AsterixDB:
A Big Data Management System
"One size fits a Bunch"

https://asterixdb.ics.uci.edu
https://asterixdb.incubator.apache.org

At a Glance

- Open Source (Apache Incubating!) Big Data Management System
- Native support for data ingestion (data feeds)
- Runs on large commodity clusters
- Designed for mass quantities of semi-structured data
- Highly scalable storage and index management
- Native support for rich data types and operations (*e.g.*, spatial & temporal data)
- Native support for similarity queries

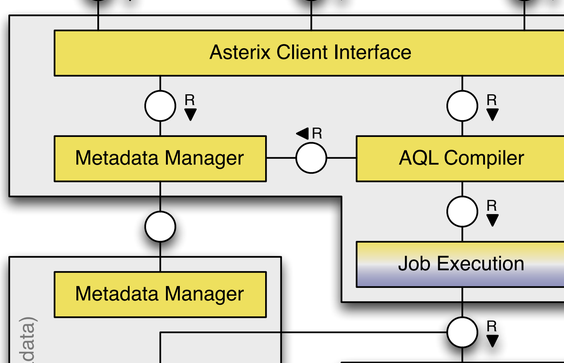

AQL (AsterixDB Query Language) + ADM

- ADM is a superset of JSON with a richer set of types (*e.g.*, bags, spatial data, temporal data, text) and optional schemas
- AQL is a powerful query language for semi-structured data, influenced by the best parts of W3C's XQuery
  - *Ex: List the user name and messages sent by those users who joined the Mugshot social network in a certain time frame:*

```
for $user in dataset MugshotUsers
where $user.user-since >= datetime('2010-07-22T00:00:00')
 and $user.user-since <= datetime('2012-07-29T23:59:59')
return {
 "uname" : $user.name,
 "messages" :
    for $message in dataset MugshotMessages
    where $message.author-id = $user.id
    return $message.message
};
```

System Architecture

- Uses the Hyracks data parallel platform as its runtime engine
- Shared-nothing storage
- Built for commodity clusters
- AQL uses Algebricks to optimize queries
- Each Node Controller stores a partitioned portion of the data stored in each index

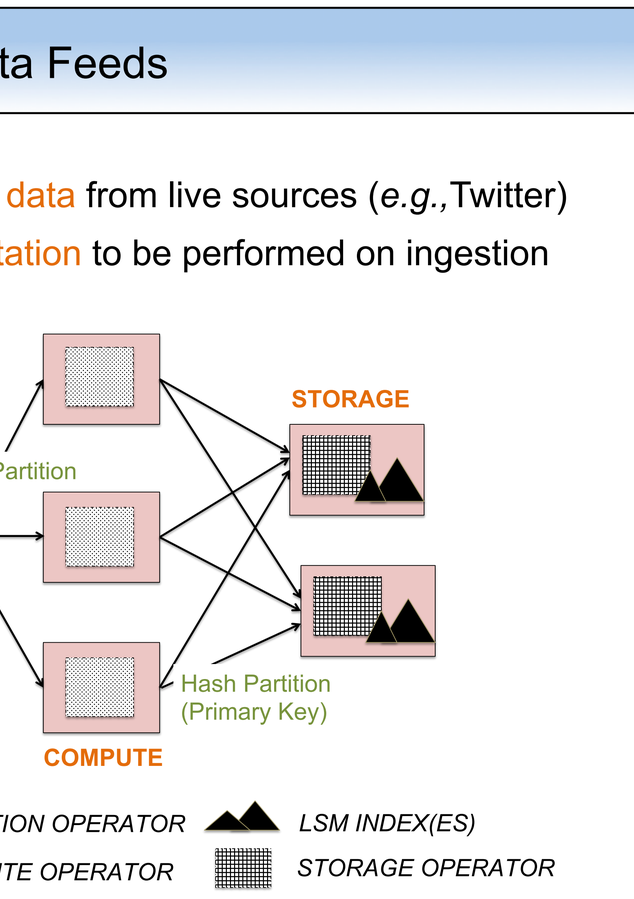

LSM-Based Storage and Indexing

- Log-Structured Merge technology to support high data ingestion rates
- *All* primary and secondary index types are LSM-ified (B+Tree, R-Tree, Inverted Indexes, …)

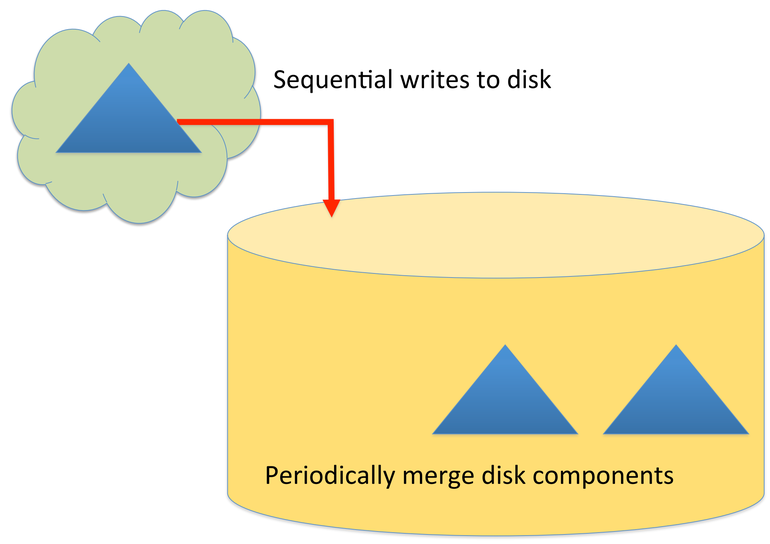

Data Feeds

- Integrated system for intake of data from live sources (*e.g.,*Twitter)
- Allows for user-defined computation to be performed on ingestion

Competitive Performance

	Users	Messages	Tweets
Asterix (*Schema*)	192	120	330
Asterix (*KeyOnly*)	360	240	600
Syst-X	290	100	495
Hive	38	12	25
Mongo	240	215	478

Table 2: Dataset sizes (in GB)

Batch Size	Asterix *Schema*	Asterix *KeyOnly*	Syst-X	Mongo
1	0.091	0.093	0.040	0.035
20	0.010	0.011	0.026	0.024

Table 4: Average insert time per record (in sec)

	Asterix *Schema*	Asterix *KeyOnly*	Syst-X	Hive	Mongo
Rec Lookup	0.03	0.03	0.12	(379.11)	0.02
Range Scan — with IX	79.47 0.10	148.15 0.10	148.33 4.90	11717.18 (11717.18)	175.84 0.05
Sel-Join (Sm) — with IX	78.03 0.51	96.76 0.55	55.01 2.13	333.56 (333.56)	66.46 0.62
Sel-Join (Lg) — with IX	79.62 2.24	99.73 2.32	56.65 10.59	350.92 (350.92)	273.52 14.97
Sel2-Join (Sm) — with IX	79.06 0.50	97.82 0.52	55.81 2.62	340.02 (340.02)	66.45 0.61
Sel2-Join (Lg) — with IX	80.18 2.32	101.24 2.32	56.10 10.70	394.11 (394.11)	313.17 15.28
Agg (Sm) — with IX	128.66 0.16	232.30 0.17	130.64 0.14	83.18 (83.18)	400.97 0.19
Agg (Lg) — with IX	128.71 5.53	232.41 5.55	132.19 4.67	94.11 (94.11)	401 8.34
Grp-Aggr (Sm) — with IX	130.20 0.45	232.77 0.46	131.18 0.17	127.85 (127.85)	398.27 0.20
Grp-Aggr (Lg) — with IX	130.62 5.96	234.10 5.91	133.02 4.72	140.21 (140.21)	400.10 9.03

Table 3: Average query response time (in sec)

Sponsored by: eBay, Facebook, Google, HTC, Microsoft, NSF, Oracle Labs, Yahoo!, Infosys, Amazon Web Services