


# Datahub GMA: Strongly Consistent Secondary Index (SCSI) Support

Jyoti Wadhvani

# Motivation

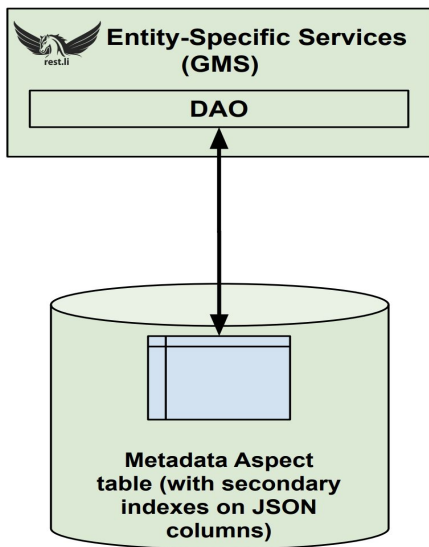
- Build a mission-critical metadata store
- Systems with UI facing elements e.g. Dynamic Configuration UI
- Search index is eventually consistent - serves stale data and leads to bad user experience



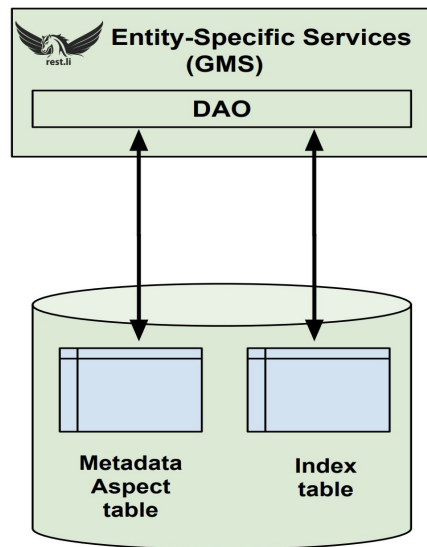
# Alternate Designs

- Generate a new column for every queryable attribute
- JSON column indexes

```
{  
  "firstName": "John",  
  "lastName": "Smith",  
  "username": "jsm"  
  "title": "Engineer"  
}
```



Design1



Design2

# Architecture

## Schema of the index table

```
CREATE TABLE metadata_index (  
  `id` BIGINT NOT NULL AUTO_INCREMENT,  
  `urn` VARCHAR(500) NOT NULL,  
  `aspect` VARCHAR(200) NOT NULL,  
  `path` VARCHAR(500) NOT NULL,  
  `longVal` BIGINT,  
  `stringVal` VARCHAR(500),  
  `doubleVal` DOUBLE,  
  
  CONSTRAINT id_pk PRIMARY KEY (id),  
  INDEX longIndex (`urn`, `aspect`, `path`, `longVal`),  
  INDEX stringIndex (`urn`, `aspect`, `path`, `stringVal`),  
  INDEX doubleIndex (`urn`, `aspect`, `path`, `doubleVal`)  
)
```

id	urn	aspect	path	longVal	stringVal	doubleVal
0	urn:li:corpuser:jsm	com.linkedin.common.urn.CorpuserUrn	/username	null	jsm	null
1	urn:li:corpuser:jsm	com.linkedin.identity.CorpUserInfo	/firstName	null	John	null
2	urn:li:corpuser:jsm	com.linkedin.identity.CorpUserInfo	/lastName	null	Smith	null

# INSERT and GET operations

## INSERT

- Update the primary document store and index table in the same transaction
- Only certain attributes of a given metadata aspect will be indexed as defined in the storage config

```
{
  "aspectStorageConfigMap": {
    "com.linkedin.identity.CorpUserInfo": {
      "pathStorageConfigMap": {
        "/username": {
          "strongConsistentSecondaryIndex": true
        }
      }
    }
  }
}
```

## GET

- *Step1*: Fetch all urns from the index table that satisfy given filter conditions
- *Step2*: Join urns with the primary document store (aspect table) to fetch metadata

# Future Work

Extend read-after-write consistency in NoSQL DBs leveraging native secondary indexes (e.g. MongoDB)

# Acknowledgements



Mars Lan



Shirshanka Das



Kerem Sahin