

# Table of Contents

<b>Chapter 1: Graphic bundle</b>	<b>1</b>
<b>Chapter 1: Introduction to Big Data and MySQL 8</b>	<b>1</b>
<b>Chapter 2: Data Query Techniques in MySQL 8</b>	<b>6</b>
<b>Chapter 3: Indexing your data for High-Performing Queries</b>	<b>8</b>
<b>Chapter 4: Using Memcached with MySQL 8</b>	<b>11</b>
<b>Chapter 5: Partitioning High Volume Data</b>	<b>12</b>
<b>Chapter 6: Replication for building highly available solutions</b>	<b>17</b>
<b>Chapter 7: MySQL 8 Best Practices</b>	<b>20</b>
<b>Chapter 8: NoSQL API for Integrating with Big Data Solutions</b>	<b>24</b>
<b>Chapter 9: Case study: Part I - Apache Sqoop for exchanging data between MySQL and Hadoop</b>	<b>25</b>
<b>Chapter 10: Case study: Part II - Real time event processing using MySQL applier</b>	<b>30</b>
<b>Index</b>	<b>34</b>

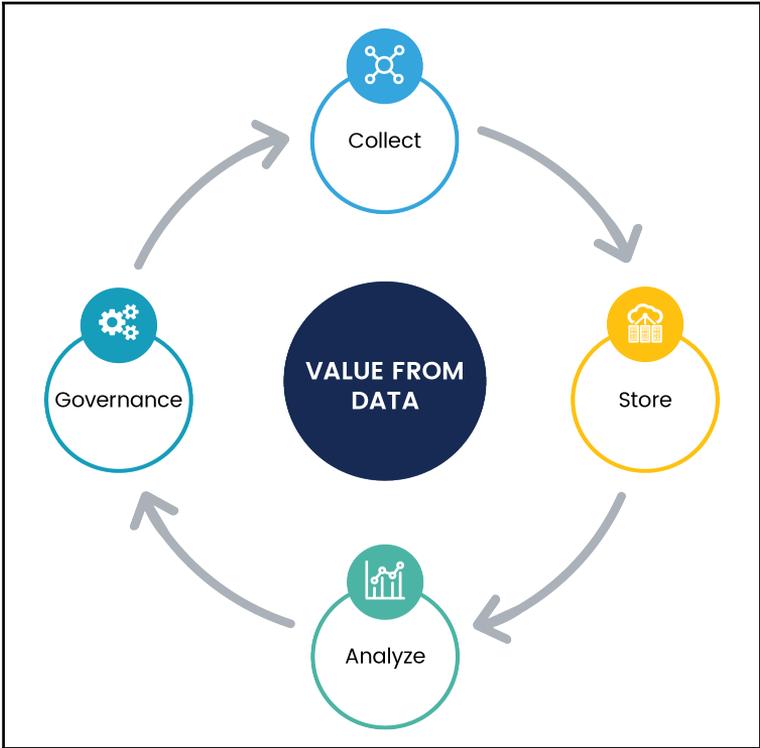
# 1

## Graphic bundle

**Chapter 1: Introduction to Big Data and  
MySQL 8**









```
[root@ip-172-31-43-204 tmp]# rpm -ivh mysql-community-common-8.0.1-0.1.dmr.el6.x86_64.rpm
warning: mysql-community-common-8.0.1-0.1.dmr.el6.x86_64.rpm: Header V3 DSA/SHA1 Signature, key ID 5072e1f5: NOKEY
Preparing...
##### [100%]
Updating / installing...
 1:mysql-community-common-8.0.1-0.1.##### [100%]
[root@ip-172-31-43-204 tmp]# rpm -ivh mysql-community-libs-8.0.1-0.1.dmr.el6.x86_64.rpm
warning: mysql-community-libs-8.0.1-0.1.dmr.el6.x86_64.rpm: Header V3 DSA/SHA1 Signature, key ID 5072e1f5: NOKEY
Preparing...
##### [100%]
Updating / installing...
 1:mysql-community-libs-8.0.1-0.1.dm##### [100%]
[root@ip-172-31-43-204 tmp]# rpm -ivh mysql-community-client-8.0.1-0.1.dmr.el6.x86_64.rpm
warning: mysql-community-client-8.0.1-0.1.dmr.el6.x86_64.rpm: Header V3 DSA/SHA1 Signature, key ID 5072e1f5: NOKEY
Preparing...
##### [100%]
Updating / installing...
 1:mysql-community-client-8.0.1-0.1.##### [100%]
[root@ip-172-31-43-204 tmp]# rpm -ivh mysql-community-server-8.0.1-0.1.dmr.el6.x86_64.rpm
warning: mysql-community-server-8.0.1-0.1.dmr.el6.x86_64.rpm: Header V3 DSA/SHA1 Signature, key ID 5072e1f5: NOKEY
Preparing...
##### [100%]
Updating / installing...
 1:mysql-community-server-8.0.1-0.1.##### [100%]
```

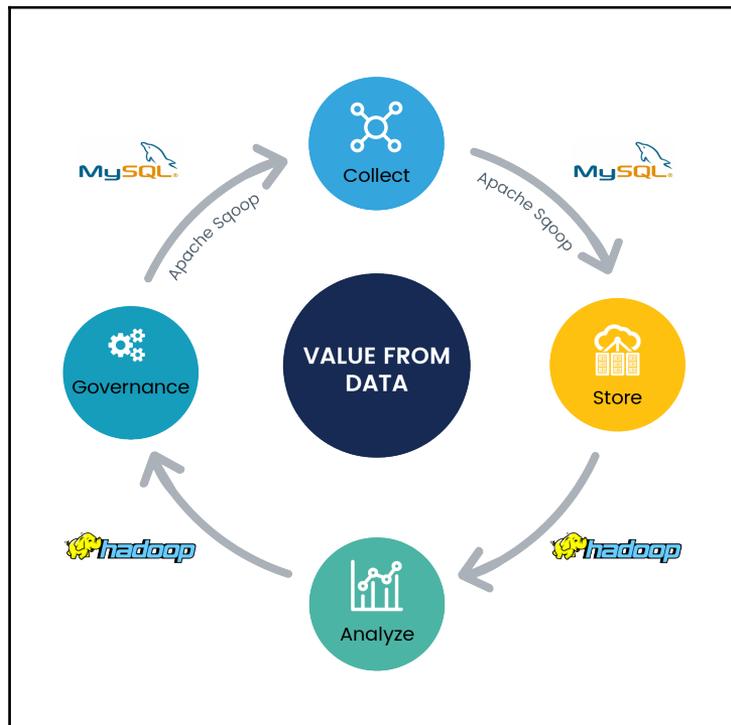
```
[root@ip-172-31-43-204 tmp]# mysql --version
mysql Ver 8.0.1-dmr for Linux on x86_64 (MySQL Community Server (GPL))
[root@ip-172-31-43-204 tmp]#
```

```
[root@ip-172-31-43-204 tmp]# mysql -u root -p
Enter password:
Welcome to the MySQL monitor.  Commands end with ; or \g.
Your MySQL connection id is 8
Server version: 8.0.1-dmr

Copyright (c) 2000, 2017, Oracle and/or its affiliates. All rights reserved.

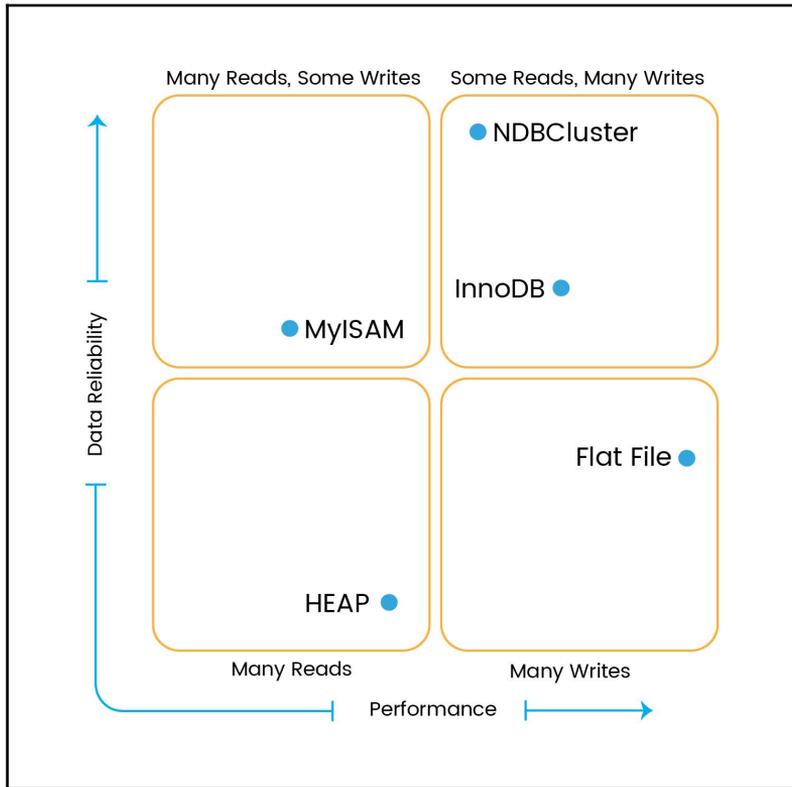
Oracle is a registered trademark of Oracle Corporation and/or its
affiliates. Other names may be trademarks of their respective
owners.

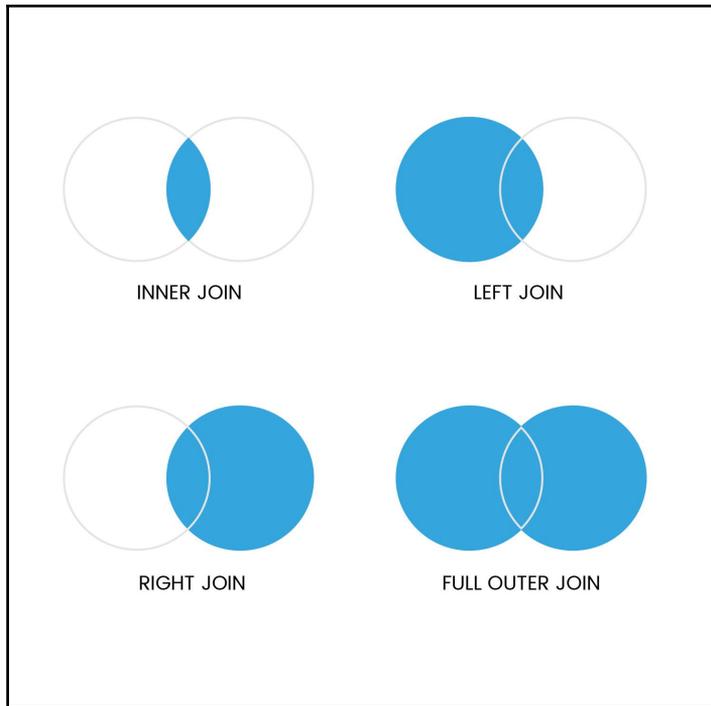
Type 'help;' or '\h' for help. Type '\c' to clear the current input statement.
mysql>
```



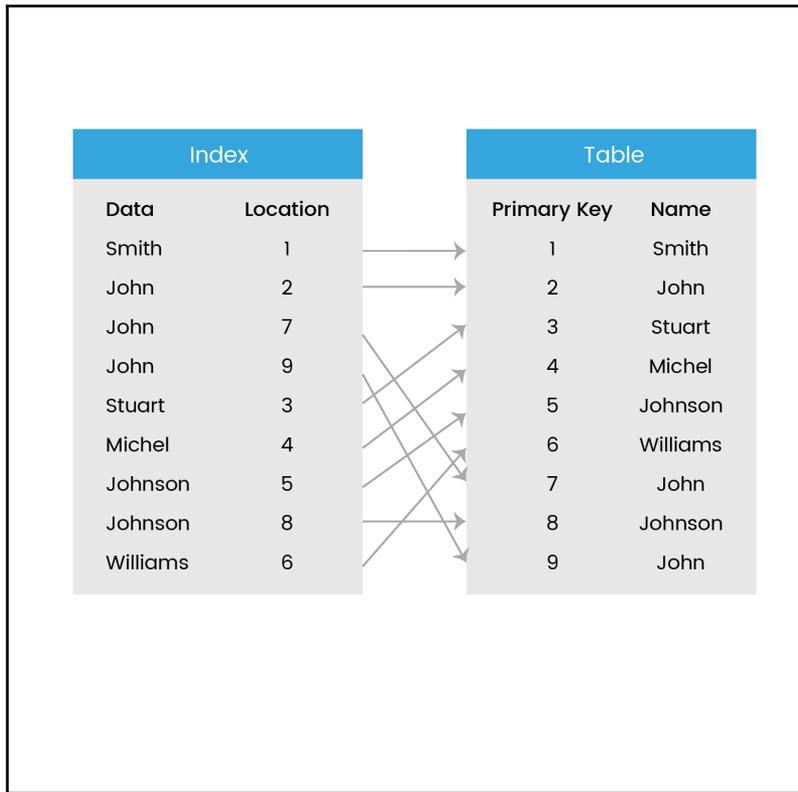
# Chapter 2: Data Query Techniques in MySQL

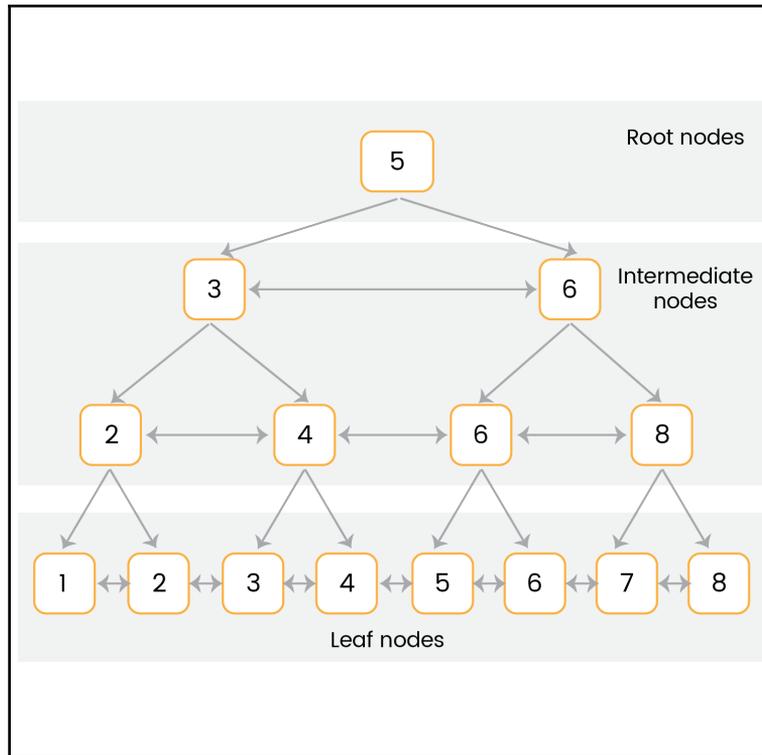
## 8

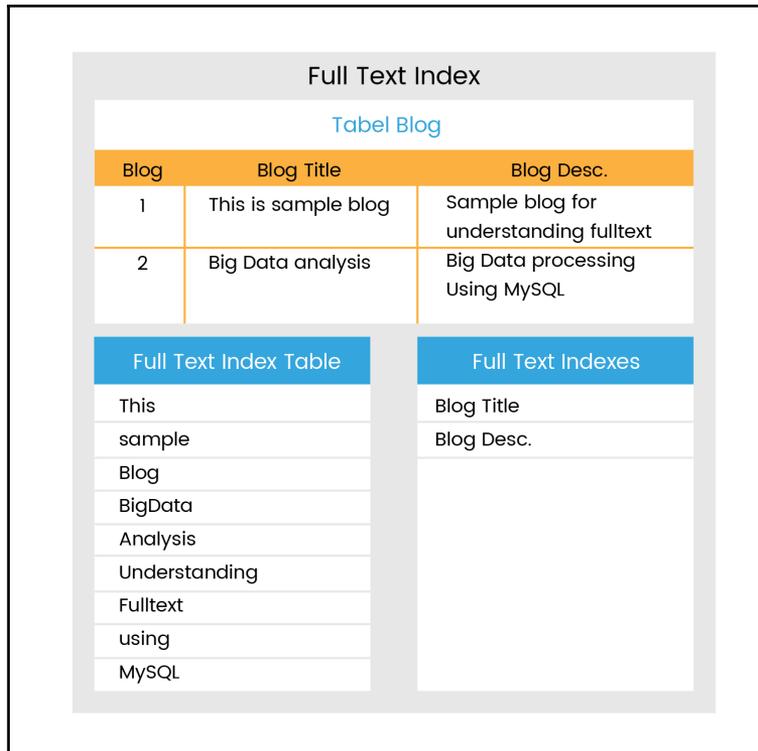




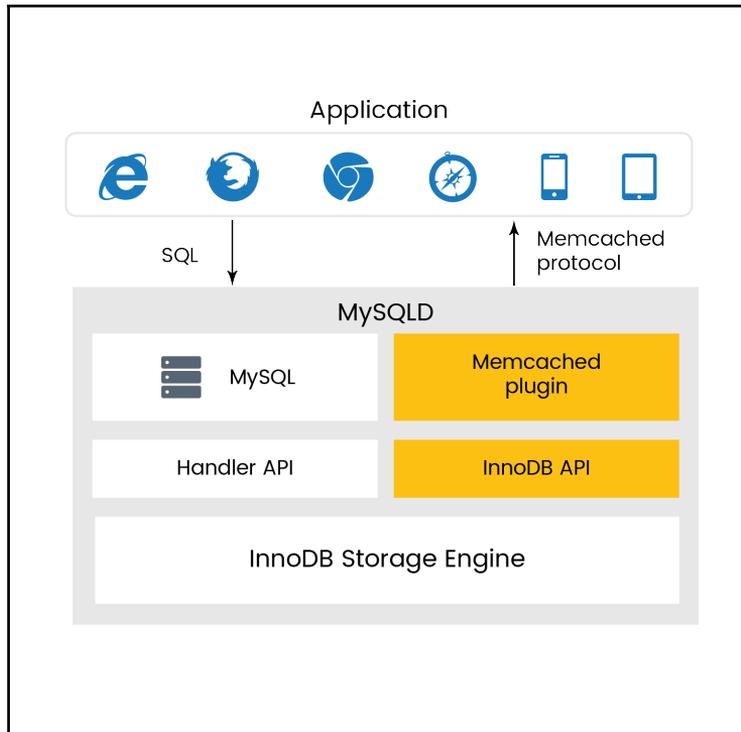
## Chapter 3: Indexing your data for High-Performing Queries







## Chapter 4: Using Memcached with MySQL 8

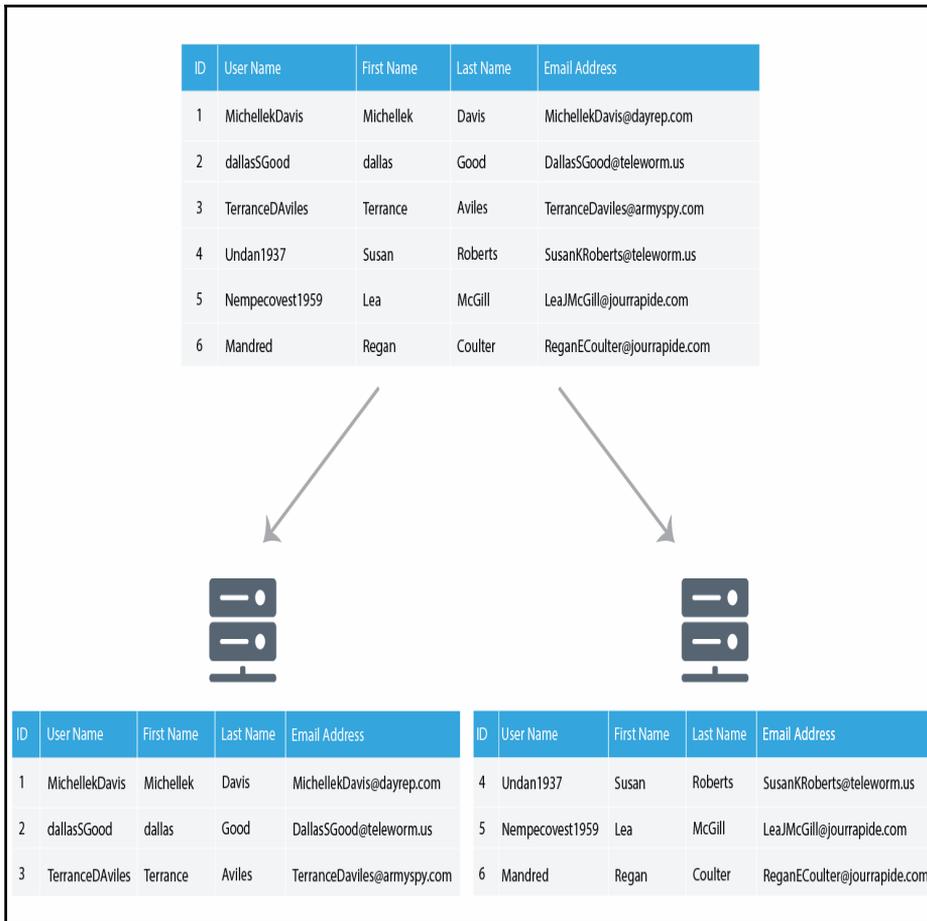


```
[root@ip-172-31-22-59 ~]# telnet localhost 11211
Trying 127.0.0.1...
Connected to localhost.
Escape character is '^]'.

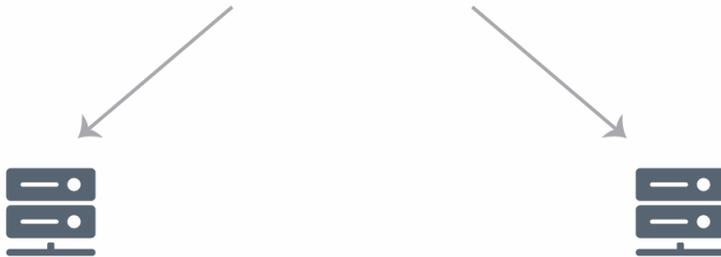
```

```
mysql> select * from test.student_result;
+-----+-----+-----+-----+-----+-----+-----+-----+
| student_id | student_name | maths | english | history | flags | cas | expiry |
+-----+-----+-----+-----+-----+-----+-----+-----+
|          101 | Jaydip      | 50    | 60     | 70     | 0     | 2   | 0     |
+-----+-----+-----+-----+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

## Chapter 5: Partitioning High Volume Data



ID	User Name	First Name	Last Name	Email Address	Phone Number	Birth Date
1	MichellekDavis	Michellek	Davis	MichellekDavis@dayrep.com	9412233379	10/10/2086
2	dallasSGood	dallas	Good	DallasSGood@teleworm.us	8329843007	30/03/1972
3	TerranceDAviles	Terrance	Aviles	TerranceDaviles@armyspy.com	7045076659	20/04/1992
4	Undan1937	Susan	Roberts	SusanKRoberts@teleworm.us	4014126825	25/12/1996
5	Nempecovest1959	Lea	McGill	LeaJMcGill@jourrapide.com	2176567191	3/6/2000
6	Mandred	Regan	Coulter	ReganECoulter@jourrapide.com	6304832993	9/3/1998
7	Begivaing	Richard	Endres	RichardJEndres@rhyrep.com	3096848793	12/8/1988
8	Biturnight	Gregory	Melton	GregoryMMelton@dayrep.com	6304236158	6/9/1977
9	Stiong	Bernice	Stevenson	BerniceGStevenson@rhyrep.com	7736461308	3/1/2001



ID	User Name	First Name	Last Name	Email Address
1	MichellekDavis	Michellek	Davis	MichellekDavis@dayrep.com
2	dallasSGood	dallas	Good	DallasSGood@teleworm.us
3	TerranceDAviles	Terrance	Aviles	TerranceDaviles@armyspy.com
4	Undan1937	Susan	Roberts	SusanKRoberts@teleworm.us
5	Nempecovest1959	Lea	McGill	LeaJMcGill@jourrapide.com
6	Mandred	Regan	Coulter	ReganECoulter@jourrapide.com
7	Begivaing	Richard	Endres	RichardJEndres@rhyrep.com
8	Biturnight	Gregory	Melton	GregoryMMelton@dayrep.com
9	Stiong	Bernice	Stevenson	BerniceGStevenson@rhyrep.com

Phone Number	Birth Date
9412233379	10/10/2086
8329843007	30/03/1972
7045076659	20/04/1992
4014126825	25/12/1996
2176567191	3/6/2000
6304832993	9/3/1998
3096848793	12/8/1988
6304236158	6/9/1977
7736461308	3/1/2001

```
mysql> SELECT * FROM student PARTITION (p2);
+-----+-----+-----+-----+
| rollNo | name   | class | marks |
+-----+-----+-----+-----+
|      1 | Jaydip | A     |    80 |
|      4 | Malav  | B     |    84 |
+-----+-----+-----+-----+
2 rows in set (0.00 sec)
```

```
mysql> SELECT * FROM student PARTITION (p1,p2);
+-----+-----+-----+-----+
| rollNo | name   | class | marks |
+-----+-----+-----+-----+
|      3 | Palak  | A     |    45 |
|      6 | Kandarp | B     |    60 |
|      1 | Jaydip | A     |    80 |
|      4 | Malav  | B     |    84 |
+-----+-----+-----+-----+
4 rows in set (0.00 sec)
```

```
mysql> SELECT * FROM student1 PARTITION(p0);
+----+-----+
| id | name |
+----+-----+
|  1 | Jaydip |
+----+-----+
1 row in set (0.00 sec)
```

```
mysql> SELECT * FROM student WHERE class = "A";
+-----+-----+-----+-----+
| rollNo | name   | class | marks |
+-----+-----+-----+-----+
|      3 | Palak  | A     |    45 |
|      1 | Jaydip | A     |    80 |
+-----+-----+-----+-----+
2 rows in set (0.00 sec)
```

```
mysql> SELECT * FROM student WHERE class = "A";
+-----+-----+-----+-----+
| rollNo | name   | class | marks |
+-----+-----+-----+-----+
|      3 | Palak | A     |    45 |
|      1 | Jaydip| A     |    80 |
+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> DELETE FROM student PARTITION(p2) WHERE class = "A";
Query OK, 1 row affected (0.01 sec)

mysql> SELECT * FROM student WHERE class = "A";
+-----+-----+-----+-----+
| rollNo | name   | class | marks |
+-----+-----+-----+-----+
|      3 | Palak | A     |    45 |
+-----+-----+-----+-----+
1 row in set (0.00 sec)
```

```
mysql> SELECT * FROM student PARTITION(p1);
+-----+-----+-----+-----+
| rollNo | name   | class | marks |
+-----+-----+-----+-----+
|      3 | Palak | A     |    45 |
|      6 | Kandarp| C     |    60 |
+-----+-----+-----+-----+
2 rows in set (0.00 sec)

mysql> UPDATE student PARTITION(P1) SET class = 'A' WHERE marks > 50;
Query OK, 1 row affected (0.00 sec)
Rows matched: 1  Changed: 1  Warnings: 0

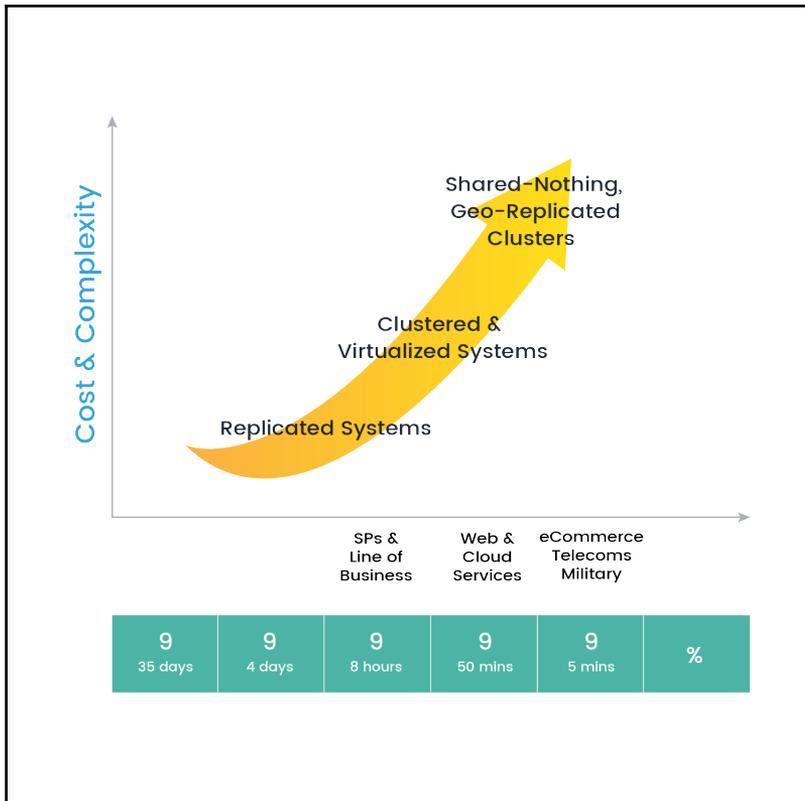
mysql> SELECT * FROM student PARTITION(p1);
+-----+-----+-----+-----+
| rollNo | name   | class | marks |
+-----+-----+-----+-----+
|      3 | Palak | A     |    45 |
|      6 | Kandarp| A     |    60 |
+-----+-----+-----+-----+
2 rows in set (0.01 sec)
```

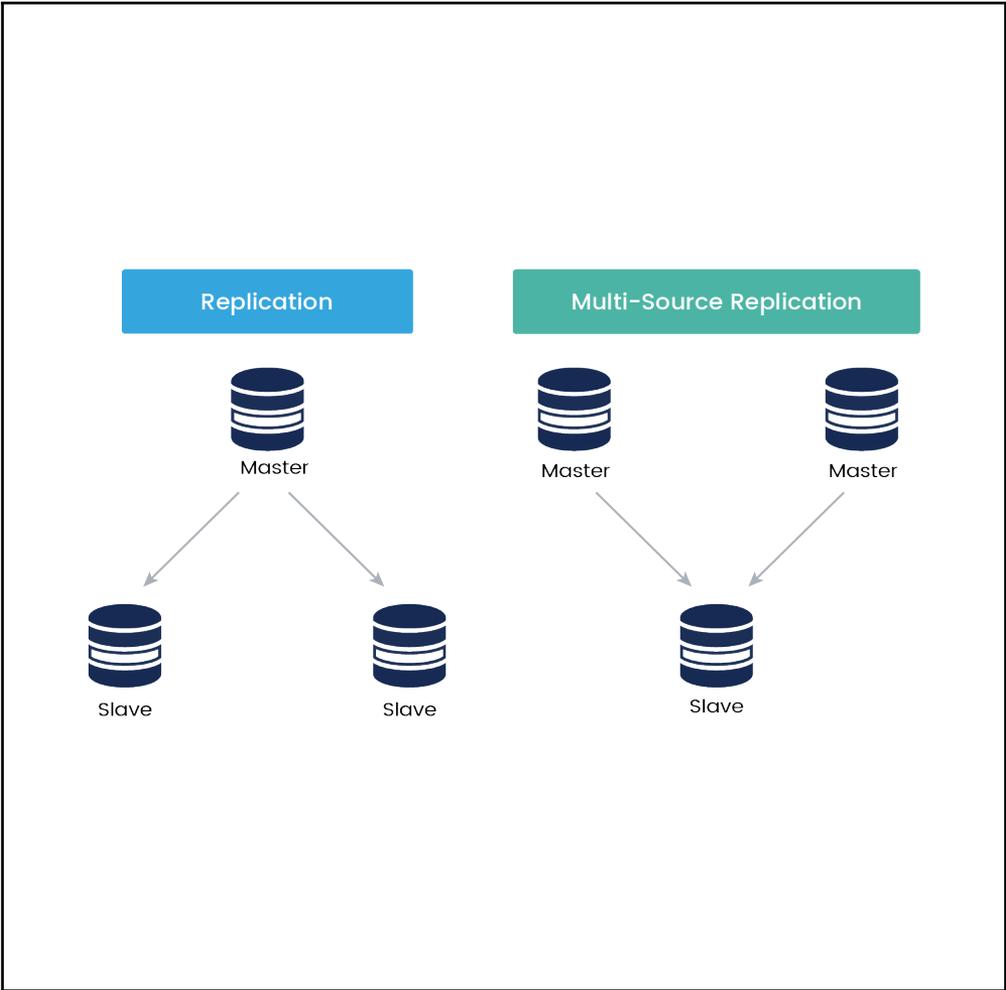
```
mysql> SELECT * FROM student PARTITION(p2);
+-----+-----+-----+-----+
| rollNo | name  | class | marks |
+-----+-----+-----+-----+
|      4 | Malav | B     |    84 |
+-----+-----+-----+-----+
1 row in set (0.00 sec)

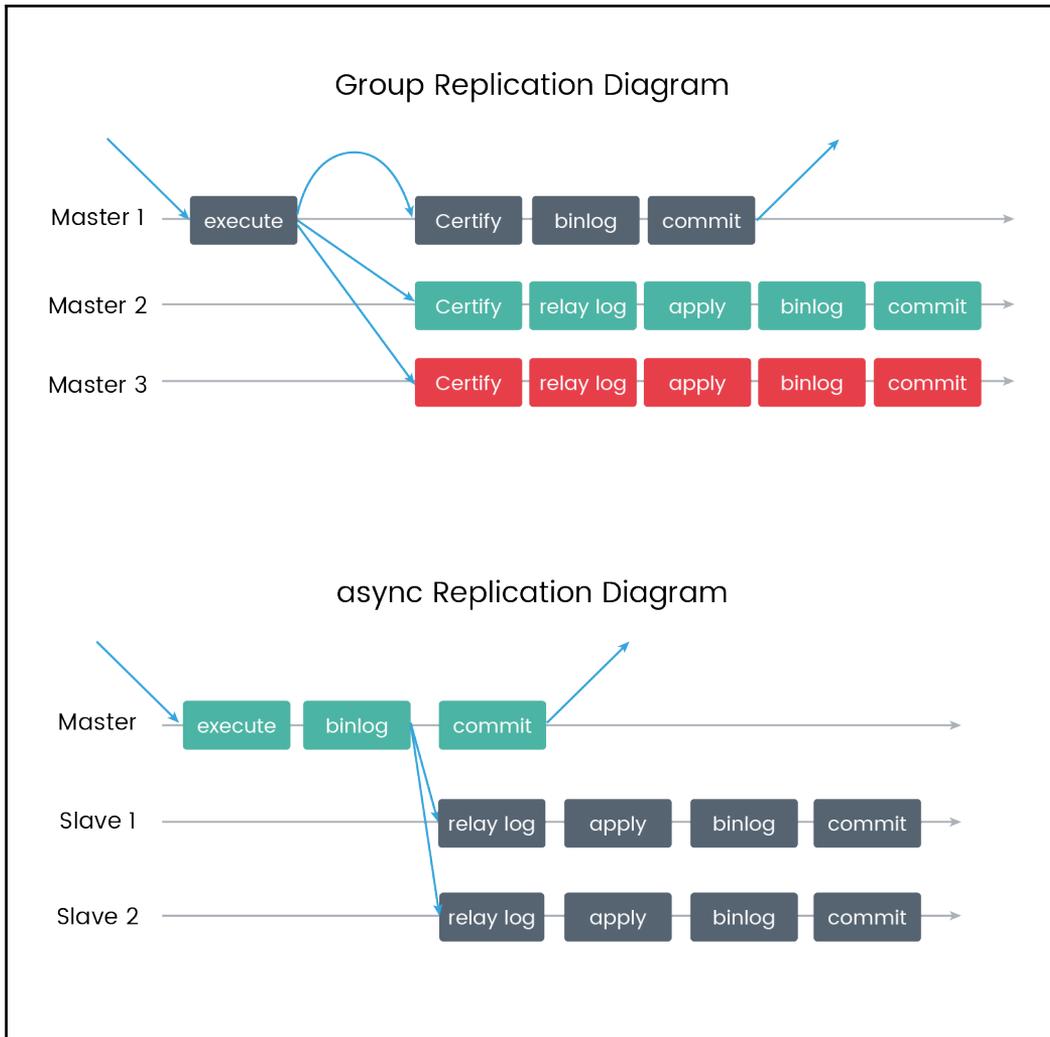
mysql> INSERT INTO student PARTITION(P2) VALUES (7,'Jaydip','A',30);
ERROR 1748 (HY000): Found a row not matching the given partition set
mysql> INSERT INTO student PARTITION(P2) VALUES (7,'Jaydip','A',70);
Query OK, 1 row affected (0.01 sec)

mysql> SELECT * FROM student PARTITION(p2);
+-----+-----+-----+-----+
| rollNo | name  | class | marks |
+-----+-----+-----+-----+
|      4 | Malav | B     |    84 |
|      7 | Jaydip | A     |    70 |
+-----+-----+-----+-----+
2 rows in set (0.00 sec)
```

## Chapter 6: Replication for building highly available solutions

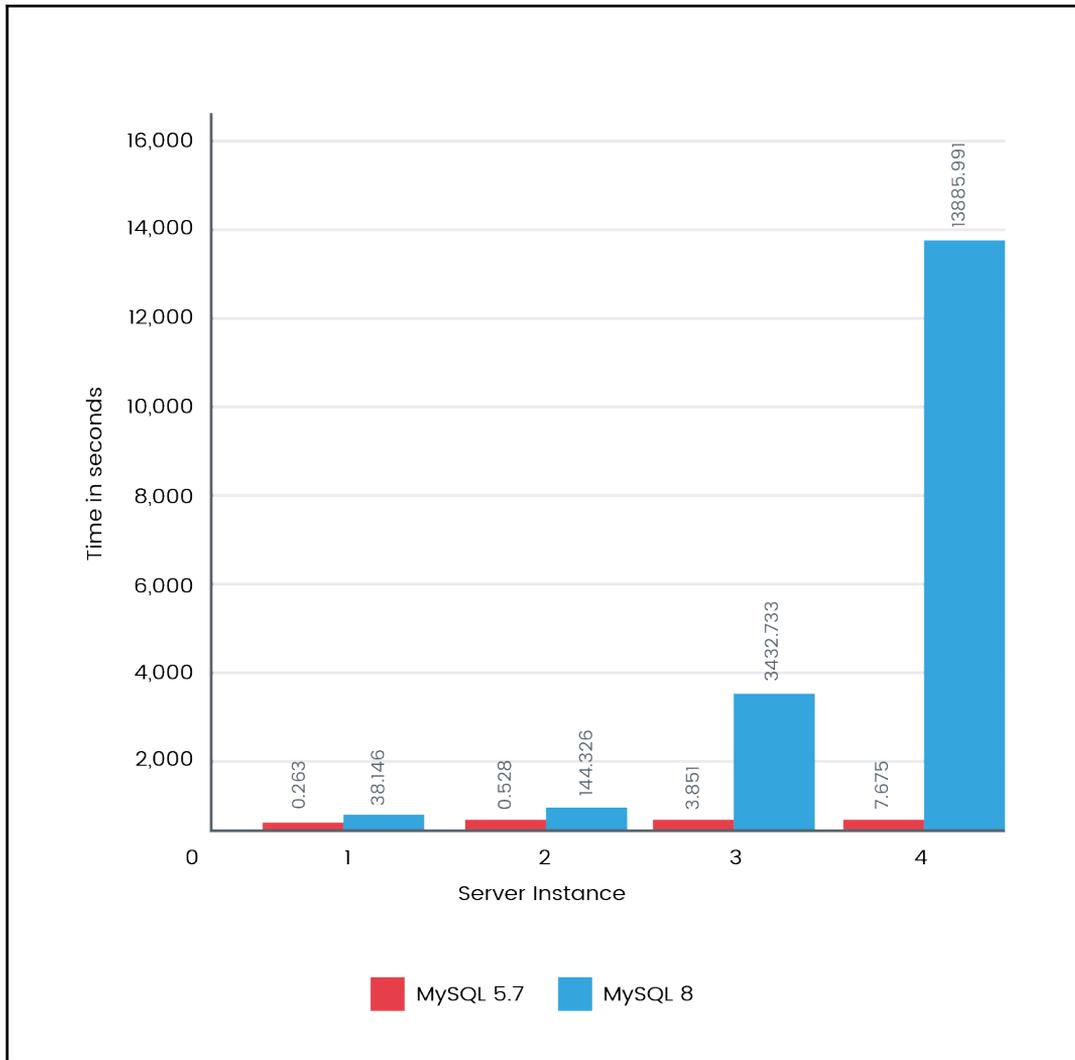


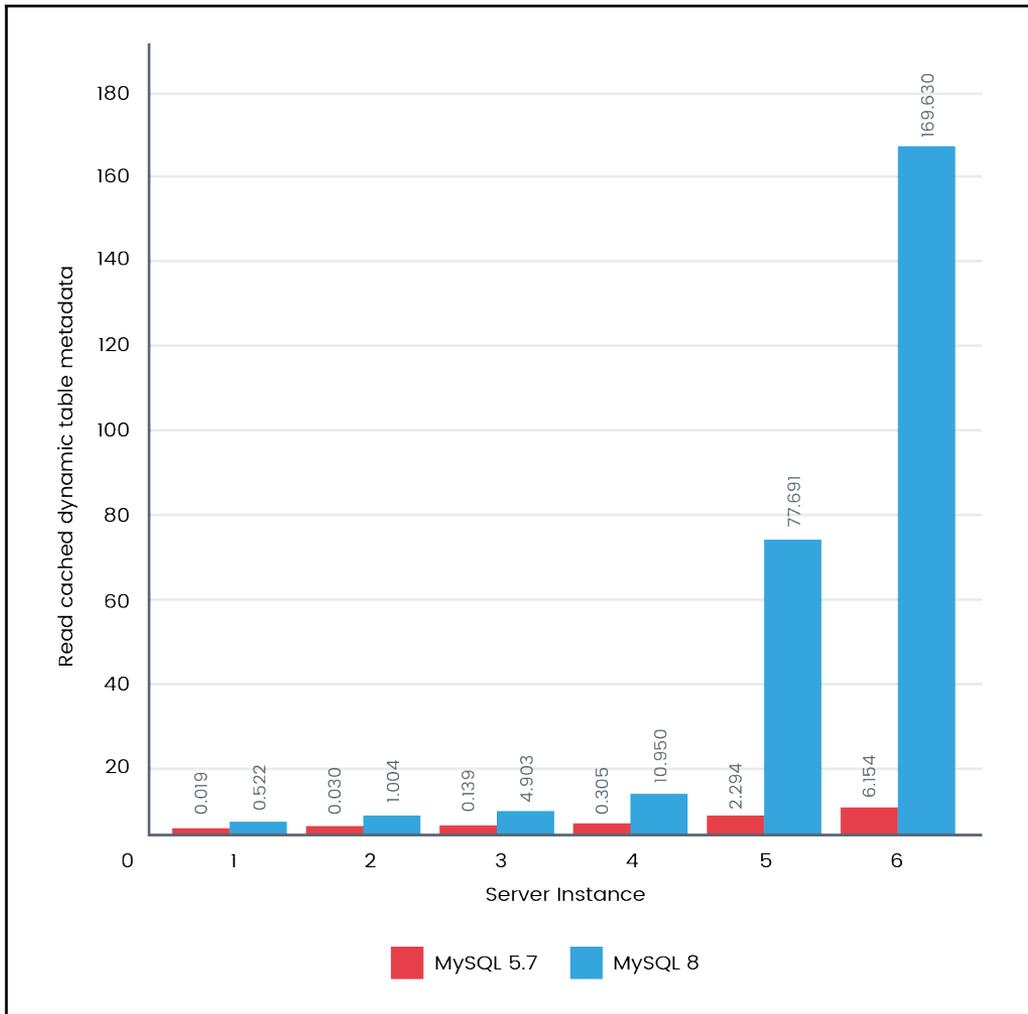


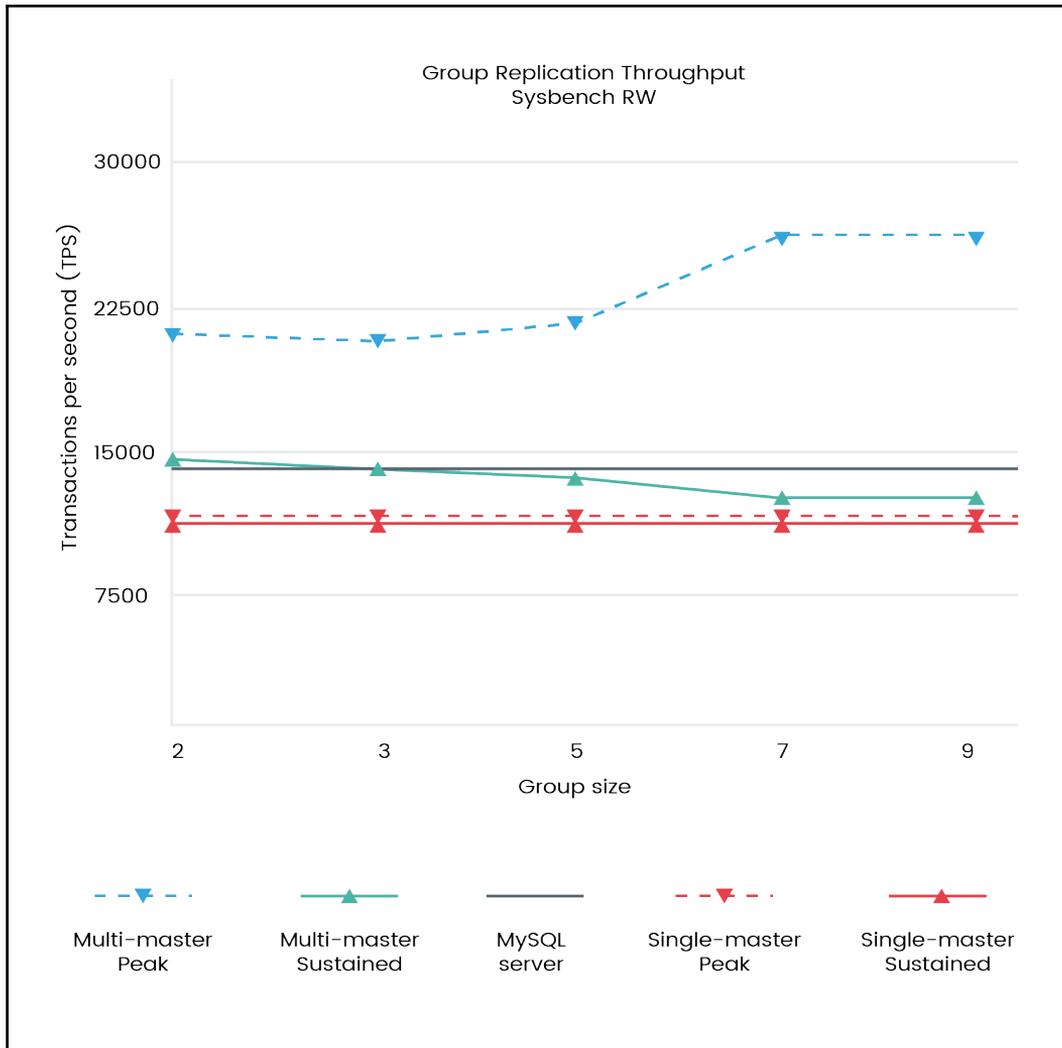


## Chapter 7: MySQL 8 Best Practices

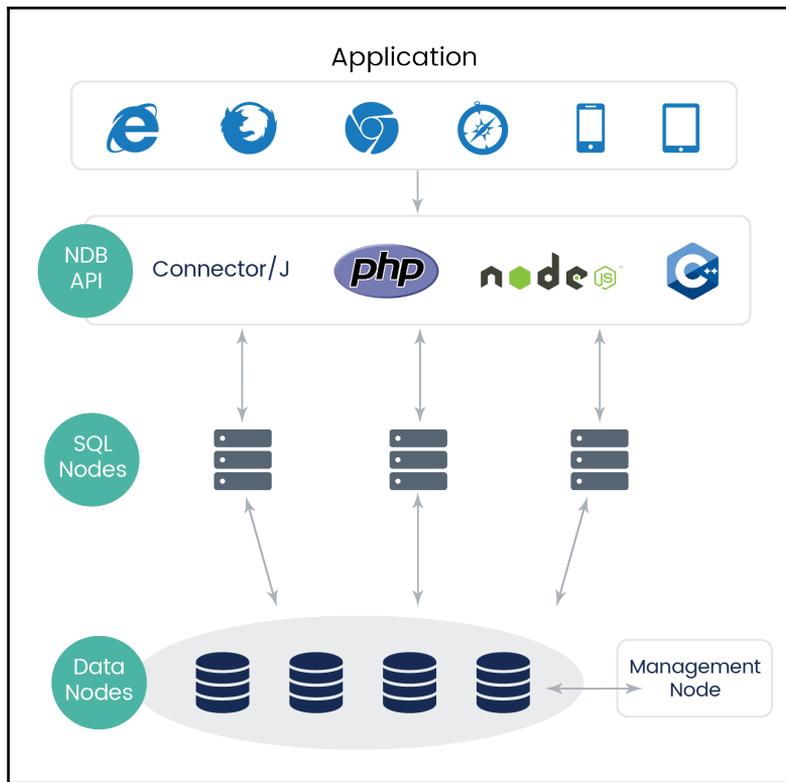




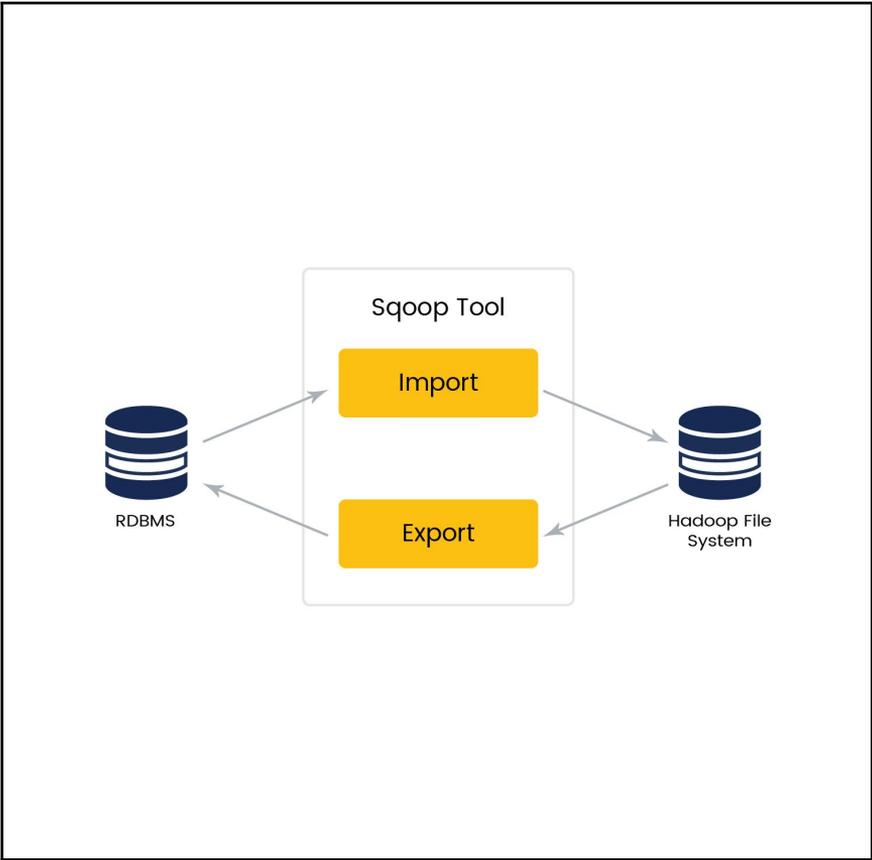


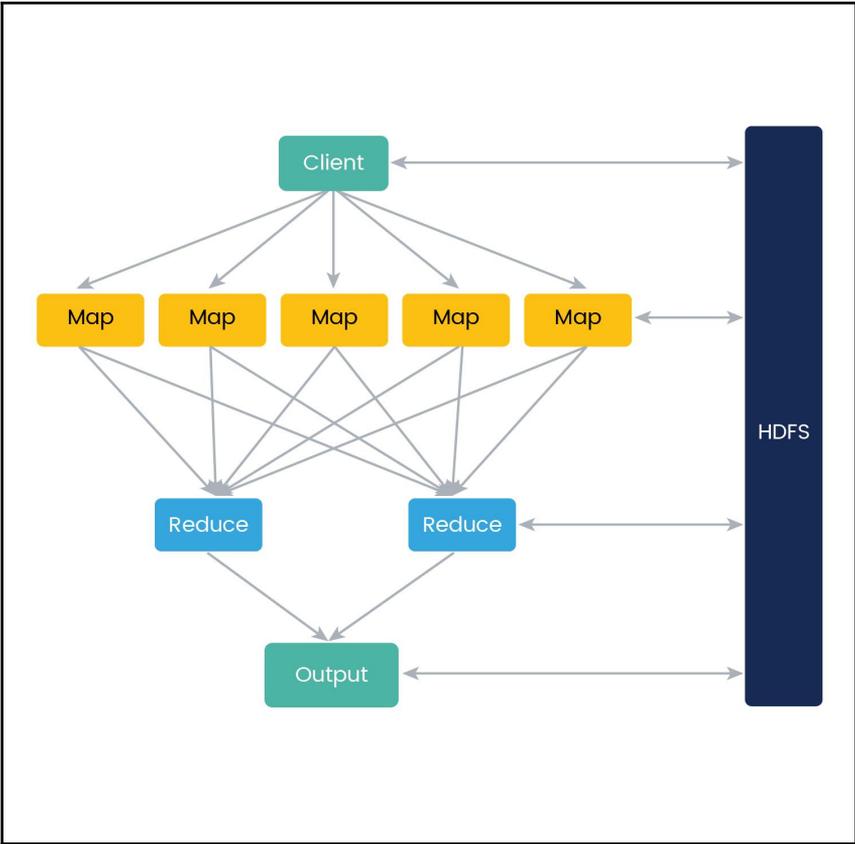


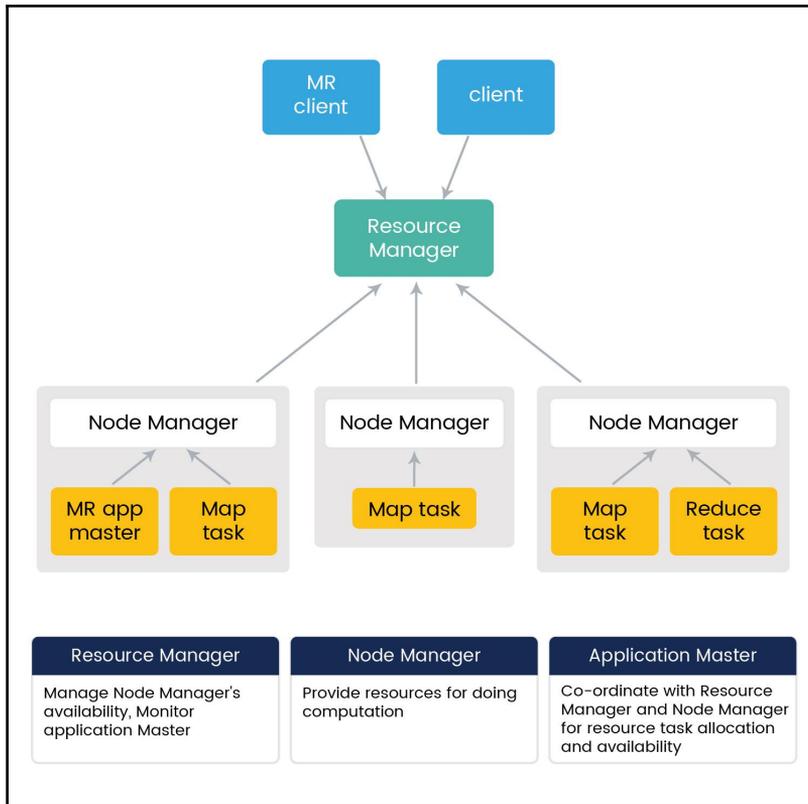
## Chapter 8: NoSQL API for Integrating with Big Data Solutions



## Chapter 9: Case study: Part I - Apache Sqoop for exchanging data between MySQL and Hadoop







```

[root@ip-172-31-22-59 local]# wget http://www-us.apache.org/dist/hadoop/common/h
adoop-2.6.5/hadoop-2.6.5.tar.gz
--2017-10-02 07:25:51-- http://www-us.apache.org/dist/hadoop/common/hadoop-2.6.
5/hadoop-2.6.5.tar.gz
Resolving www-us.apache.org (www-us.apache.org)... 140.211.11.105
Connecting to www-us.apache.org (www-us.apache.org)|140.211.11.105|:80... connec
ted.
HTTP request sent, awaiting response... 200 OK
Length: 199635269 (190M) [application/x-gzip]
Saving to: 'hadoop-2.6.5.tar.gz'

hadoop-2.6.5.tar.gz 100%[=====>] 190.39M  7.93MB/s   in 14s

2017-10-02 07:26:06 (13.3 MB/s) - 'hadoop-2.6.5.tar.gz' saved [199635269/1996352
69]

[root@ip-172-31-22-59 local]#
  
```

```

[root@ip-172-31-22-59 local]# mv hadoop-2.6.5 hadoop
[root@ip-172-31-22-59 local]# export HADOOP_HOME=/usr/local/hadoop
[root@ip-172-31-22-59 local]# export HADOOP_MAPRED_HOME=$HADOOP_HOME
[root@ip-172-31-22-59 local]# export HADOOP_COMMON_HOME=$HADOOP_HOME
[root@ip-172-31-22-59 local]# export HADOOP_HDFS_HOME=$HADOOP_HOME
[root@ip-172-31-22-59 local]# export YARN_HOME=$HADOOP_HOME
[root@ip-172-31-22-59 local]# export HADOOP_COMMON_LIB_NATIVE_DIR=$HADOOP_HOME/lib/native
[root@ip-172-31-22-59 local]# export PATH=$PATH:$HADOOP_HOME/sbin:$HADOOP_HOME/bin
[root@ip-172-31-22-59 local]# source ~/.bashrc
[root@ip-172-31-22-59 local]# cd $HADOOP_HOME/etc/hadoop
  
```

```
<configuration>
  <property>
    <name>fs.default.name</name>
    <value>hdfs://localhost:9000 </value>
  </property>
</configuration>
```

```
<configuration>
  <property>
    <name>dfs.replication</name>
    <value>1</value>
  </property>

  <property>
    <name>dfs.name.dir</name>
    <value>file:///home/hadoop/hadoopinfra/hdfs/namenode </value>
  </property>

  <property>
    <name>dfs.data.dir</name>
    <value>file:///home/hadoop/hadoopinfra/hdfs/datanode </value>
  </property>
</configuration>
```

```
<!--
<configuration>
<!-- Site specific YARN configuration properties -->
<property>
  <name>yarn.nodemanager.aux-services</name>
  <value>mapreduce_shuffle</value>
</property>
</configuration>
```

```
<configuration>
<property>
  <name>mapreduce.framework.name</name>
  <value>yarn</value>
</property>
</configuration>
```

**hadoop** All Applications Logged in as

Cluster Metrics

Apps Submitted	Apps Pending	Apps Running	Apps Completed	Containers Running	Memory Used	Memory Total	Memory Reserved	VCores Used	VCores Total	VCores Reserved	Active Nodes	Decommissioned Nodes	Lost Nodes	Unhealthy Nodes	Ret N
0	0	0	0	0	0 B	0 B	0 B	0	0	0	0	0	0	0	0

Show 20 entries

ID	User	Name	Application Type	Queue	StartTime	FinishTime	State	FinalStatus	Progress	Tracking UI	Blacklisted Node
No data available in table											

Showing 0 to 0 of 0 entries First Previous Next

```
[root@ip-172-31-22-59 local]# wget http://www-us.apache.org/dist/sqoop/1.4.6/sqoop-1.4.6.bin_hadoop-2.0.4-alpha.tar.gz
--2017-10-08 13:26:33-- http://www-us.apache.org/dist/sqoop/1.4.6/sqoop-1.4.6.bin_hadoop-2.0.4-alpha.tar.gz
Resolving www-us.apache.org (www-us.apache.org)... 140.211.11.105
Connecting to www-us.apache.org (www-us.apache.org)|140.211.11.105|80... connected.
HTTP request sent, awaiting response... 200 OK
Length: 16870735 (16M) [application/x-gzip]
Saving to: 'sqoop-1.4.6.bin_hadoop-2.0.4-alpha.tar.gz'

sqoop-1.4.6.bin_hadoop-2.0.4-alpha.tar.g 100%[=====] 16.09M 13.6MB/s in 1.2s
2017-10-08 13:26:34 (13.6 MB/s) - 'sqoop-1.4.6.bin_hadoop-2.0.4-alpha.tar.gz' saved [16870735/16870735]

[root@ip-172-31-22-59 local]# tar xf sqoop-1.4.6.bin_hadoop-2.0.4-alpha.tar.gz
[root@ip-172-31-22-59 local]# mv sqoop-1.4.6.bin_hadoop-2.0.4-alpha sqoop
[root@ip-172-31-22-59 local]#
```

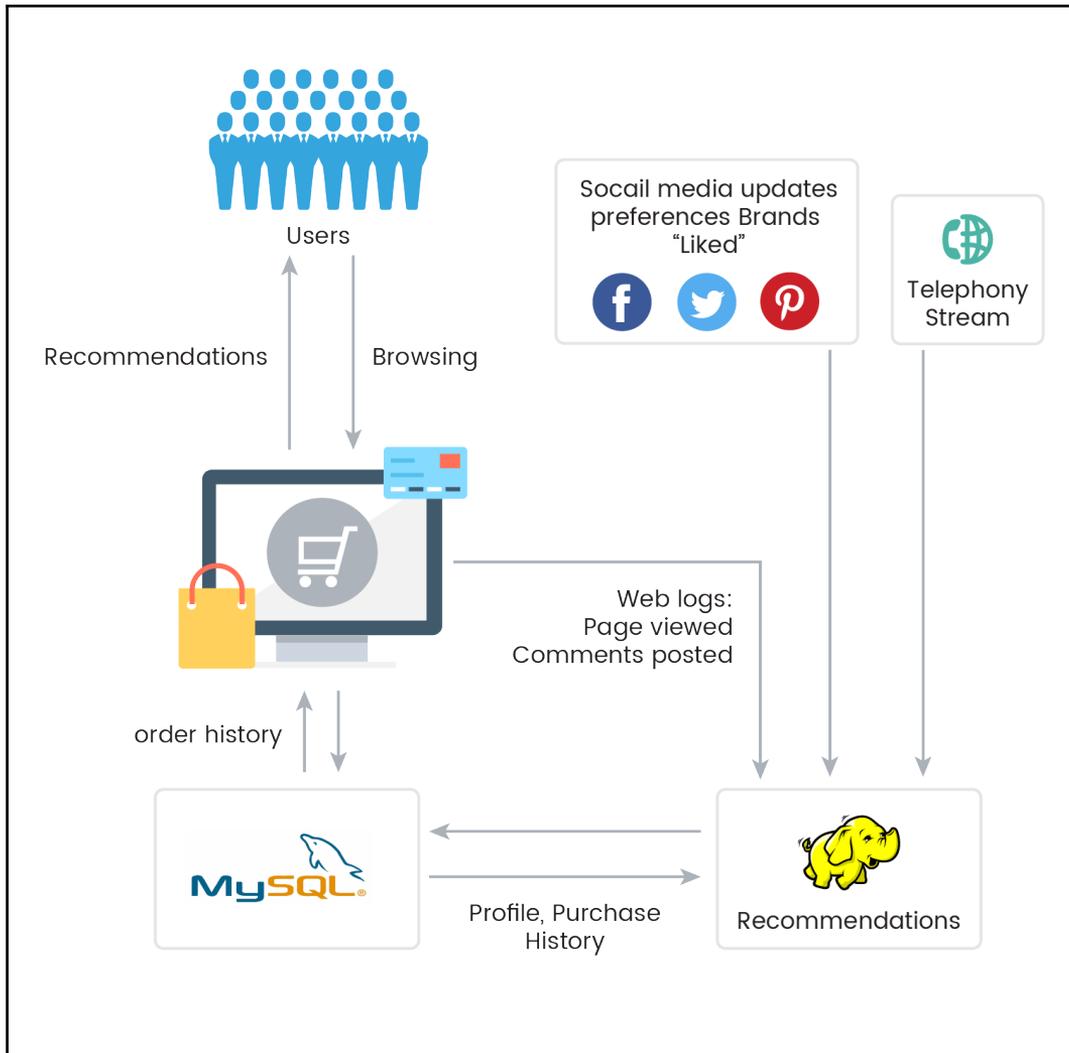
```
[root@ip-172-31-22-59 bin]# export HADOOP_COMMON_HOME=/usr/local/hadoop
[root@ip-172-31-22-59 bin]# export HADOOP_MAPRED_HOME=/usr/local/hadoop
```

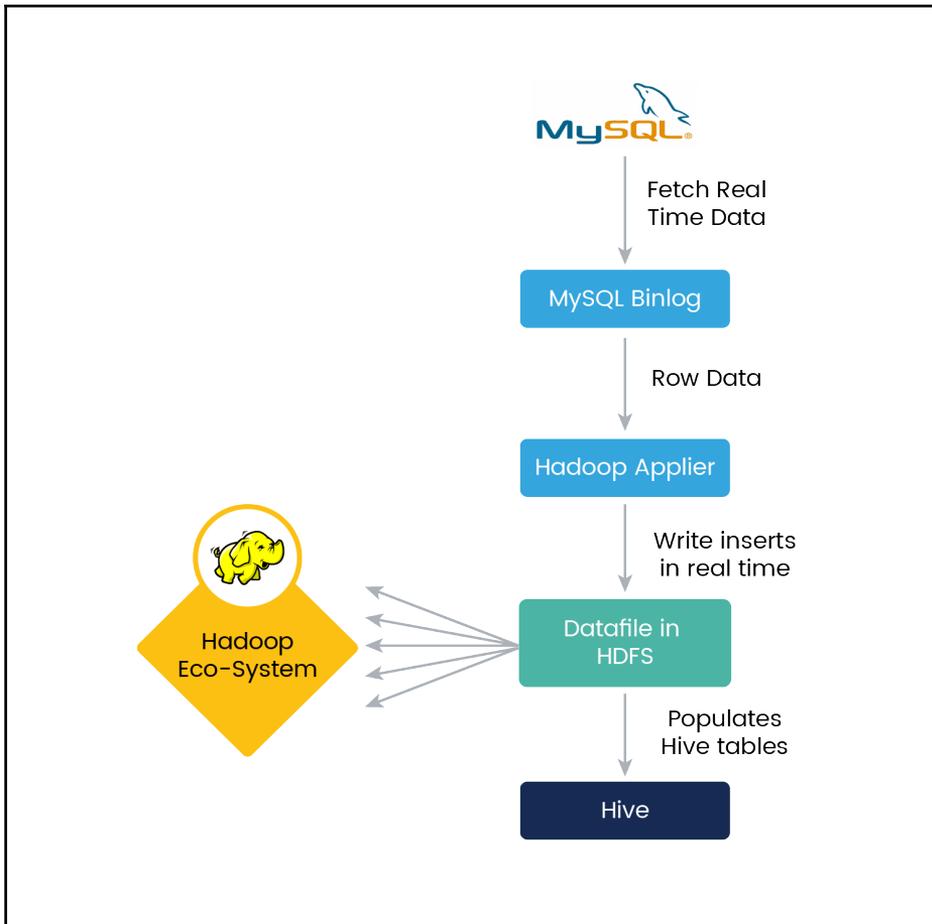
```
mysql> create database hadoop_bigdata;
Query OK, 1 row affected (0.00 sec)

mysql>
mysql> CREATE TABLE
-> `hadoop_bigdata`.`users` (
-> `user_id` INT NOT NULL AUTO INCREMENT ,
-> `email` VARCHAR(200) NOT NULL ,
-> `date_of_joining` DATE NOT NULL ,
-> `date_of_birth` DATE NOT NULL ,
-> `first_name` VARCHAR(200) NOT NULL ,
-> PRIMARY KEY (`user_id`)
-> ) ENGINE = InnoDB;
Query OK, 0 rows affected (0.02 sec)

mysql>
```

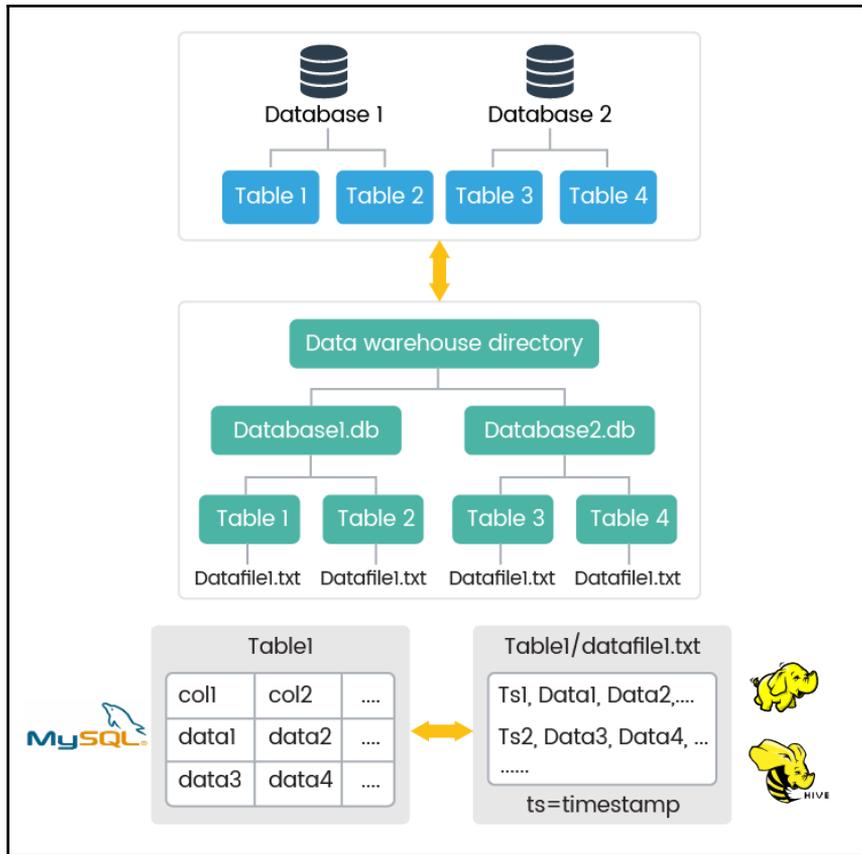
## Chapter 10: Case study: Part II - Real time event processing using MySQL applier





```

[root@ip-172-31-27-59 hive]# hive
17/10/11 01:56:27 INFO Configuration.deprecation: mapred.input.dir.recursive is deprecated. Instead, use mapreduce.input.fileinputformat.input.dir.recursive
17/10/11 01:56:27 INFO Configuration.deprecation: mapred.max.split.size is deprecated. Instead, use mapreduce.input.fileinputformat.split.maxsize
17/10/11 01:56:27 INFO Configuration.deprecation: mapred.min.split.size is deprecated. Instead, use mapreduce.input.fileinputformat.split.minsize
17/10/11 01:56:27 INFO Configuration.deprecation: mapred.min.split.size.per.rack is deprecated. Instead, use mapreduce.input.fileinputformat.split.minsize.per.rack
17/10/11 01:56:27 INFO Configuration.deprecation: mapred.min.split.size.per.node is deprecated. Instead, use mapreduce.input.fileinputformat.split.minsize.per.node
17/10/11 01:56:27 INFO Configuration.deprecation: mapred.reduce.tasks is deprecated. Instead, use mapreduce.job.reduces
17/10/11 01:56:27 INFO Configuration.deprecation: mapred.reduce.tasks.speculative.execution is deprecated. Instead, use mapreduce.reduce.speculative
Logging initialized using configuration in jar:file:/usr/local/hadoop/hive/lib/hive-common-0.12.0.jar!/hive-log4j.properties
SLF4J: Class path contains multiple SLF4J bindings.
SLF4J: Found binding in [jar:file:/usr/local/hadoop/share/hadoop/common/lib/slf4j-log4j12-1.7.5.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: Found binding in [jar:file:/usr/local/hadoop/hive/lib/slf4j-log4j12-1.6.1.jar!/org/slf4j/impl/StaticLoggerBinder.class]
SLF4J: See http://www.slf4j.org/codes.html#multiple_bindings for an explanation.
SLF4J: Actual binding is of type [org.slf4j.impl.Log4jLoggerFactory]
hive>
  
```



# Index