

TiDB 101

Fernando Ipar / Percona University Montevideo / April 2019

Pythian

About me

- Internal Principal Consultant @ Pythian.
- Pager-carrying, buck stops at my team.
- OSDB Practice (MySQL/MariaDB, MongoDB, Cassandra, ...).
- 19 years of MySQL experience.

pe- riodo	1A	2A	3B	4B	5B	6B	7B
1	1 1 +1 H 1,01						
2	2 3 +1 Li 6,94	2 4 +2 Be 9,01					
3	2 11 +1 Na 22,99	2 12 +2 Mg 24,312					
4	8 19 +1 K 39,10	8 20 +2 Ca 40,08	8 21 +3 Sc 44,96	8 22 +2 Ti 47,90	8 23 +2 V 50,94 ⁺⁵	8 24 +2 Cr 52	8 25 +2 Mn 54,94 ⁺⁷
5	18 37 +1 Rb 85,47	18 38 +2 Sr 87,62	18 39 +3 Y 88,91	18 40 +4 Zr 91,22	18 41 +3 Nb 92,91	18 42 +6 Mo 95,94	18 43 +4 Tc 99
6	18 55 +1 Cs 132,91	18 56 +2 Ba 137,34	18 57 +3 La 138,91	32 72 +4 Hf 178,49	32 73 +5 Ta 180,95	32 74 +6 W 183,85	32 75 +4 Re 186,2
7	18 87 +1 Fr 223	18 88 +2 Ra 226	18 89 +3 Ac 227				
tierras							
8	19 58 +3 Ce 140,12	20 59 +3 Pr 140,91	22 60 +3 Nd 144,24	23 61 +3 Pm 147	24 62 +2 Sm 150,35	25 63 +2 Eu 151,96	26 64 +2 Gd 157,25
9	20 72 +3 Er 167,26	21 73 +3 Tm 168,93	23 74 +3 Yb 173,05	24 75 +3 Lu 174,97	26 76 +2 Ho 164,93	27 77 +2 Dy 162,50	28 78 +2 Tb 158,93

SISTEMA

TiDB

Created by PingCAP
(Founded in April 2015
by 3 Infrastructure
Engineers)

TiDB Overview

- SQL-based (MySQL[0] protocol).
- Distributed (ACID).
- Transactional (Snapshot Isolation, advertised as REPEATABLE-READ).
- Inspired on Google's Spanner and F1.
- HTAP workloads.
- Open Source (Open Core).
- RocksDB for storage (though Storage Engine API exists)[1].
- "15+ Petabytes in 300+ Companies"[2]

[0] <https://github.com/pingcap/docs/blob/master/sql/mysql-compatibility.md>

[1] <https://blog.pythian.com/tag/myrocks/>

[2] <https://www.pingcap.com/success-stories/https://www.pingcap.com/success-stories/>

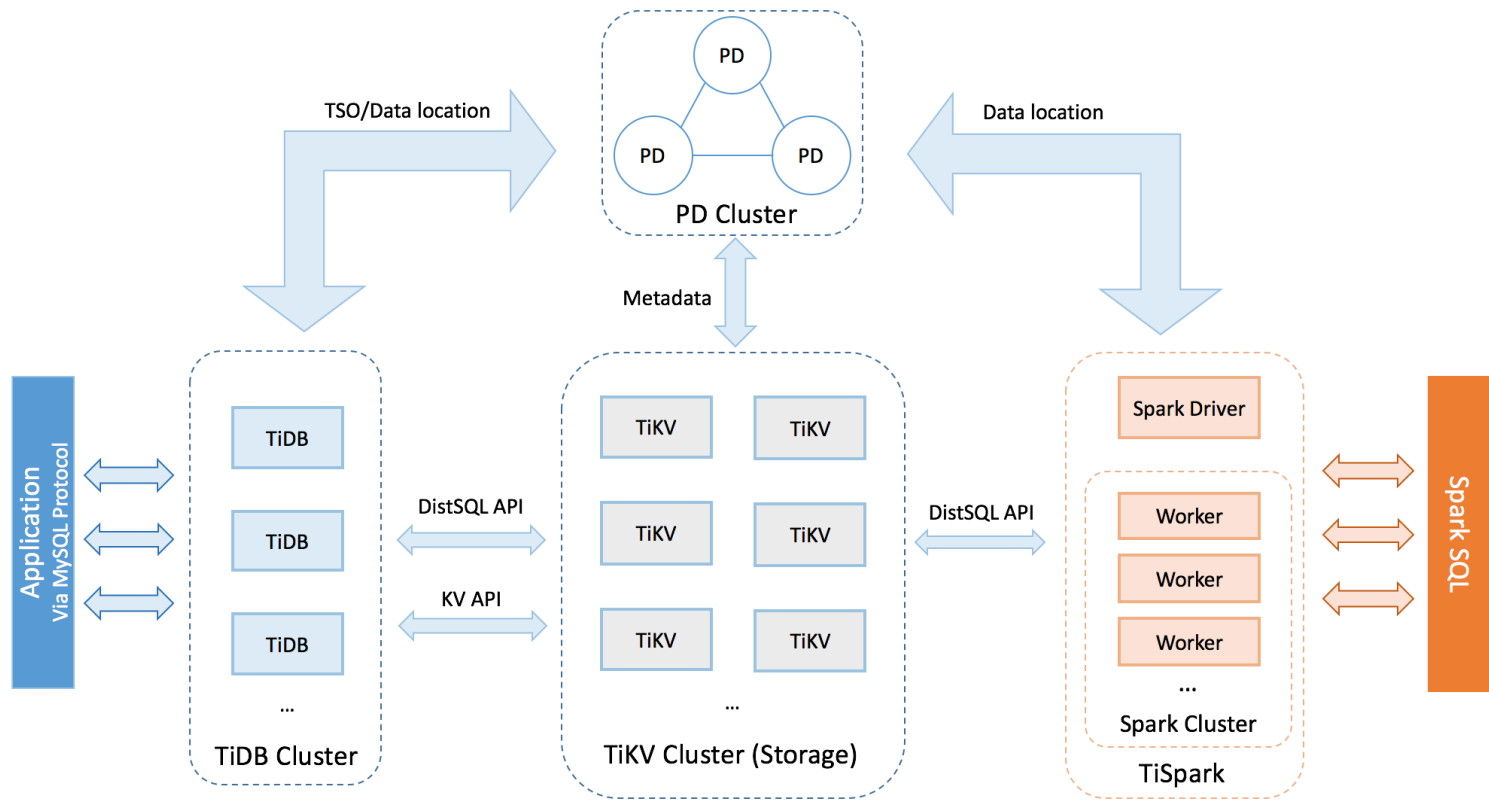
Good Use Cases

- You need a sharded MySQL.
- You have a sharded MySQL but have found out the shard key does not match all query use cases (who would have thought?!).
- Your data set does not fit on a single machine and your query volume is rather low.

Poor Use Cases

- Your workload is very latency-sensitive.
- Your data fits in a single server.
- You're closer to OLTP than OLAP.
- You have very large transactions.

TiDB Overview - Architecture



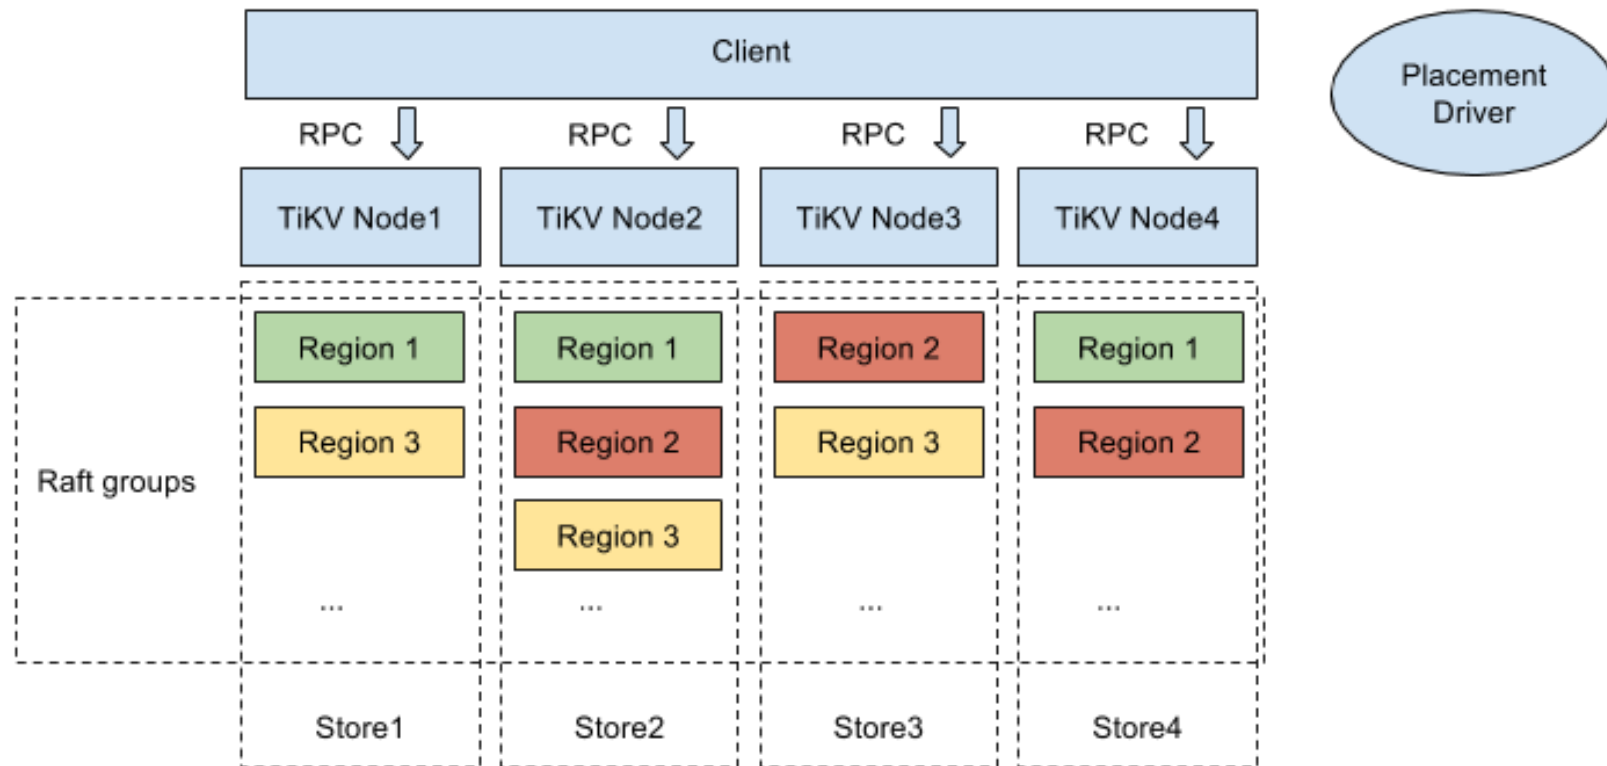
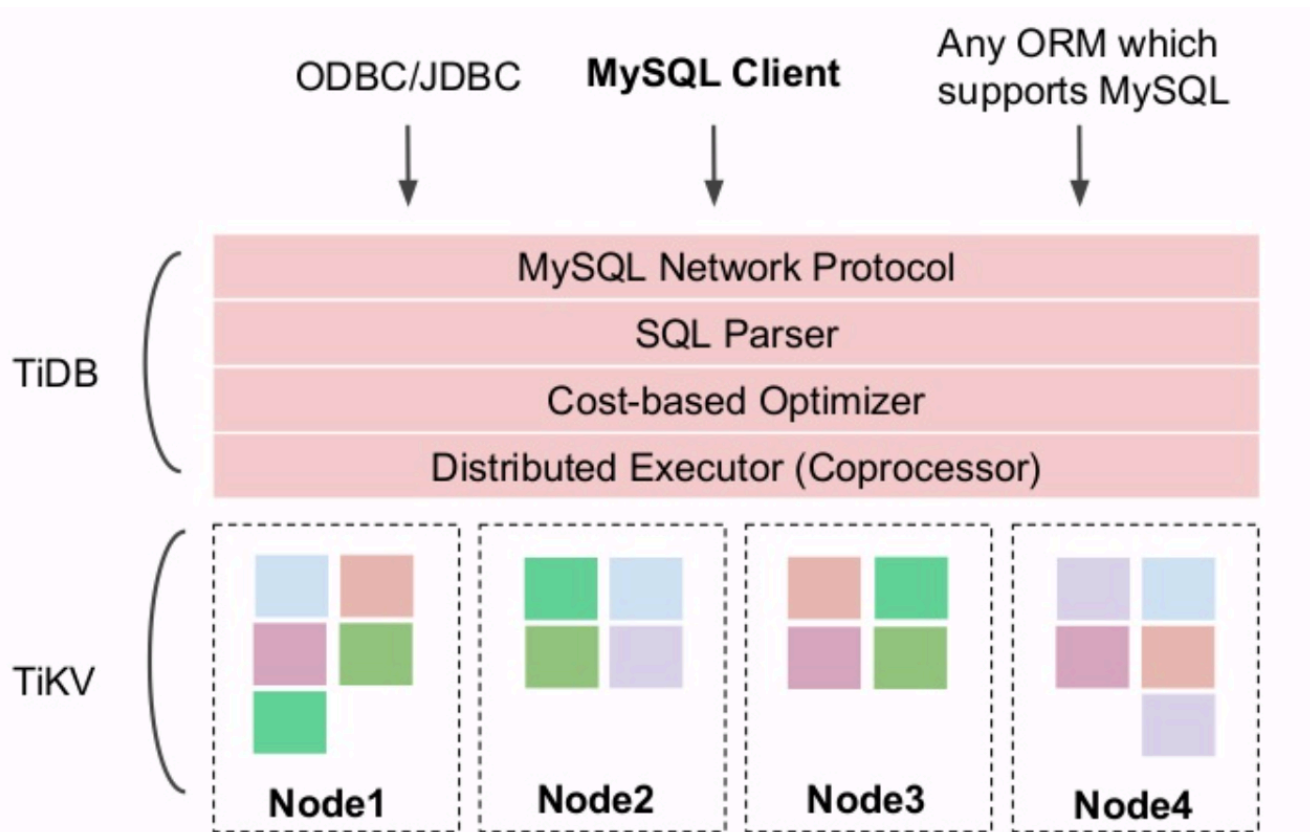


Image Credit: <https://github.com/tikv/tikv>

TiDB: OLTP + Ad Hoc OLAP



TiDB: SQL Parser and Coprocessor.

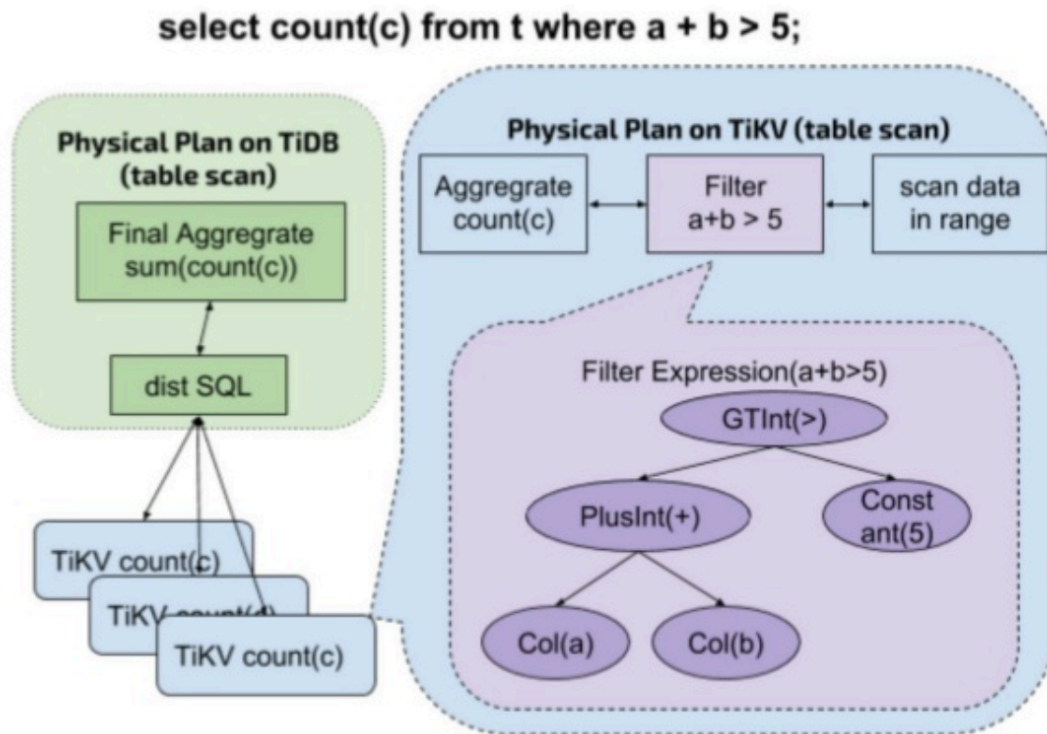



Image Credit: <https://www.slideshare.net/morgo/tidb-introduction>

TiDB: Replicating from MySQL



<https://pingcap.com/tidb-academy/>



I have certification vouchers for questions I can't answer right now or otherwise interesting questions.



THANK YOU