Merging patches for Spider into MariaDB 10.4

Kentoku SHIBA

The list of patches

The status of merging (in review) for 10.4

01. MDEV-7717 019_mariadb-10.0.15.vp_ref.diff
02. MDEV-7719 021_mariadb-10.0.15.merge_table.diff
03. MDEV-7724 027_mariadb-10.0.15.force_bulk_update.diff
04. MDEV-7735 039_mariadb-10.0.15.child_partition_pruning.diff
05. MDEV-7743 047_mariadb-10.0.15.child_partition_pruning2.diff
06. MDEV-7744 048_mariadb-10.0.15.vp_partition.diff
07. MDEV-7750 054_mariadb-10.0.15.for_vp_pruning.diff
08. MDEV-12970 056_mariadb-10.2.0.partition_top_table_fields.diff
09. MDEV-12971 060_mariadb-10.2.0.partition_reset_top_table_fields.diff
10. MDEV-13000 055_mariadb-10.2.0.for_vp_same_columns.diff

Merging VP storage engine

Merging VP storage engine

VP means Vertical Partitioning.

VP merges multiple child tables into a single View. VP chooses efficiently child tables for Each query.

VP structure sample of using different partitioning rules (1/2)

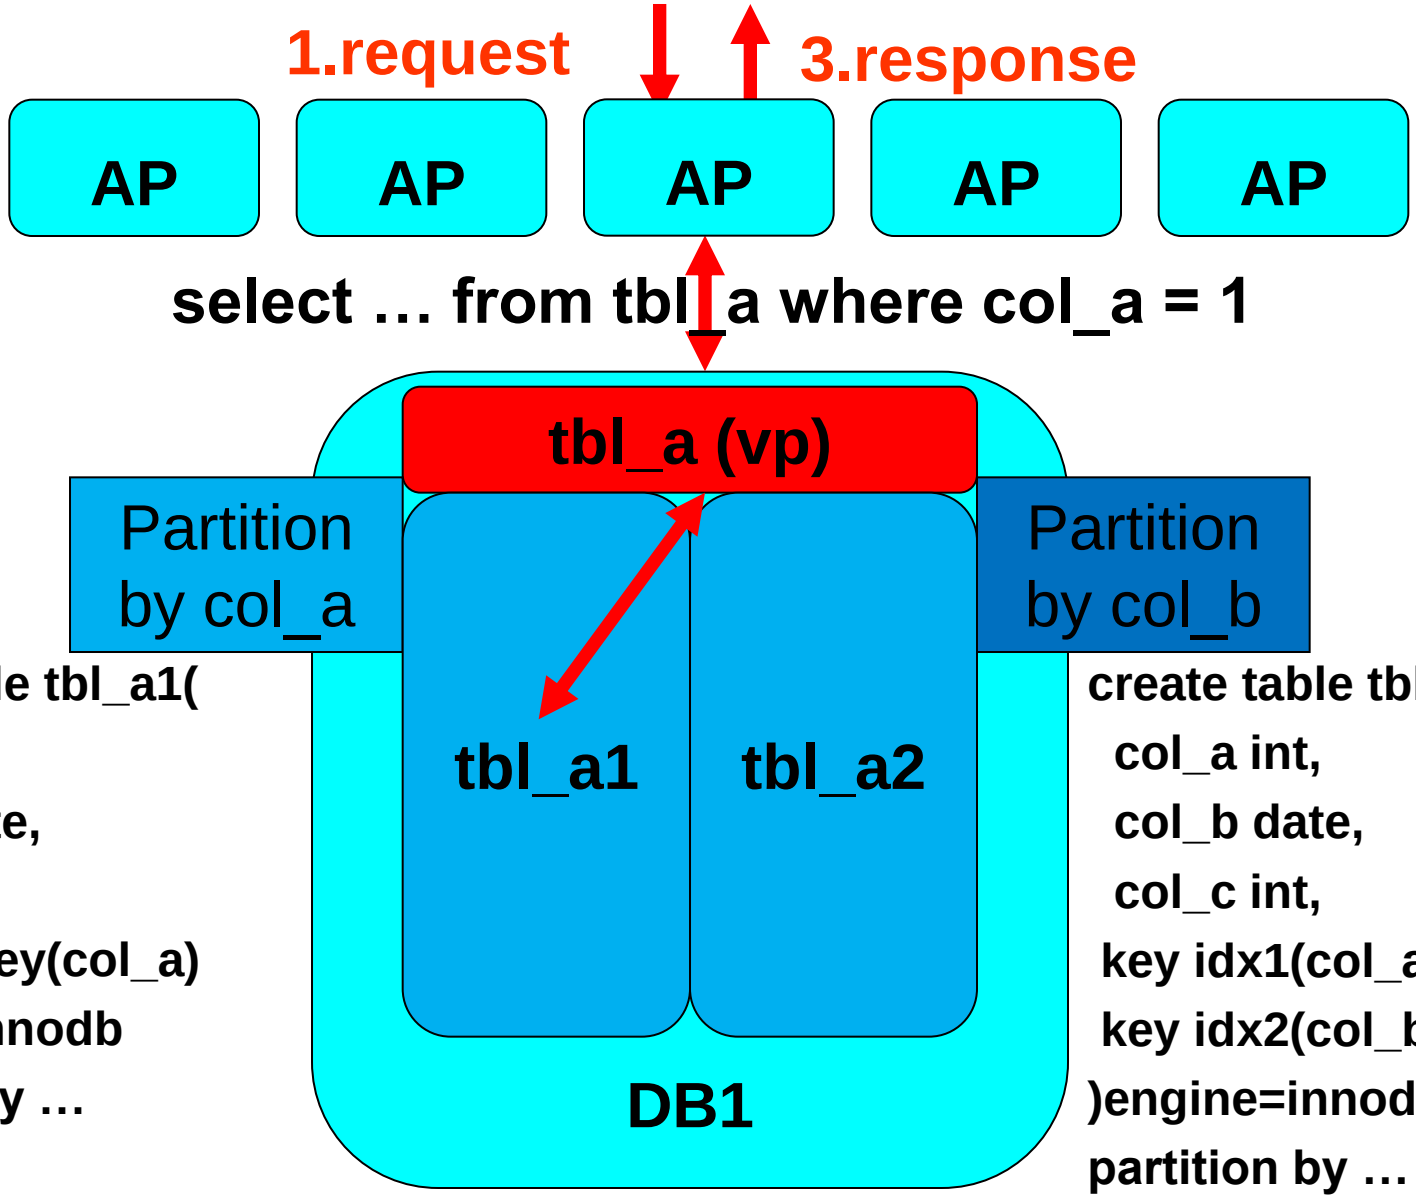

VP structure sample of using different partitioning rules (2/2)

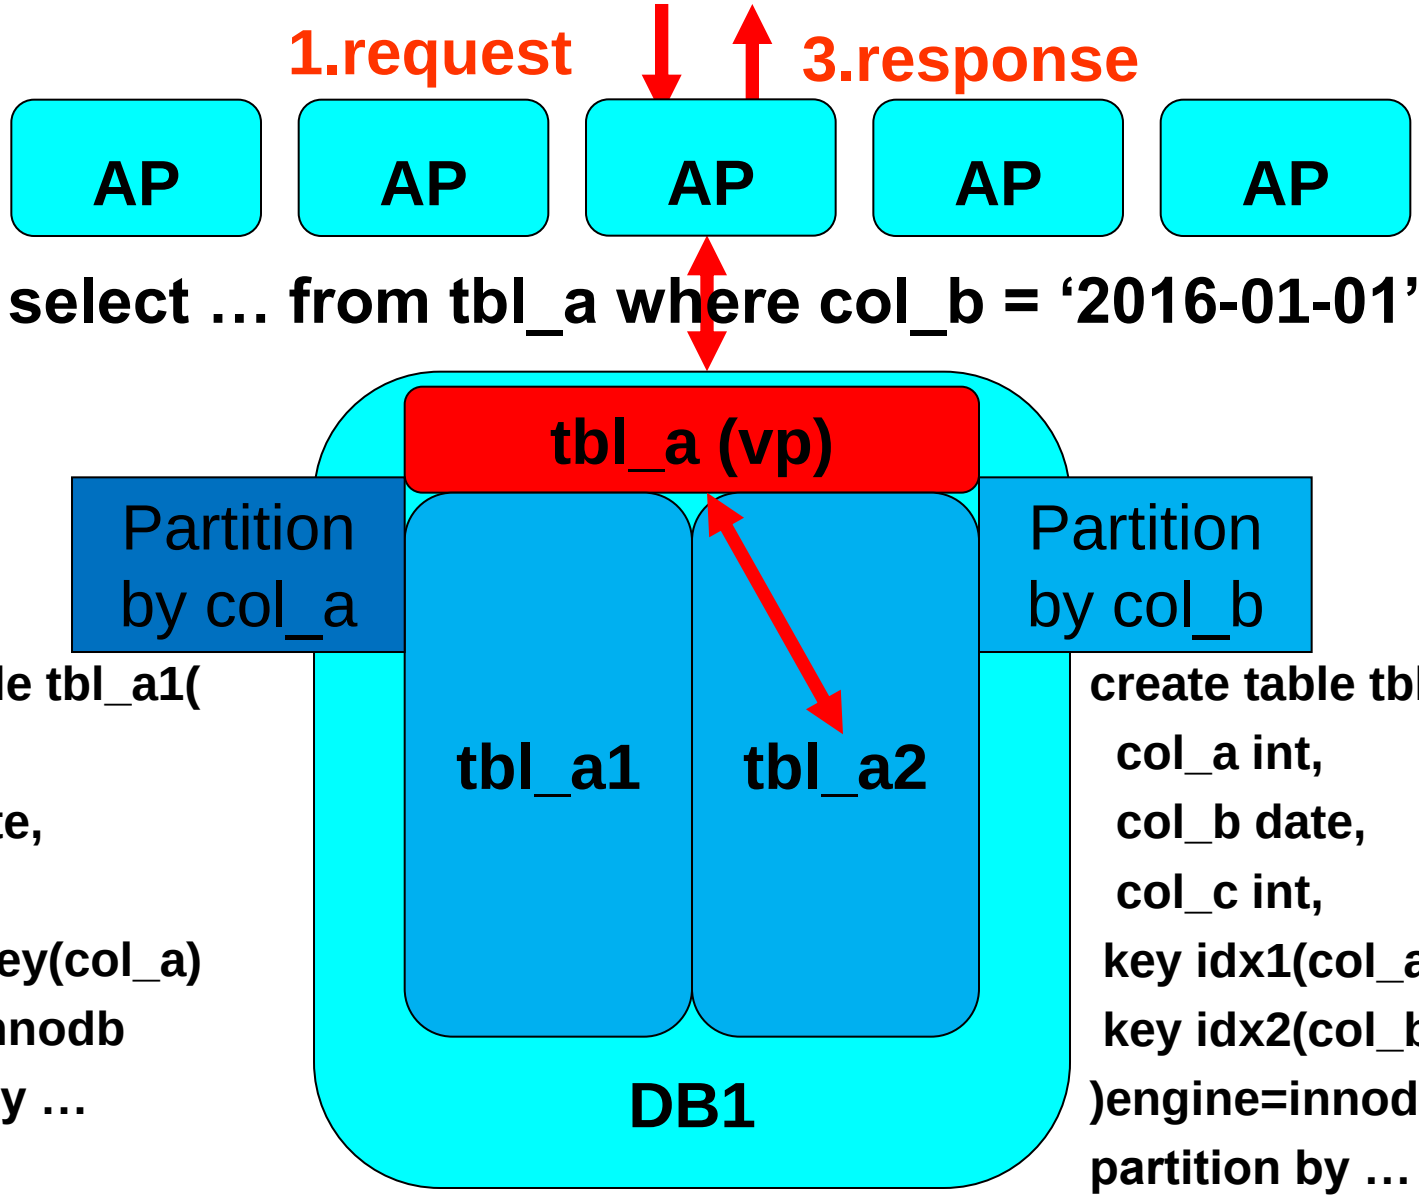

So, when you use sharded Spider tables which have different partitioning rules for VP child tables, VP chooses sharded Spider tables efficiently.

Structure sample of using different sharding rules (1/2)

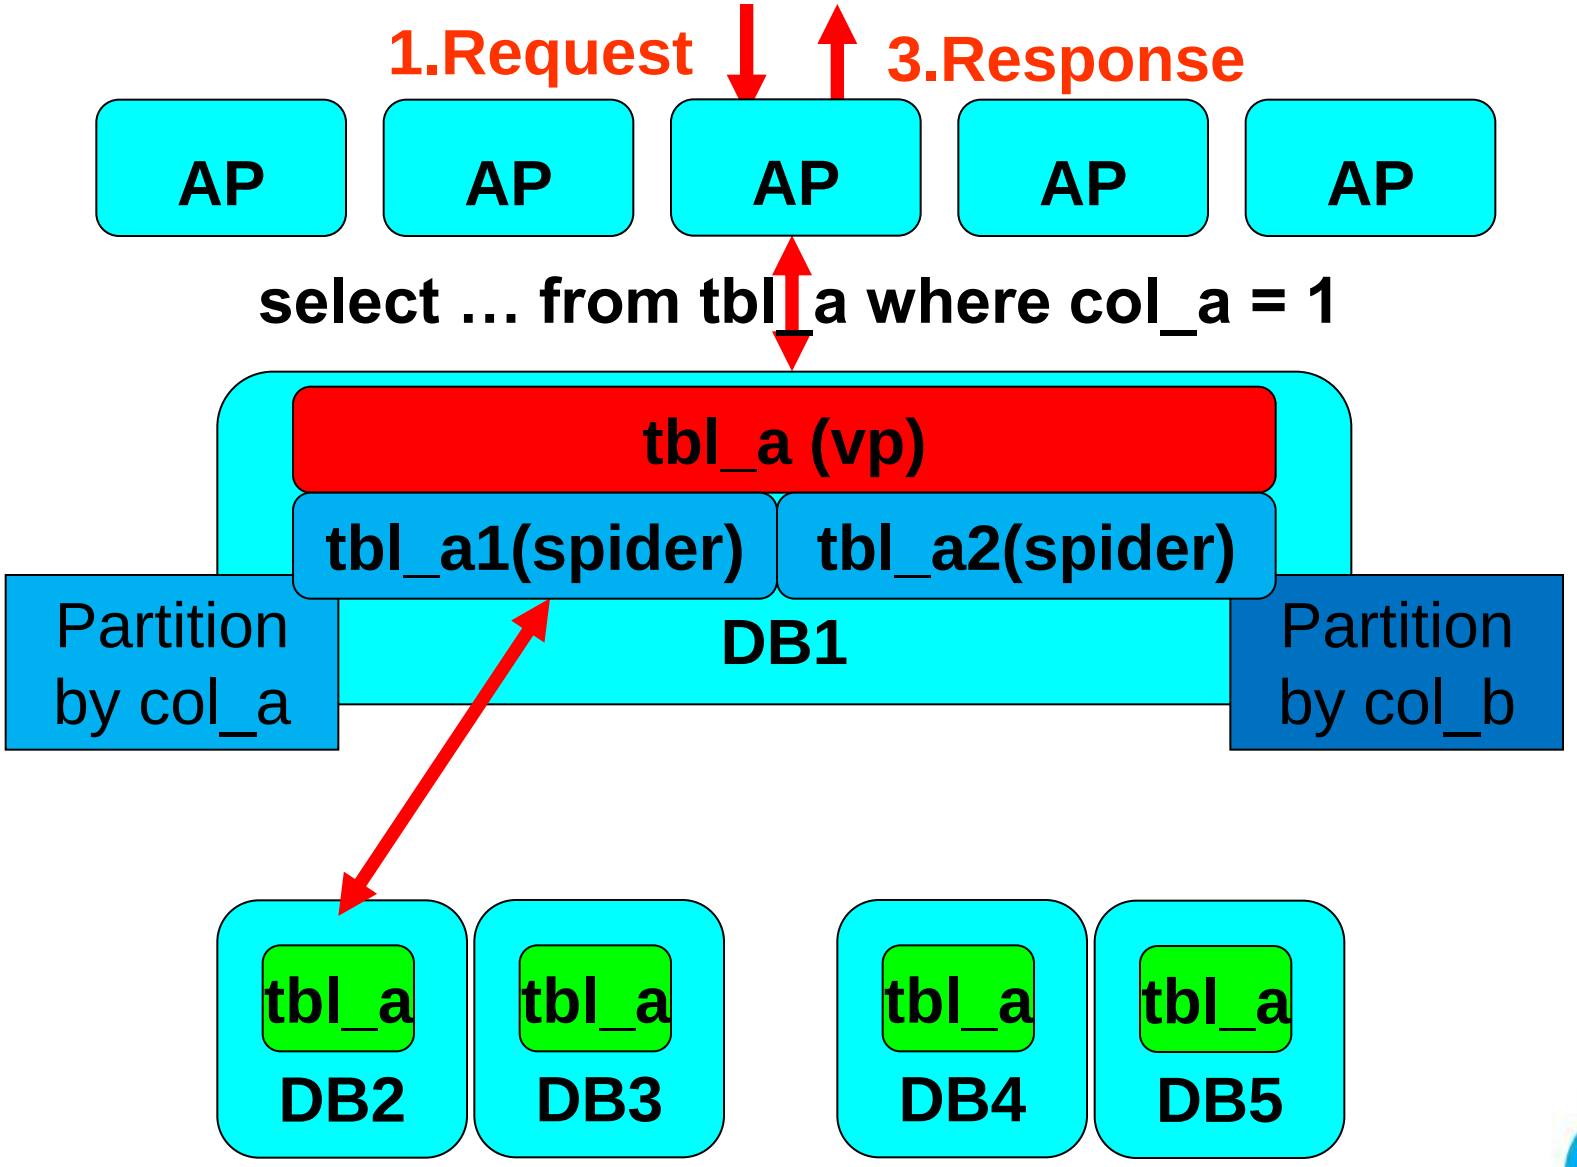

Structure sample of using different sharding rules (2/2)