NoSQL memcached API



- Reads and writes
- MySQL 5.6
 - plugin
 - InnoDB
- MySQL Cluster 7.2
 - built-in

InnoDB NoSQL memcached API



InnoDB NoSQL memcached API



- Can use memcached API
 - Fast
 - Low latency
- Can use SQL
 - ACID compliance
 - Complex JOINs

Plugin Support



- Linux
- Solaris
- Mac OS X

InnoDB memcached Plugin



- Prerequisites:
 - libevent
 - Required by memcached
 - yum, apt-get, port-install

```
sudo yum install libevent
```

```
sudo apt-get install libevent-dev
```

InnoDB memcached Plugin



- Prerequisites:

- Schema
- innodb_memcached_config.sql
- Scripts/share directory

```
mysql < /usr/share/mysql/innodb_memcached_config.sql
```

Install the Plugin



- **INSTALL PLUGIN name SONAME "filename.so"**
 - File is in the plugin directory

```
SHOW GLOBAL VARIABLES LIKE 'plugin_dir';
```

```
mysql> install plugin daemon_memcached soname  
"libmemcached.so";
```



Configure the Plugin

- daemon_memcached_option
- In my.cnf or on mysqld commandline

```
daemon_memcached_option="-p11222"
```

Initial Configuration



- `daemon_memcached_lib_name`
 - Default `innodb_engine.so`
 - Used with memcached -E
- `daemon_memcached_lib_path`
 - Default `NULL = plugin_dir`

Initial Configuration



- `daemon_memcached_r_batch_size`
 - Commit after x reads
 - Default 1
- `daemon_memcached_w_batch_size`
 - Commit after x writes
 - Default is 1

Verify Things Work



- memcached on localhost:11211

```
telnet 127.0.0.1 11211
```

```
set a11 10 0 9
```

```
123456789
```

```
STORED
```

```
get a11
```

```
VALUE a11 0 9
```

```
123456789
```

```
END
```

```
quit
```



Verify in MySQL

- test.demo_test by default

```
mysql> select * from test.demo_test;
```

More Configuration



- innodb_memcache schema
- Created by innodb_memcached_config.sql
- 3 tables
 - containers
 - cache_policies
 - config_options

containers Table Requirements



- db_schema – default test
- db_name – default demo_test
- key_columns – lookup field
- unique_idx_name_on_key – unique index name for key
 - Preferably primary key
- value_columns – field1 or field1|field2|...

containers Table Optional



- Set to 0 if unused, otherwise plugin will fail
- flags – fields for memcached flags
- expiration_time_column – exp value for memcached
- cas_column – compare-and-swap value for memcached

containers Table Optional



- Set to 0 if unused, otherwise plugin will fail
- `expiration_time_column`
 - exp value for memcached
- `cas_column`
 - compare-and-swap value for memcached

containers Table Optional



- Set to 0 if unused, otherwise plugin will fail
- flags
 - fields for memcached flags
- unique_idx_name_on_key
 - unique index name for key

cache_policies



- Read at startup
- Uninstall and reinstall plugin when changed
- innodb_only
- cache_only
- caching

cache_policies



- get
- set
- delete
- flush

config_options



- separator
 - default |
 - Same as value_columns separator
- table_map_delimiter
 - @@schema.key_name

config_options



- separator
 - default |
 - Same as value_columns separator
- table_map_delimiter
 - @@schema.key_name

Access



- Underlying table, with SQL
- Via memcached:
 - memcp
 - memcat

Access Example



```
MEMCACHED_SERVERS=127.0.0.1:11211
```

```
memcp /etc/apache2/mime.types
```

```
memcat mime.types
```

memcached and MySQL Cluster



- Built-in
- Auto-sharding
- Writes committed directly to Cluster
 - No cache invalidation
 - No extra sync needed

Schema or No Schema?



- No Schema
 - Key/value written to one row
- Schema
 - Can use prefix to specify table field
 - Online schema change

No Prerequisites



- libevent already installed
- No schema needed

memcached Configuration



- **Engine (-E)**
 - E ndb_engine.so
 - Innodb was innodb_engine.so
 - **Options (-e)**
 - e "connectstring=sheeri:1186;role=dev"

memcached Configuration



- Worker threads (-t)
- TCP port (-p)
 - Default 11211
- UDP port (-U)
 - Default 11211
 - 0 means disable UDP

NDB Configuration



- In a schema
 - ndbmemcache
 - ndb_memcache_metadata.sql
- 6 tables

ndbmempool Tables



- meta
 - Version
 - Read-only
- ndb_clusters
 - cluster-id
 - Connection string
 - microsec_rtt column – auto tuning

cache_policies Table



- get_policy
- set_policy
- delete_policy
 - cache_only
 - ndb_only
 - caching
 - disabled

cache_policies Table



- flush_from_db
 - True
 - False

ndbmempool Tables



- containers
 - Similar to InnoDB plugin
 - Defines what table to use
- Complex key mapping tables
 - memcache_server_roles
 - key_prefixes

ndbmemcache Issues



- **AUTO_INCREMENT**
 - Use a generated sequence instead
- NDB schema changes
 - Restart memcached when schema changes

Q+A

ORACLE®

Copyright © 2014 Oracle and/or its affiliates. All rights reserved. |

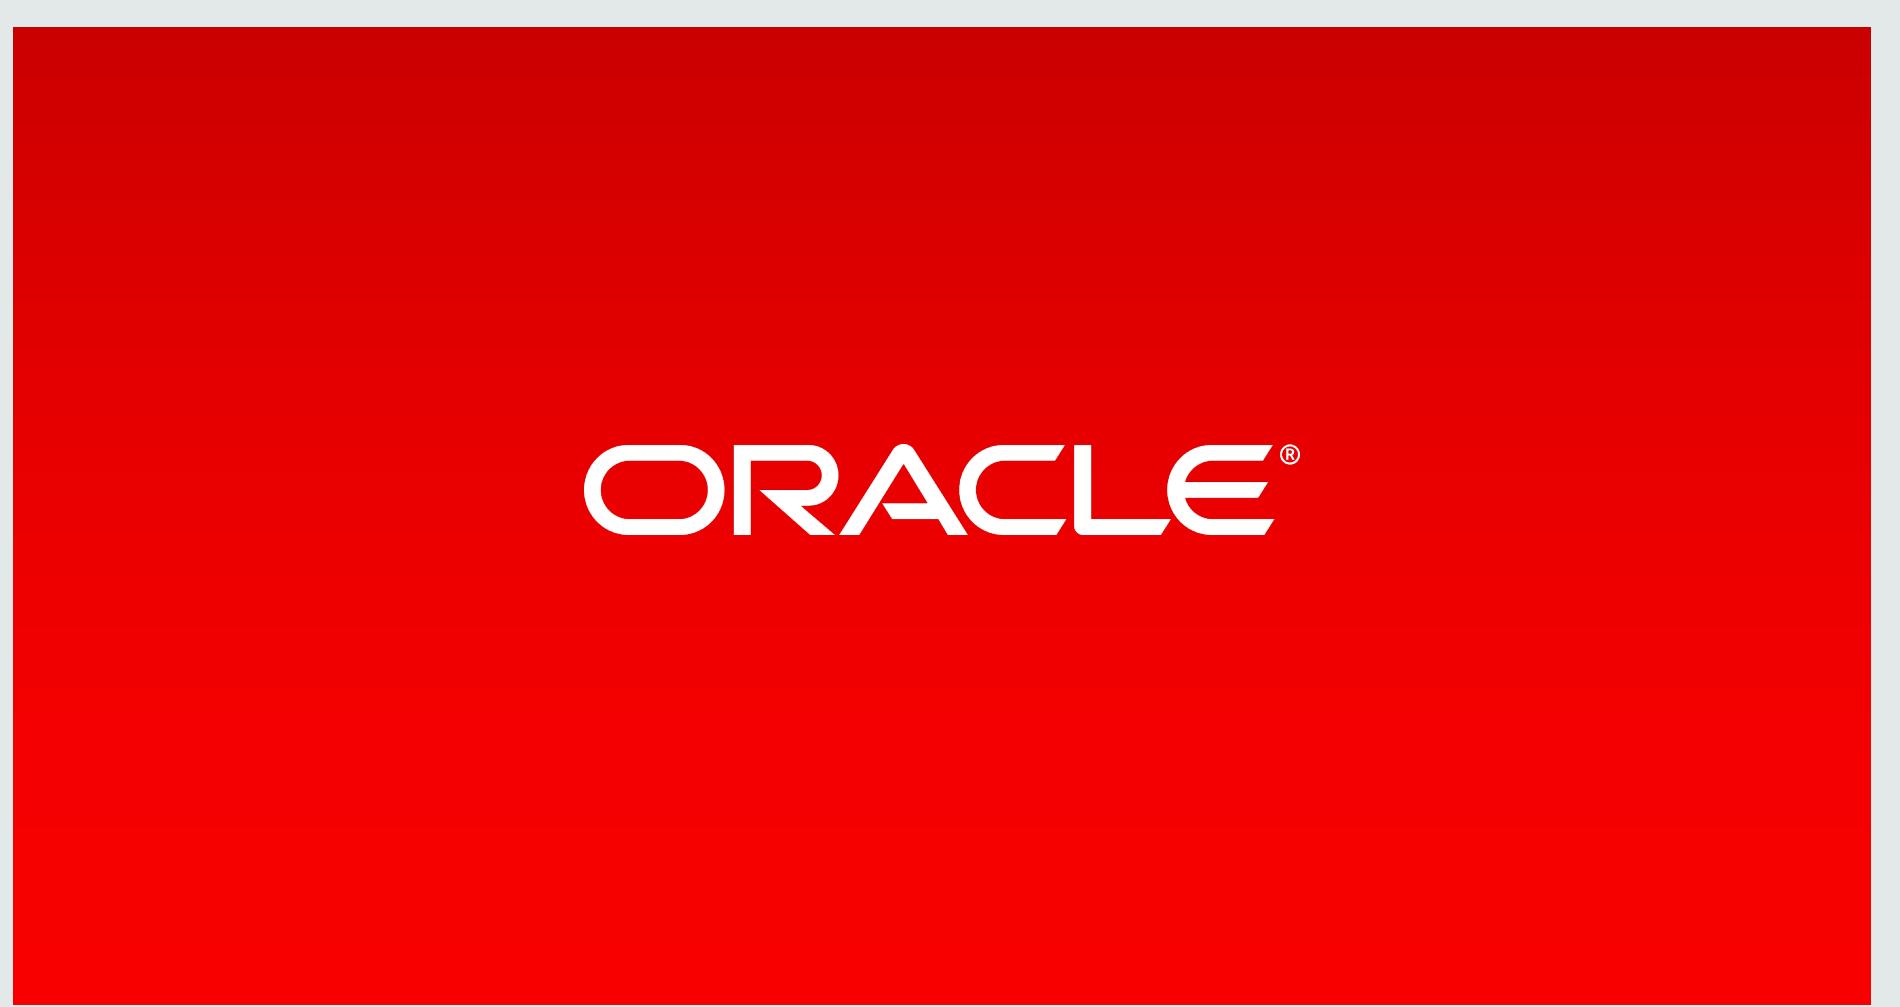


Thank You for Joining Us Today



Hardware and Software Engineered to Work Together





ORACLE®