

## RapidsDB Release Notes

**Release Date: December 22nd, 2016**

### Software

#### RapidsDB Cluster:

/home/dave/software/rapidsdb/release\_3.0\_beta/rdp-gb90aea8.run

RELEASE ID: 2.4-145-gb90aea8

#### RapidsDB Manager:

/home/dave/software/rapidsdb/release\_2.4/rdp-opsconsole-1.0-45.run

**Supported Linux Systems: Linux CentOS or RHEL release 6.5 or later.**

**Java version: THIS RELEASE REQUIRES JAVA8 JDK**

### Documentation

/home/dave/software/rapidsdb/release\_3.0\_beta/RapidsDB\_Installation\_and\_Management\_Guide\_Release\_3.0\_Beta.pdf

/home/dave/software/rapidsdb/release\_3.0\_beta/RapidsDB\_User\_Guide\_Release\_3.0\_Beta.pdf

/home/dave/software/rapidsdb/release\_3.0\_beta/RapidsDB\_Streaming\_Guide\_Release\_3.0\_Beta.pdf

### Installation

Refer to the RapidsDB Installation and Management Guide Release 3.0 Beta for details on how to install the RapidsDB system.

### New Features

1. This is the initial release for the RapidsDB in-memory Storage Engine, RapidsSE. This release supports a single node. The user can CREATE and DROP tables from the rapids-shell (see DDL support below). The user can load data into RapidsSE using either the RapidsSE load utility, ld, or by using INSERT or INSERT ... SELECT statements from the rapids-shell (see DML support below).
2. DDL support. The user issue can issue create or drop table commands in the rapids-shell against either a MemSQL, VoltDB or RapidsSE Connector.
3. DML support. The user can issue an INSERT or INSERT ... SELECT command against either a MemSQL, VoltDB or RapidsSE Connector, where the SELECT include tables managed by a MemSQL, VoltDB, RapidsSE or Stream Connector.

## Problems Fixed

Problem Id	Description
RDP-186	<p>Type mismatch error reported incorrectly with nested sub-queries when the innermost subquery refers to a variable from the outermost query. The system will incorrectly report a "type mismatch." error. This error only occurs when: A) the reference is discontinuous (i.e. not a reference to the immediate outer query); and B) the reference is a non-predicate reference (i.e. is not a WHERE, ON or HAVING predicate).</p> <p>Below is an example:</p> <pre>SELECT * FROM t1 WHERE k = ( SELECT k FROM t2 WHERE i = ( SELECT j / t1.d FROM t3 ) );</pre> <p>In this example, the reference to t1.d is what causes the problem.</p>
RDP-194	<p>The system will report an exception when there is GROUP BY within a nested scalar subquery.</p> <pre>rapids &gt; SELECT * FROM public.region WHERE ( SELECT x FROM ( SELECT 0 AS x FROM public.nation GROUP BY n_name ) ) = 0 ;</pre> <p>System exception on node NODE1:</p> <pre>java.lang.NullPointerException     at com.google.common.base.Preconditions.checkNotNull(Preconditions.java:212)     at com.rapidsdata.plan.ProjectOp.asBound(ProjectOp.java:965)     at com.rapidsdata.planner.PostProjectOp.asBound(PostProjectOp.java:187)     at com.rapidsdata.planner.PostProjectOp.asBound(PostProjectOp.java:26)     at com.rapidsdata.planner.Planner.newPlanFor(Planner.java:195)     at com.rapidsdata.dqx.DqxProcessor.process(DqxProcessor.java:130)     at com.rapidsdata.dqx.DqxProcessor.process(DqxProcessor.java:83)     at com.rapidsdata.mesh.AbstractWorker.run(AbstractWorker.java:47)     at java.lang.Thread.run(Thread.java:745)</pre>
RDP-207	<p>With this release of RapidsDB, the partitioning support used to parallelize queries across multiple VoltDB instances has been disabled. The net effect of this is that all queries pushed down to VoltDB will be sent to a single node, and we will rely on VoltDB to parallelize the (pushed down) query, and the RapidsDB Execution Engine will not parallelize the execution of any remaining parts of the query, instead any remaining processing will be carried out on the DQC node. This will be the same as the processing model used for MemSQL. This restriction will be lifted in the very near future for queries going against VoltDB (which will be consistent with the processing model for R2.3 and prior releases). This could negatively impact the performance of queries involving VoltDB.</p>

	The Stream Connector will report an error that partitioning is not enabled when attempting to execute a query against a stream.
--	---

### Known Problems Remaining

Problem Id	Description
RDP-202	<p>Multiple intersect query with corresponding by returns exception.</p> <pre> rapids &gt; select * from public.table1 where colb1 in ('violet','pink','purple','indigo') intersect corresponding by (colc1,ucol) select * from public.table1 where ucolvc not in ('eee','fff') intersect corresponding by (colc1,pcol1) select * from public.table1 where pcol1 &lt; 50 order by ucol; System exception on node NODE1: java.lang.IllegalStateException: set op column list fault     at com.google.common.base.Preconditions.checkState(Preconditions.java:174)     at com.rapidsdata.planner.Planner\$Generator.validateSetOperands(Planner.java:3085)     at com.rapidsdata.planner.Planner\$Generator.visitSetExpr(Planner.java:2879)     at com.rapidsdata.planner.Planner\$Generator.visitSetExpr(Planner.java:218)     at com.rapidsdata.parser.DqsParser\$SetExprContext.accept(DqsParser.java:1346)     at org.antlr.v4.runtime.tree.AbstractParseTreeVisitor.visitChildren(AbstractParseTreeVisitor. java:70)     at com.rapidsdata.planner.Planner\$Generator.visitSetQuery(Planner.java:2804)     at com.rapidsdata.planner.Planner\$Generator.visitSetQuery(Planner.java:218)     at com.rapidsdata.parser.DqsParser\$SetQueryContext.accept(DqsParser.java:1235)     at org.antlr.v4.runtime.tree.AbstractParseTreeVisitor.visitChildren(AbstractParseTreeVisitor. java:70)     at com.rapidsdata.planner.Planner\$Generator.visitQuery(Planner.java:622)     at com.rapidsdata.planner.Planner\$Generator.visitQuery(Planner.java:218)     at com.rapidsdata.parser.DqsParser\$QueryContext.accept(DqsParser.java:336)     at org.antlr.v4.runtime.tree.AbstractParseTreeVisitor.visitChildren(AbstractParseTreeVisitor. java:70)     at com.rapidsdata.parser.DqsBaseVisitor.visitQueryStatement(DqsBaseVisitor.java:69) </pre>

	<pre> at com.rapidsdata.parser.DqsParser\$QueryStatementContext.accept(DqsParser.java:219) at org.antlr.v4.runtime.tree.AbstractParseTreeVisitor.visitChildren(AbstractParseTreeVisitor.java:70) at com.rapidsdata.parser.DqsBaseVisitor.visitStatement(DqsBaseVisitor.java:805) at com.rapidsdata.parser.DqsParser\$StatementContext.accept(DqsParser.java:155) at org.antlr.v4.runtime.tree.AbstractParseTreeVisitor.visit(AbstractParseTreeVisitor.java:42) at com.rapidsdata.planner.Planner\$Generator.generateFor(Planner.java:244) at com.rapidsdata.planner.Planner.newPlanFor(Planner.java:108) at com.rapidsdata.dqx.DqxProcessor.process(DqxProcessor.java:130) at com.rapidsdata.dqx.DqxProcessor.process(DqxProcessor.java:83) at com.rapidsdata.mesh.AbstractWorker.run(AbstractWorker.java:47) at java.lang.Thread.run(Thread.java:745) </pre>
RDP-212	<p>Creating a table with a Boolean data type using the MemSQL Connector can fail with the following error:</p> <pre> System exception on node NODE1: java.lang.IllegalStateException at com.google.common.base.Preconditions.checkState(Preconditions.java:159) at com.rapidsdata.memsql.MemsqlCreateOp.executeInSitu(MemsqlCreateOp.java:119) at com.rapidsdata.plan.Plan.execute(Plan.java:622) at com.rapidsdata.planner.RefreshConnectorOp.executeInSitu(RefreshConnectorOp.java:108) at com.rapidsdata.plan.Plan.execute(Plan.java:622) at com.rapidsdata.dqx.DqxProcessor.process(DqxProcessor.java:144) at com.rapidsdata.dqx.DqxProcessor.process(DqxProcessor.java:88) at com.rapidsdata.mesh.AbstractWorker.run(AbstractWorker.java:47) at java.lang.Thread.run(Thread.java:745) </pre>