

# Gora Cassandra Module

## Table of contents

1 Overview .....	2
2 gora.properties .....	2
3 Gora Cassandra mappings .....	2

## 1 Overview

This is the main documentation for the gora-cassandra module. gora-cassandra module enables [Apache Cassandra](#) backend support for Gora.

## 2 gora.properties

To configure the CassandraStore one would typically set the following:  
`gora.cassandra.servers=localhost:9160` - This value should specify the host:port for a running Cassandra server or node. In this case the server happens to be running on localhost at port 9160 which is the default Cassandra server configuration.

## 3 Gora Cassandra mappings

Say we wished to map some Employee data and store it into the CassandraStore.

```
<gora-orm>
  <keyspace name="Employee" host="localhost" cluster="Gora Cassandra Test Cluster">
    <family name="p" />
    <family name="f" />
    <family name="sc" type="super" />
  </keyspace>

  <class name="org.apache.gora.examples.generated.Employee" keyClass="java.lang.String"
  keyspace="Employee">
    <field name="name" family="p" qualifier="info:nm"/>
    <field name="dateOfBirth" family="p" qualifier="info:db"/>
    <field name="ssn" family="p" qualifier="info:sn"/>
    <field name="salary" family="p" qualifier="info:sl"/>
  </class>
  ...
</gora-orm>
```

Here you can see that within the `keyspace` section we specify the Cassandra keyspace schema name `Employee`, the host `localhost` and the cluster name `Gora Cassandra Test Cluster`. Additionally we are required to specify the type of column families we wish to create within Cassandra. In this case we create three columns; `p`, `f` and `sc` the last of which is further defined as a super column.

Further to this, within the Cassandra mappings we are required to specify persistent fields which values should map to. In the above example we see that within the `class` element the generated class name `org.apache.gora.examples.generated.Employee` specifies the Avro generated persistent class, the `keyClass java.lang.String` which specifies the keys which map to the field values and finally the `keyspace Employee` which matches to the above keyspace definition.

Finally fields which are to be persisted into Cassandra need to be configured such that they receive a name (name, dateOfBirth, ssn and salary respectively), the column family to which they belong (all p in this case) and additionally a qualifier, which enables more granular control over the data to be persisted into Cassandra.