

# Updatable Views

## Proposal

Hive will consider a view updatable if:

- The view refers to exactly one base table or updatable view in the FROM clause without a WHERE clause.
- Each column in the view is a column in the underlying table/updatable view with no underlying columns duplicated.
- Views must have the same partition columns as the underlying table/updatable view.

When inserting into a view:

- If a view does not specify all underlying columns, NULL will be inserted for each column not specified.

Given

```
create table t1 (id int, key string, value string) partitioned by (ds string, hr string);  
  
create view v partitioned on (ds, hr) as select id, value, ds, hr from t1;
```

we can insert into v where NULL will be used for key.

Notes:

- Should we try to support views like (create view v partitioned on (ds, hr) as select id, value, ds where ds='2011-01-01' and hr='12') for updating, where we infer values for ds and hr?
- With non-dynamic partitioning, do we require the partition be on each view in the updatable view chain? This seems burdensome if you don't have write access to all the views?
- When we specify dynamic partitions for a view, do we create partitions on each view in a chain of updatable views? If we don't, there may be strange behavior where SHOW PARTITIONS may not show anything on a view, but we can insert into such partitions of a view. If we do, drop partition on the view actually does nothing to the data.

See [Hive Views](#) for general information about views.