

Merge Engine (functionality test)

merge storage engine is a very simple engine that allows to "merge" multiple identical MyISAM tables. It has largely been superseded by partitioning engine which is more flexible (and works with non-MyISAM engines). The advantages of merge engine:

- partitioning engine requires the partitioning key to be part of primary key. If we divide by filterId, that means that our primary key in Source and Object would have to be composite: (objectId, filterId)
- it is not immediately clear how to turn existing MyISAM tables into partitioned table.

Advantages of partitioning:

- it is more clever, eg, if data is partitioned by filterId and query selects rows by filterId, it'll only examine the partition with a given filterId

The test described below was run on lsst-dbdev5 at ncsa (which has 16 GB of RAM), using 10 Source tables, total size of these table: ~30GB (2.8G, 1.1G, 62M, 6.0G, 4.5G, 164M, 1.5G, 6.2G, 5.7G, 1.3G). So, data does not fit in memory.

```
create table mergeTest (  
  `sourceId` bigint(20) NOT NULL,  
  `scienceCcdExposureId` bigint(20) DEFAULT NULL,  
  `filterId` tinyint(4) NOT NULL,  
  `objectId` bigint(20) DEFAULT NULL,  
  `movingObjectId` bigint(20) DEFAULT NULL,  
  `procHistoryId` int(11) NOT NULL,  
  `ra` double NOT NULL,  
  `raErrForDetection` float DEFAULT NULL,  
  `raErrForWcs` float NOT NULL,  
  `decl` double NOT NULL,  
  `declErrForDetection` float DEFAULT NULL,  
  `declErrForWcs` float NOT NULL,  
  `xFlux` double DEFAULT NULL,  
  `xFluxErr` float DEFAULT NULL,  
  `yFlux` double DEFAULT NULL,  
  `yFluxErr` float DEFAULT NULL,  
  `raFlux` double DEFAULT NULL,  
  `raFluxErr` float DEFAULT NULL,  
  `declFlux` double DEFAULT NULL,  
  `declFluxErr` float DEFAULT NULL,  
  `xPeak` double DEFAULT NULL,  
  `yPeak` double DEFAULT NULL,  
  `raPeak` double DEFAULT NULL,  
  `declPeak` double DEFAULT NULL,  
  `xAstrom` double DEFAULT NULL,  
  `xAstromErr` float DEFAULT NULL,  
  `yAstrom` double DEFAULT NULL,  
  `yAstromErr` float DEFAULT NULL,  
  `raAstrom` double DEFAULT NULL,  
  `raAstromErr` float DEFAULT NULL,  
  `declAstrom` double DEFAULT NULL,  
  `declAstromErr` float DEFAULT NULL,  
  `raObject` double DEFAULT NULL,  
  `declObject` double DEFAULT NULL,  
  `taiMidPoint` double NOT NULL,  
  `taiRange` float DEFAULT NULL,  
  `psfFlux` double NOT NULL,
```

```

`psfFluxErr` float NOT NULL,
`apFlux` double NOT NULL,
`apFluxErr` float NOT NULL,
`modelFlux` double NOT NULL,
`modelFluxErr` float NOT NULL,
`petroFlux` double DEFAULT NULL,
`petroFluxErr` float DEFAULT NULL,
`instFlux` double NOT NULL,
`instFluxErr` float NOT NULL,
`nonGrayCorrFlux` double DEFAULT NULL,
`nonGrayCorrFluxErr` float DEFAULT NULL,
`atmCorrFlux` double DEFAULT NULL,
`atmCorrFluxErr` float DEFAULT NULL,
`apDia` float DEFAULT NULL,
`Ixx` float DEFAULT NULL,
`IxxErr` float DEFAULT NULL,
`Iyy` float DEFAULT NULL,
`IyyErr` float DEFAULT NULL,
`Ixy` float DEFAULT NULL,
`IxyErr` float DEFAULT NULL,
`snr` float NOT NULL,
`chi2` float NOT NULL,
`sky` float DEFAULT NULL,
`skyErr` float DEFAULT NULL,
`extendedness` float DEFAULT NULL,
`flux_PS` float DEFAULT NULL,
`flux_PS_Sigma` float DEFAULT NULL,
`flux_SG` float DEFAULT NULL,
`flux_SG_Sigma` float DEFAULT NULL,
`sersicN_SG` float DEFAULT NULL,
`sersicN_SG_Sigma` float DEFAULT NULL,
`e1_SG` float DEFAULT NULL,
`e1_SG_Sigma` float DEFAULT NULL,
`e2_SG` float DEFAULT NULL,
`e2_SG_Sigma` float DEFAULT NULL,
`radius_SG` float DEFAULT NULL,
`radius_SG_Sigma` float DEFAULT NULL,
`flux_flux_SG_Cov` float DEFAULT NULL,
`flux_e1_SG_Cov` float DEFAULT NULL,
`flux_e2_SG_Cov` float DEFAULT NULL,
`flux_radius_SG_Cov` float DEFAULT NULL,
`flux_sersicN_SG_Cov` float DEFAULT NULL,
`e1_e1_SG_Cov` float DEFAULT NULL,
`e1_e2_SG_Cov` float DEFAULT NULL,
`e1_radius_SG_Cov` float DEFAULT NULL,
`e1_sersicN_SG_Cov` float DEFAULT NULL,
`e2_e2_SG_Cov` float DEFAULT NULL,
`e2_radius_SG_Cov` float DEFAULT NULL,
`e2_sersicN_SG_Cov` float DEFAULT NULL,
`radius_radius_SG_Cov` float DEFAULT NULL,
`radius_sersicN_SG_Cov` float DEFAULT NULL,
`sersicN_sersicN_SG_Cov` float DEFAULT NULL,
`flagForAssociation` smallint(6) DEFAULT NULL,
`flagForDetection` smallint(6) DEFAULT NULL,
`flagForWcs` smallint(6) DEFAULT NULL,
`chunkId` int(11) DEFAULT NULL,
`subChunkId` int(11) DEFAULT NULL,
INDEX (`sourceId`),
KEY `IDX_scienceCcdExposureId` (`scienceCcdExposureId`),
KEY `IDX_filterId` (`filterId`),
KEY `IDX_movingObjectId` (`movingObjectId`),
KEY `IDX_objectId` (`objectId`),

```

```

KEY `IDX_procHistoryId` (`procHistoryId`),
KEY `IDX_Source_decl` (`decl`)

) ENGINE=MERGE UNION (Source_6630, Source_6631, Source_6797, Source_6800,
                      Source_6801, Source_6802, Source_6968, Source_6970,
                      Source_6971, Source_6972);

## note above, the schema is a verbatim copy of one of the Source tables,
## but the primary key has been turned into non-unique index

select count(*) from mergeTest;
+-----+
| count(*) |
+-----+
| 63646581 |
+-----+
1 row in set (0.00 sec)

select count(*) from mergeTest where flux_PS > 0.01; # 4 min 54.10 sec
select count(*) from Source_6630 where flux_PS > 0.01; # 1.84
select count(*) from Source_6631 where flux_PS > 0.01; # 12.08
select count(*) from Source_6797 where flux_PS > 0.01; # 0.04
select count(*) from Source_6800 where flux_PS > 0.01; # 58.94
select count(*) from Source_6801 where flux_PS > 0.01; # 39.12
select count(*) from Source_6802 where flux_PS > 0.01; # 1.57
select count(*) from Source_6968 where flux_PS > 0.01; # 17.09
select count(*) from Source_6970 where flux_PS > 0.01; # 1 min 9.75
select count(*) from Source_6971 where flux_PS > 0.01; # 1 min 7.44
select count(*) from Source_6972 where flux_PS > 0.01; # 15.32
# total for 10 separate tables: 4 min 53.19

## queries like
select * from mergeTest where objectId = 1233;
select * from mergeTest where sourceId = 2867930259261118;

## come back instantaneously

```