



What's new in Hive 2.0

Sergey Shelukhin

What is Hive 2.0?

- **Split in 2015**
- **Hive 1.* is the "more stable" line**
 - Receives the bugfixes, some features and improvements
 - Keeps everything backward compatible
- **Hive 2.* is the "more ambitious" line**
 - Receives the bugfixes and improvements
 - Also receives all the major new features
 - Deprecates the support for some older features
- **Doesn't mean Hive 2 is unstable**
- **Where is Hive 1.3?!!**

When is Hive 2.0 coming?

- **The original plan was Dec 2015**
 - Unrealistic – too many blockers, too many features wanting to get in
- **2016-01-21**
 - 1 blocker left (hello Eugene 😊)!
 - Some features and improvements about to get in
 - RC 0 expected this week

Hive 2.0 at a (rather blurry) glance

- `project = HIVE AND fixVersion in (2.0.0, llap, spark-branch, hbase-metastore-branch) AND fixVersion not in (1.3.0, 1.2.2, 1.2.1, 1.0.1, 1.1.1, 1.2.0, 1.1.0) AND resolution = Fixed`
- 764 tickets (Hive 2.0 only)
 - 333 Sub-tasks (remember all those new features?)
 - 313 bugs (but we mark everything as Bug)
 - 99 Improvements and Tasks
- `project = HIVE AND fixVersion in (2.0.0, llap, spark-branch, hbase-metastore-branch) AND fixVersion not in (1.2.1, 1.0.1, 1.1.1, 1.2.0, 1.1.0) AND resolution = Fixed`
- 1193 tickets (Hive 2.0 + future Hive 1.3/1.2.2)

Upgraded versions

- Upgraded versions
 - Log4j 1 -> Slf4j/log4j 2 (perf gain – logging doesn't block the thread!)
 - Calcite 1.2 -> 1.5 (new features for CBO)
 - Tez 0.5 -> 0.8.2 (perf gains, new features, plugins)
 - Spark 1.3.1 -> 1.5 (perf gains, new features) (also in Hive 1.3)
 - DataNucleus 3 -> 4, Kryo 2 -> 3, Hbase 0.98 -> 1.1
 - Parquet 1.6 -> 1.8 (1.7 is also in Hive 1.3)
 - Thrift 0.9.2 -> 0.9.3, Avro 1.7.5 -> 1.7.7 (also in 1.3), etc.

Breaking things

- Java 6 no longer supported
- Hadoop 1 no longer supported on Hive 2 line (is it older than Java 6?)
- MR is deprecated, but still supported (use Spark or Tez!)
- Better defaults (enforce.bucketing, metastore schema verification, etc. on by default)
- Tightened safety settings (fails on some unsafe casts, etc.)

New features #1

- HPLSQL
- LLAP (beta)
- HBase metastore (alpha)
- Improvements to Hive-on-Spark
- Improvements to CBO

New features #2

- SQL Standard Auth is the default authorization (actually works)
- CLI mode for beeline (WIP to replace and deprecate CLI in Hive 2.*)
- Codahale-based metrics (also in 1.3)
- HS2 Web UI
- Stability Improvements and bugfixes for ACID (almost production ready now)
- Native vectorized mapjoin, vectorized reducesink, improved vectorized GBY, etc.
- Improvements to Parquet performance (PPD, memory manager, etc.)
- ORC schema evolution (beta)
- Improvement to windowing functions, refactoring ORC before split, SIMD optimizations, new LIMIT syntax, parallel compilation in HS2, improvements to Tez session management, many more

- Did I forget something?

HPLSQL

- HPL/SQL is a hybrid and heterogeneous language that understands syntaxes and semantics of almost any existing procedural SQL dialect
- Compatible with Oracle PL/SQL, ANSI/ISO SQL/PSM (IBM DB2, MySQL, Teradata etc.), PostgreSQL PL/pgSQL (Netezza), Transact-SQL (Microsoft SQL Server and Sybase)
- Key SQL features
 - Flow of Control Statements
 - Built-in Functions
 - Stored Procedures, Functions and Packages
 - Exception and Condition Handling
- Merged into Hive as hplsql module
- See **hplsql** command, docs at <http://www.hplsql.org/doc>

LLAP (beta in 2.0)

- Sub-second query execution in Hive via persistent daemons
 - Parallel execution and IO optimizations, JIT, etc.
 - Reduces fixed costs like container scheduling
 - Data caching
- Some limitations in 2.0 (mostly worked around gracefully)
 - Not tested well in secure clusters
 - Tez only (API and Spark integration in progress)
- User guide shortly after release
- Demo (in 25 seconds at the end)

HBase metastore (alpha in 2.0)

- Getting rid of DataNucleus/RDBMS
 - Writes that actually scale!
 - Reads that actually scale without "direct SQL"!
 - No more bizarre errors from 10000 different RDBMSes and 10000 different JDBC drivers!
 - No need for separate backup solution for metadata
 - No need to maintain 10000 upgrade scripts in future
- New features in progress
 - File metadata cache in HBase with PPD inside HBase, etc.
- Limitations on 2.0 – rough around the edges
 - Major limitation - no cross-entity transactions (future work with Omid)
- See <https://cwiki.apache.org/confluence/display/Hive/HBaseMetastoreDevelopmentGuide>

Hive-on-Spark improvements

- Dynamic partition pruning
- Make use of spark persistence for self-join union
- Vectorized mapjoin and other mapjoin improvements
- Parallel order by
- Container pre-warm

- Did I miss anything?

CBO

- New optimizations
 - More join improvements
 - LIMIT pushdown
- CBO now supplants many native Hive optimizers
 - PPD, constant propagation, etc.
- Performance improvements
- Calcite return path – avoid repeated op tree conversions (alpha)

30-second demo (in case you missed the previous meetup)

```
hive (tpcds_bin_partitioned_orc_1000)> set hive.llap.execution.mode;  
hive.llap.execution.mode=all  
hive (tpcds_bin_partitioned_orc_1000)> set hive.llap.execution.mode
```

[cn105-10] [(0*\$bash) 1\$ bash 2\$ bash 3\$ bash 4-\$ bash 5\$ bash 6\$ bash] [10/11 16:32]

Questions?