

Creating a multimaster MariaDB replica

This topic will cover the process to create a two node Maria DB cluster. The cluster will be configured to allow Soffid console to use either database node, which in turn will replicate data changes to the other one.

Node 1 action	Node 2 action
Create and setup a Maria DB in node 1.	
Configure Maria DB to generate binary log files. Add the following lines to /etc/mysql/my.cnf: <pre>server-id = 1 log-bin binlog-format=row expire_logs_days = 15 max_binlog_size = 1000M replicate-ignore-table = soffid.SC_SEQUENCE slave-skip-errors = 1032,1053,1062</pre>	
Restart MariaDB: <div data-bbox="123 1432 782 1477" style="border: 1px solid #ccc; padding: 5px; width: 415px;">service mysql restart</div>	
	Create and setup a Maria DB in node 2.
	Configure Maria DB to generate binary log files. Add the following lines to /etc/mysql/my.conf: <pre>server-id = 2 log-bin binlog-format=row expire_logs_days = 15 max_binlog_size = 1000M replicate-ignore-table = soffid.SC_SEQUENCE slave-skip-errors = 1032,1053,1062</pre>

Node 1 action	Node 2 action								
	Restart MariaDB: <div style="border: 1px solid #ccc; padding: 5px; display: inline-block;">service mysql restart</div>								
Dump current database contents: <code>mysqldump soffid -u soffid -p >soffid.data</code>	Load current database contents <code>mysql -u soffid -p < soffid.data</code>								
	Create a user for node 1 to fetch data from node 2. From mysql, execute: <code>grant replication slave on *.* to <u>replication_user</u>@<NODE1-IP></code> <code>set password for <u>replication_user</u>@1<NODE1-IP> = password('<NODE1-PASS>')</code>								
Create a user for node 2 to fetch data from node 1. From mysql, execute: <code>grant replication slave on *.* to <u>replication_user</u>@<NODE2-IP></code> <code>set password for <u>replication_user</u>@1<NODE2-IP> = password('<NODE2-PASS>')</code>									
Query current binary log position: <code>MariaDB [(none)]> show master status;</code>									
The result should look like this: <table border="1" style="width: 100%; border-collapse: collapse;"> <thead> <tr> <th style="text-align: left; padding: 2px;">File</th> <th style="text-align: left; padding: 2px;">Position</th> <th style="text-align: left; padding: 2px;">Binlog_Do_DB</th> <th style="text-align: left; padding: 2px;">Binlog_Ignore_DB</th> </tr> </thead> <tbody> <tr> <td style="padding: 2px;">mysqld-bin.000030</td> <td style="padding: 2px;">68175</td> <td style="padding: 2px;"></td> <td style="padding: 2px;"></td> </tr> </tbody> </table>	File	Position	Binlog_Do_DB	Binlog_Ignore_DB	mysqld-bin.000030	68175			
File	Position	Binlog_Do_DB	Binlog_Ignore_DB						
mysqld-bin.000030	68175								
The got values will be used on node 2 to start replica process.									

Node 1 action	Node 2 action								
	<p>Start replication from node 1 to node 2. From mysql, execute the following sentence, replacing proper values:</p> <pre>CHANGE MASTER TO MASTER_HOST='<NODE1-IP>', MASTER_USER='replication_user', MASTER_PASSWORD='<NODE2-PASS>', MASTER_PORT=3306, MASTER_LOG_FILE='<NODE1-FILE>', /* i.e. mysql-bin.000030 */ MASTER_LOG_POS=<NODE1-POSITION>, /* i.e. 68175 */ MASTER_CONNECT_RETRY=10;</pre>								
	<p>Verify replica is working right, by executing <code>SHOW SLAVE STATUS G</code></p> <p>Check following lines:</p> <ul style="list-style-type: none"> Slave_IO_Running: Yes Slave_SQL_Running: Yes Seconds_Behind_Master: 0 								
	<p>Query current binary log position: <code>MariaDB [(none)]> show master status;</code></p> <p>The result should look like this:</p> <table border="1" data-bbox="806 1185 1489 1343"> <thead> <tr> <th data-bbox="806 1185 933 1275">File</th><th data-bbox="933 1185 1108 1275">Position</th><th data-bbox="1108 1185 1298 1275">Binlog_Do_DB</th><th data-bbox="1298 1185 1489 1275">Binlog_Ignore_DB</th></tr> </thead> <tbody> <tr> <td data-bbox="806 1275 933 1343">mysqld-bin.000060</td><td data-bbox="933 1275 1108 1343">98325</td><td data-bbox="1108 1275 1298 1343"></td><td data-bbox="1298 1275 1489 1343"></td></tr> </tbody> </table> <p>The got values will be used on node 1 to start replica process.</p>	File	Position	Binlog_Do_DB	Binlog_Ignore_DB	mysqld-bin.000060	98325		
File	Position	Binlog_Do_DB	Binlog_Ignore_DB						
mysqld-bin.000060	98325								
<p>Now, start replication from node 2 to node 1. From mysql, execute the following sentence, replacing proper values:</p> <pre>CHANGE MASTER TO MASTER_HOST='<NODE2-IP>', MASTER_USER='replication_user', MASTER_PASSWORD='<NODE1-PASS>', MASTER_PORT=3306, MASTER_LOG_FILE='<NODE2-FILE>', /* i.e. mysql-bin.000060 */ MASTER_LOG_POS=<NODE2-POSITION>, /* i.e. 98325 */ MASTER_CONNECT_RETRY=10;</pre>									

Node 1 action	Node 2 action
<p>Verify replica is working right, by executing <code>SHOW SLAVE STATUS \G</code></p> <p>Check following lines: Slave_IO_Running: Yes Slave_SQL_Running: Yes Seconds_Behind_Master: 0</p>	
<p>Now, create and start SC_SEQUENCE table in node 1. This sequence will generate values 1, 11, 21, 31, 41, and so on:</p> <pre><code>CREATE TABLE `SC_SEQUENCE` (`SEQ_NEXT` bigint(20) NOT NULL, `SEQ_CACHE` bigint(20) NOT NULL, `SEQ_INCREMENT` bigint(20) NOT NULL); INSERT INTO SC_SEQUENCE VALUES (1, 100, 10);</code></pre>	
	<p>Now, create and start SC_SEQUENCE table in node 2. This sequence will generate values 2, 12, 22, 32, 42, and so on::</p> <pre><code>CREATE TABLE `SC_SEQUENCE` (`SEQ_NEXT` bigint(20) NOT NULL, `SEQ_CACHE` bigint(20) NOT NULL, `SEQ_INCREMENT` bigint(20) NOT NULL); INSERT INTO SC_SEQUENCE VALUES (2, 100, 10);</code></pre>

Now, configure the Console to use the following jdbc URL:

`jdbc:mariadb:sequential://mariadb-host-1,mariadb-host-2/soffid`

Revision #7

Created 17 March 2021 09:32:17 by pgarcia@soffid.com

Updated 1 June 2022 14:30:18 by pgarcia@soffid.com