



From Batch to Streaming ET(L) with Apache Apex

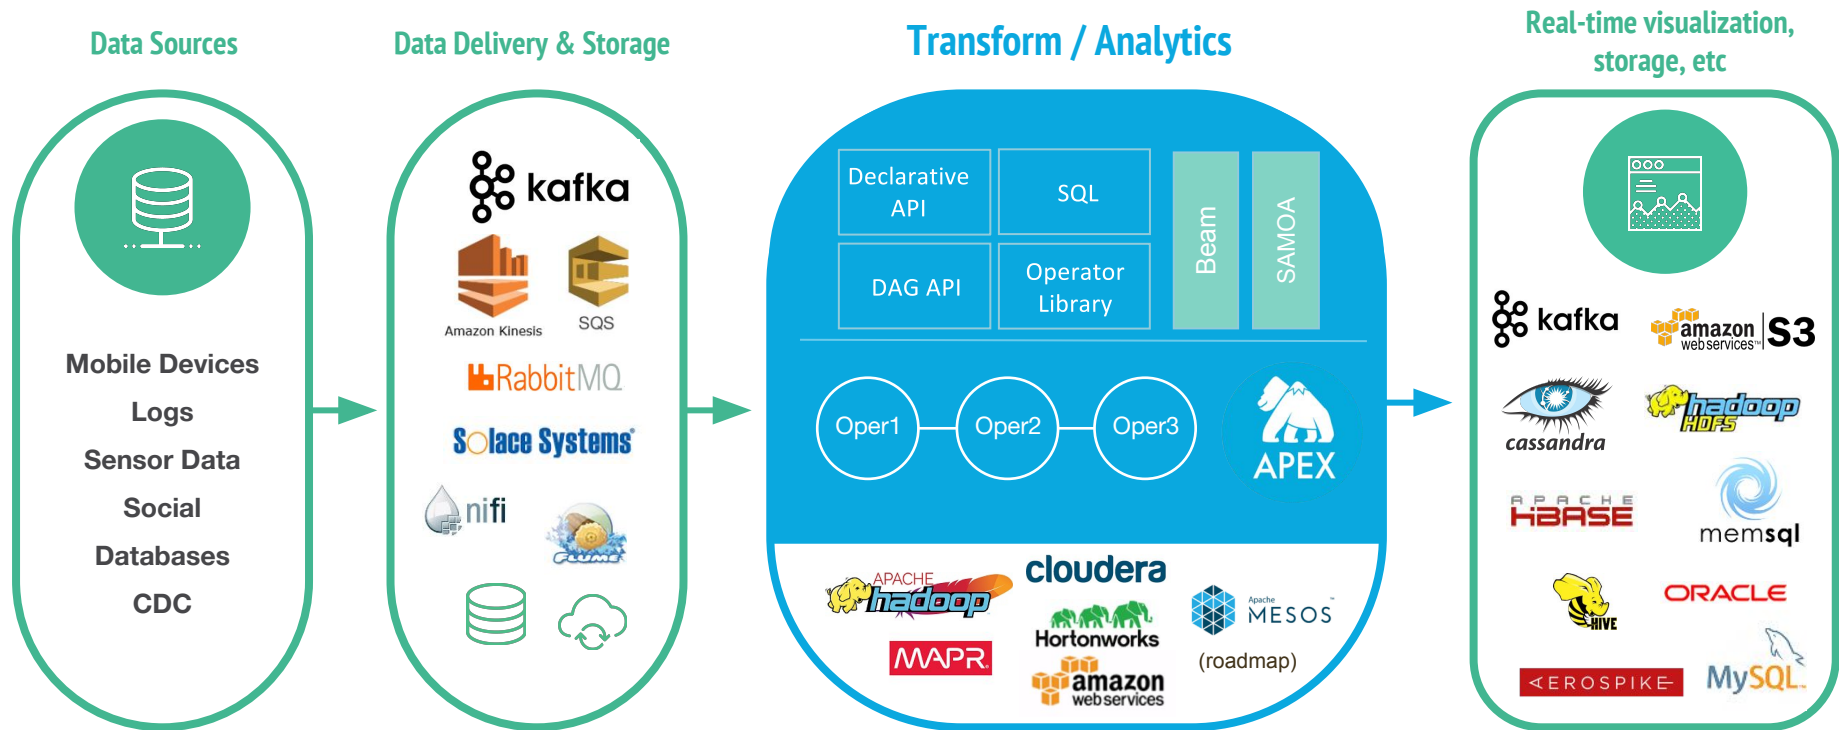
Thomas Weise

Apache Apex PMC Chair

thw@apache.org

[@thweise](#) [@atrato_io](#)

Stream Data Processing with Apache Apex



Use Case

✓ Real-time reporting

- Reporting of critical metrics around campaign monetization
- Revenue, impression & click info
- Aggregate counters & reporting on top N metrics
- Low latency querying using Kafka

in pub-sub model.

✓ Real-time Monitoring

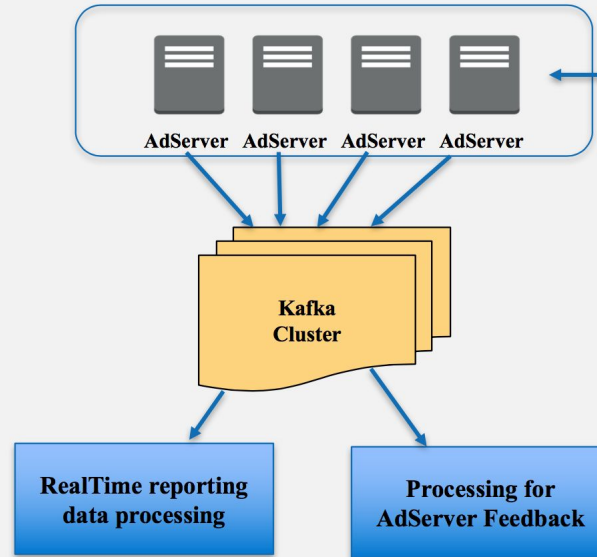
- Alerts on deal tracking & monetization
- Campaign & deal health

✓ Real-time Learning

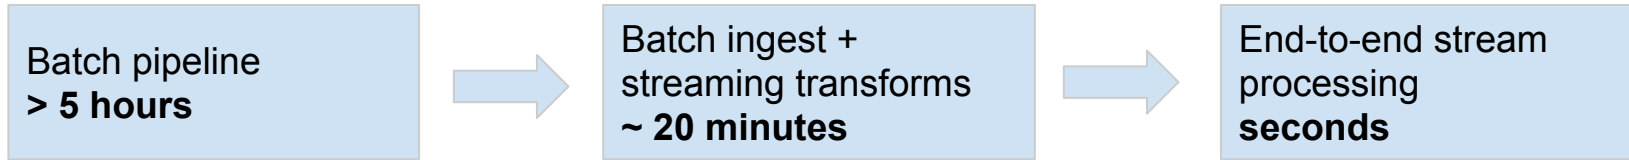
- Using the lost bid insights for price recommendations.

✓ Allocation Engine

- Feedback to ad serving for guaranteed delivery & line item pacing



Phased Transition



Batch processing with several hours till insight:

- Available data stale, does no longer apply to current situation
- Current data stuck in batch pipeline
- Complex batch processing orchestration with many different components
- Hours of delay translate to high cost due to inability to make timely campaign adjustments.

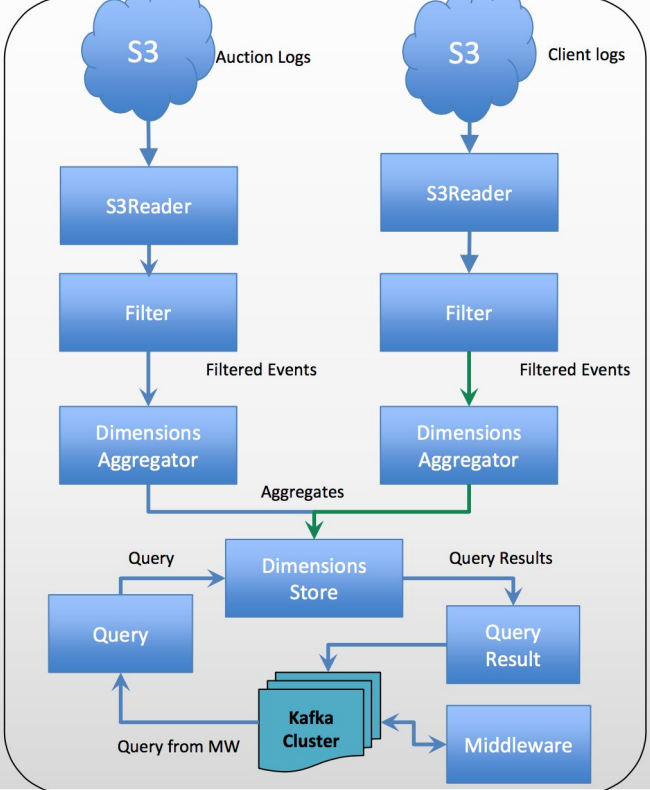


Processing with Apex, reuse batch ingestion:

- Existing ingestion mechanism (files in S3, shared with other pipelines)
- Migrate transform logic to Apex
- Enable reporting from application state (“Queryable State”)
- Reduced latency, valuable as intermediate step.

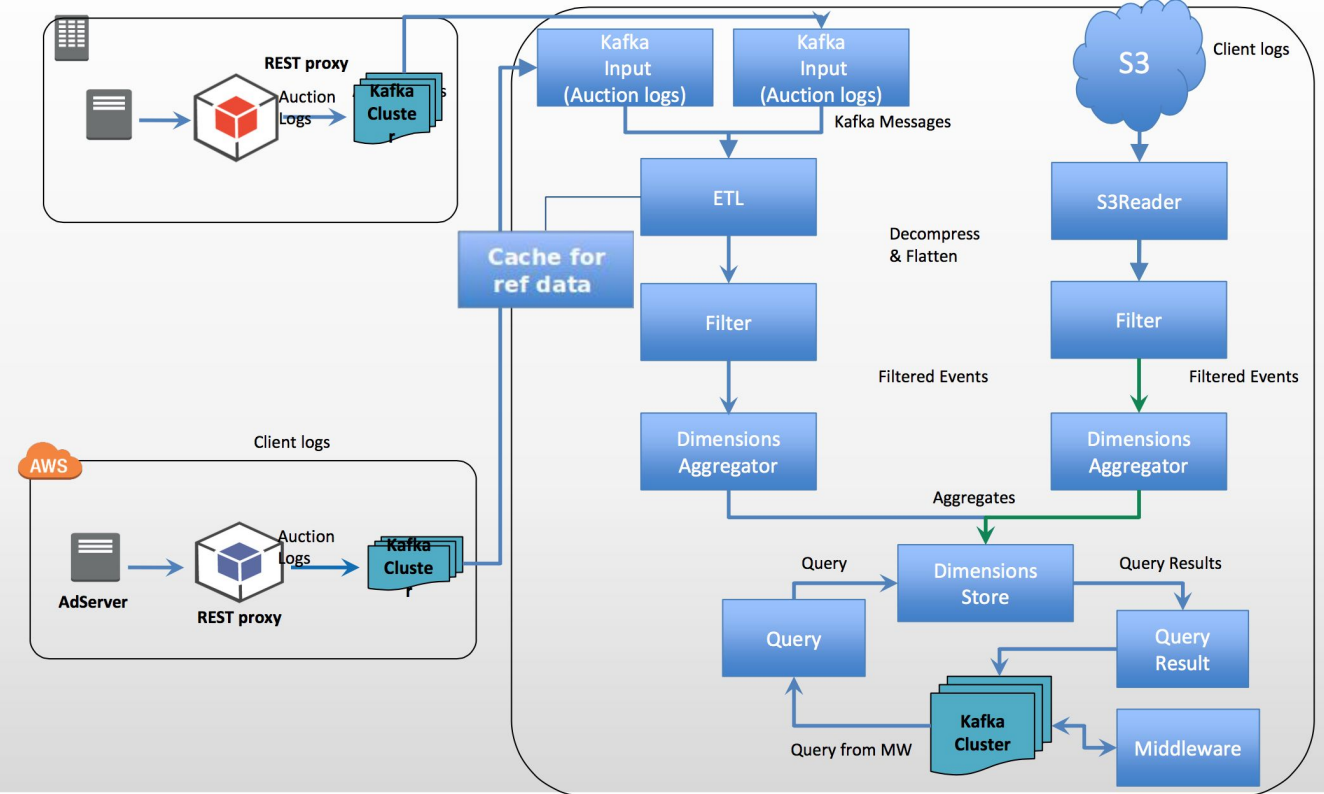
Streaming source and processing:

- Data comes directly from Kafka clusters
- Significantly reduced latency
- Balanced load (no ingestion spikes)
- Reduced resource consumption with Apex support for multi-cluster Kafka consumers
- Reporting meets SLA requirements

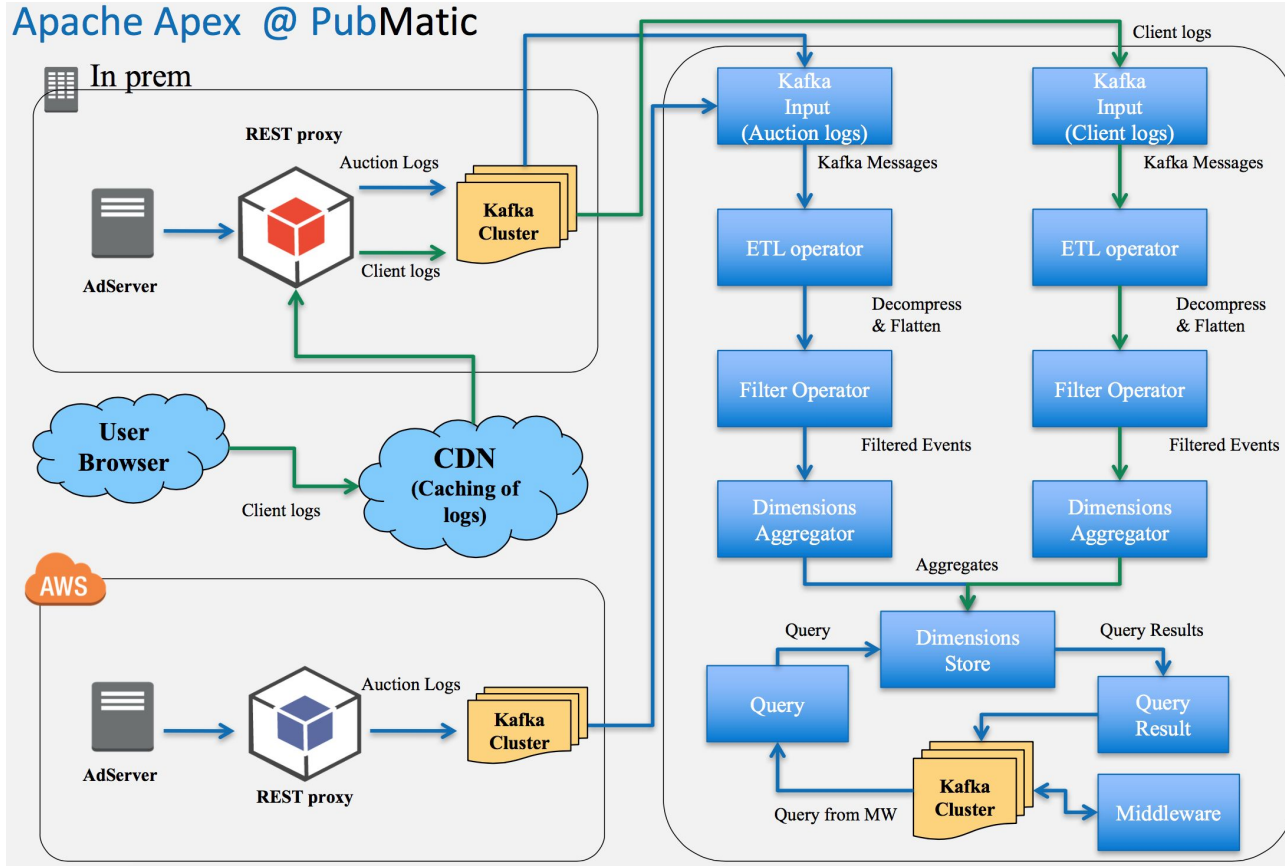
Phase 1 (batch ingest)




Phase 2 (hybrid)



Real-time Streaming



Real-time Dashboard


Dashboard Analytics Inventory & Rules Transac

Dashboard

View: All Sites | All Deals | Date Range: Real Time | Live Data: On

Key Performance Indicators


Revenue



Impressions



eCPM



PubMatic Dashboard Analytics Inventory & Rules Transactions Ad Quality

Data below is from 12:00 am to now.

Top Deals

Revenue Impressions eCPM

Deals by % of Total	Revenue	% of Total
AO_DBM_Blocklist_ROS-BTF_10.30.15	\$15.76	3.76%
AO_MediaMath_MediaMath_Blocklist_R...	\$15.55	3.72%
AO_DataXu_DataXu_Blocklist_ROS-AT...	\$5.34	1.27%
AO_Maxpoint_Maxpoint_Blocklist_BTF...	\$4.29	1.02%
AO_MP_DBM_Blocklist_ROS-Homepag...	\$4.01	0.95%
AO_DBM_Blocklist_ROS-300x250-ATF...	\$3.73	0.89%
AO_YAMP_YAMP_Blocklist_ROS-BTF...	\$3.61	0.86%
AO_MediaMath_MediaMath_Homepage...	\$3.56	0.85%
AO_TTD_TTD_Blocklist_ROS-BTF_12.1...	\$3.23	0.77%
AO_TTD_TTD_Blocklist_ROS-300x250-...	\$3.01	0.71%

Top 10 deals 14.80% All deals

[View All Deals](#)

Top Demand Sources

View: Buyers

Revenue Impressions eCPM

Buyers by % of Total	Impressions	% of Total
Not Provided	55,103	25.76%
RocketFuel	52,069	24.34%
Simplifi	29,748	13.90%
Quantcast	10,469	4.89%
Yahoo Ad Manager Plus	7,768	3.63%
ANY BUYER	6,665	3.11%
DBM2	6,151	2.87%
ZiffDavis, Inc	5,666	2.64%
Kepler (Kayak, Comwave, Progressiv...	2,906	1.35%
Adobe	2,768	1.29%

Top 10 Buyers 83.78% All Buyers

[View Detailed Report](#)

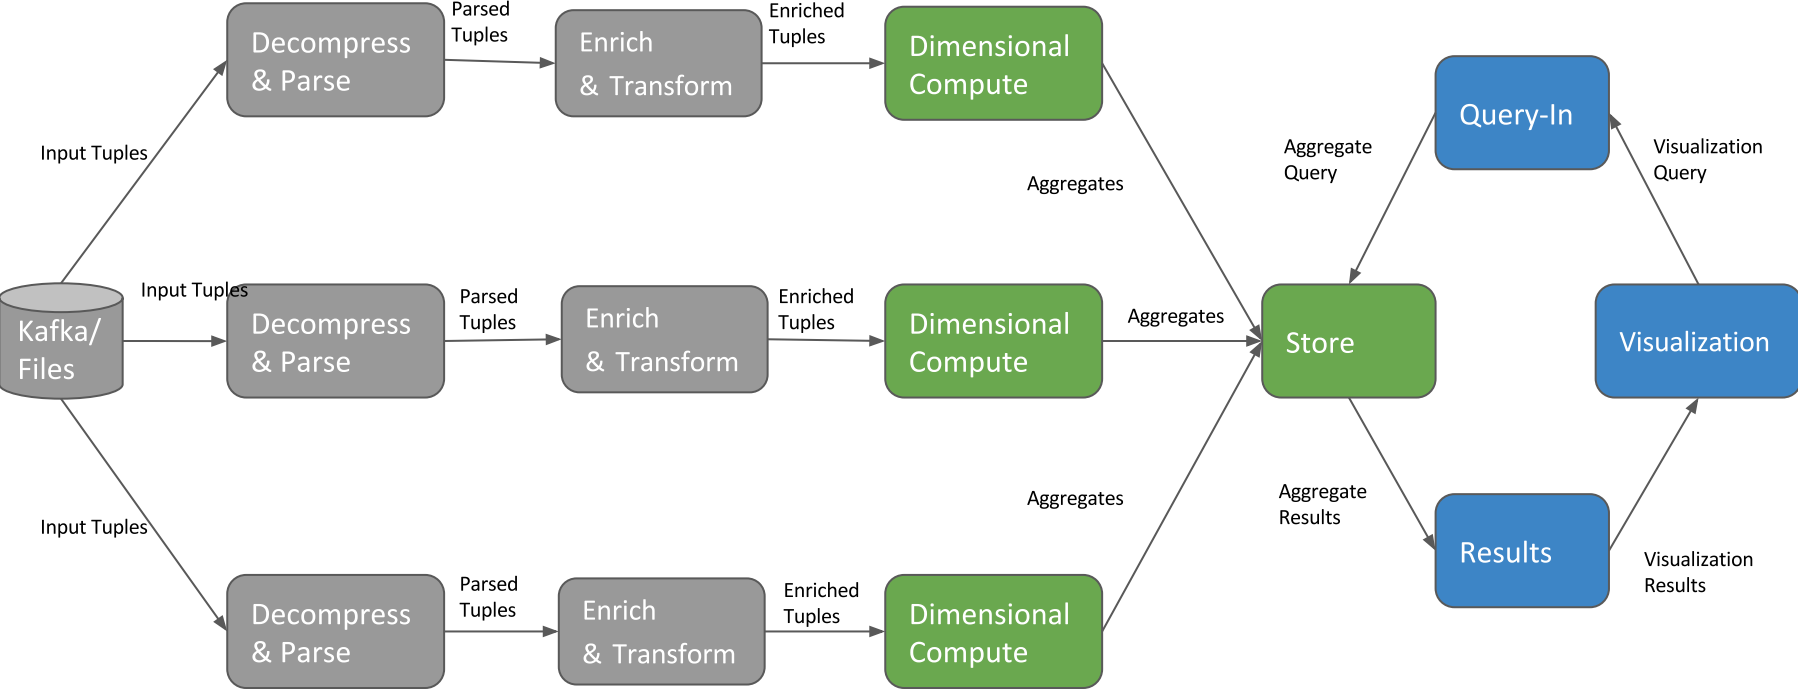
Paid Impressions

25,129

eCPM

\$1.96

Pipeline Transformations



<https://www.slideshare.net/ApacheApex/actionable-insights-with-apache-apex-at-apache-big-data-2017-by-devendra-tagare>

Dimension Computation

Advertiser: Subway
Location: WA
Cost: 2
Revenue: 1
Impressions: 5
Clicks: 1
Time: 10:15:30

hour	advertiser	location	cost	revenue	impr	clicks
10:00			6	9 => 10	22	3
10:00	Burger King		4	6	12	2
10:00	Subway		2	3 => 4	10	2
10:00		CA	4	6	15	3
10:00		WA	2	3 => 4	7	1
10:00	Burger King	CA	2	3	5	1
10:00	Burger King	WA	2	3	7	1
10:00	Subway	CA	2	3	10	2
10:00	Subway	WA		0 => 1		

Scale

- 6 geographically distributed data centers
- 10 PB of data under management
- 50 TB/day of data generated from auction & client logs
- 40+ billion ad impressions and 350+ billion bids per day
- Average data inflow of 450K events/sec
- 64 Kafka Input partitions, 32 instances of in-memory distributed store
- 1.2 TB of memory for the Apex application

Why Apex

- State Management & Fault tolerance
 - Exactly-once, Checkpointing and Windowing
 - Fine grained recovery, low-latency SLA support
 - Queryable state
- Processing based on event time
 - Accuracy, Repeatable/Replay
- Native Streaming
 - Low latency + high throughput, efficient resource utilization
 - Pipelined processing (data in motion)
- Scalability
 - Process more data by adding compute resources, no platform/architecture limits
 - Dynamic scaling and resource allocation, elasticity
- Library of connectors and transformations
 - Time to value

Apex Library

Messaging

- [Kafka](#)
- JMS (ActiveMQ etc.)
- Kinesis, SQS
- Flume, NiFi
- MQTT

File Systems

- [HDFS](#) / Hive
- Local File
- S3
- FTP

NoSQL

- Cassandra, HBase
- Aerospike, Accumulo
- Couchbase, CouchDB
- Redis, MongoDB
- Geode, Kudu

RDBMS

- [JDBC](#)
- MySQL
- Oracle
- MemSQL

Other

- Elastic Search
- Solr
- Twitter
- WebSocket / HTTP
- SMTP

Stateless Transformations

- Parsers: XML, JSON, CSV, Avro
- Filter
- Enrich
- Configurable POJO schema
- Map, FlatMap (custom Java function)
- Script (JavaScript, Jython)


Stateful Transformations

- Windowing: sliding, tumbling, session
- Accumulations: sum, merge, join, sort, top n, ...
- Triggering, Watermarks
- Dimensional Aggregations (with state management for historical data + query)
- Deduplication

How to build it



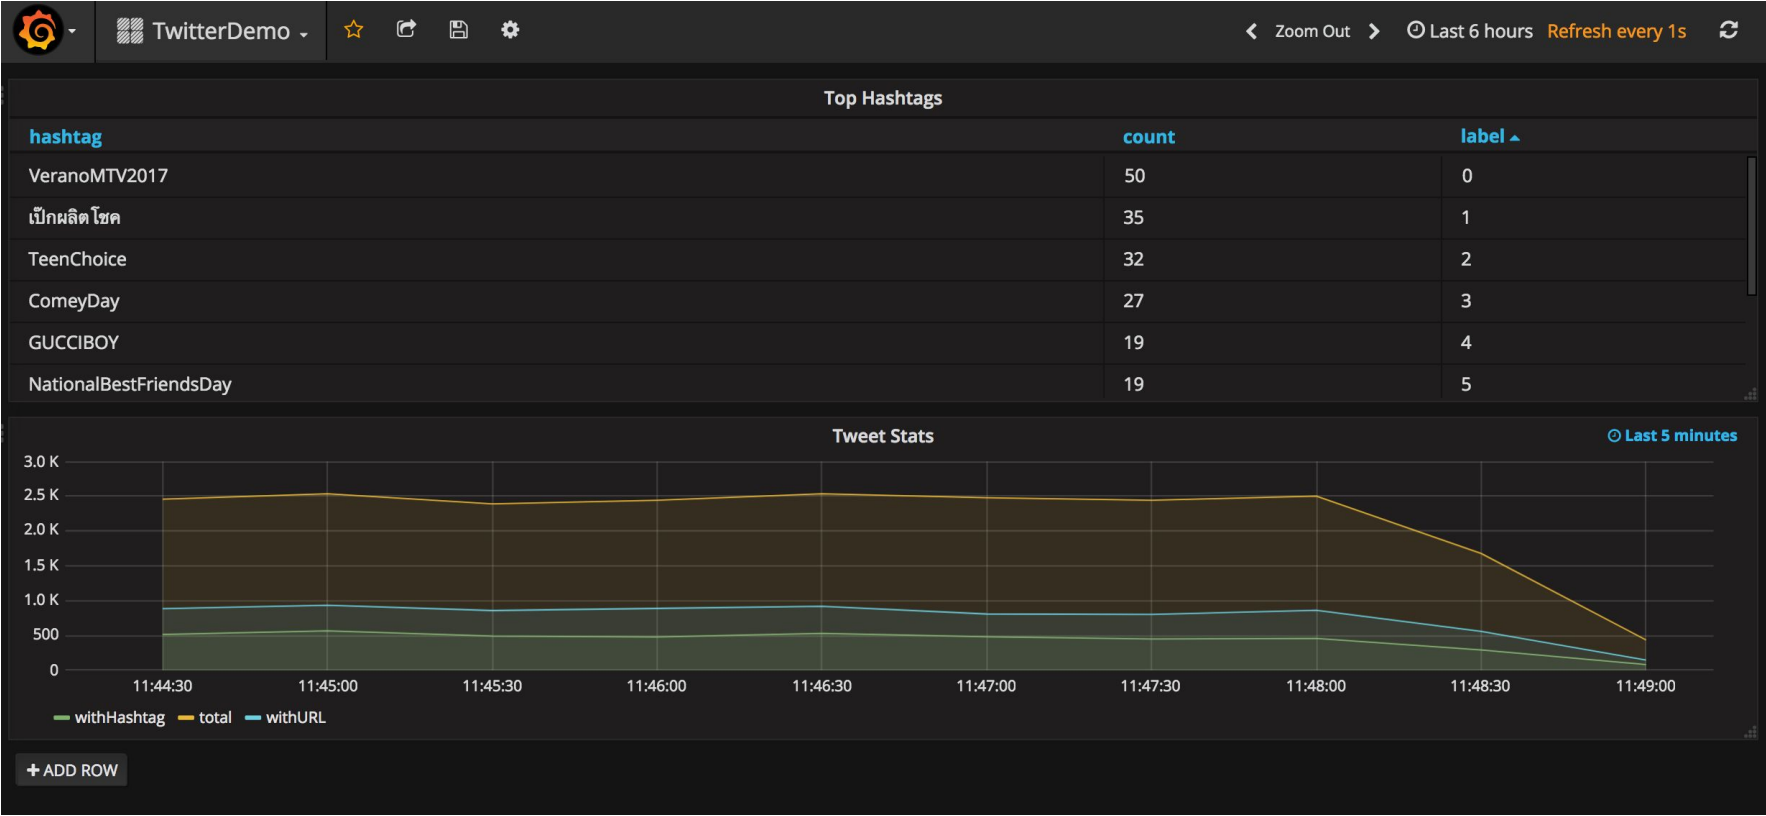
Example Application (Twitter)



- Top N hashtags
- Tweet stats time series
- Queryable state
- WebSocket Pub/Sub
- Visualization with Grafana

Source code: <https://github.com/tweise/apex-samples/tree/master/twitter>

Real-time Visualization



Top Hashtags

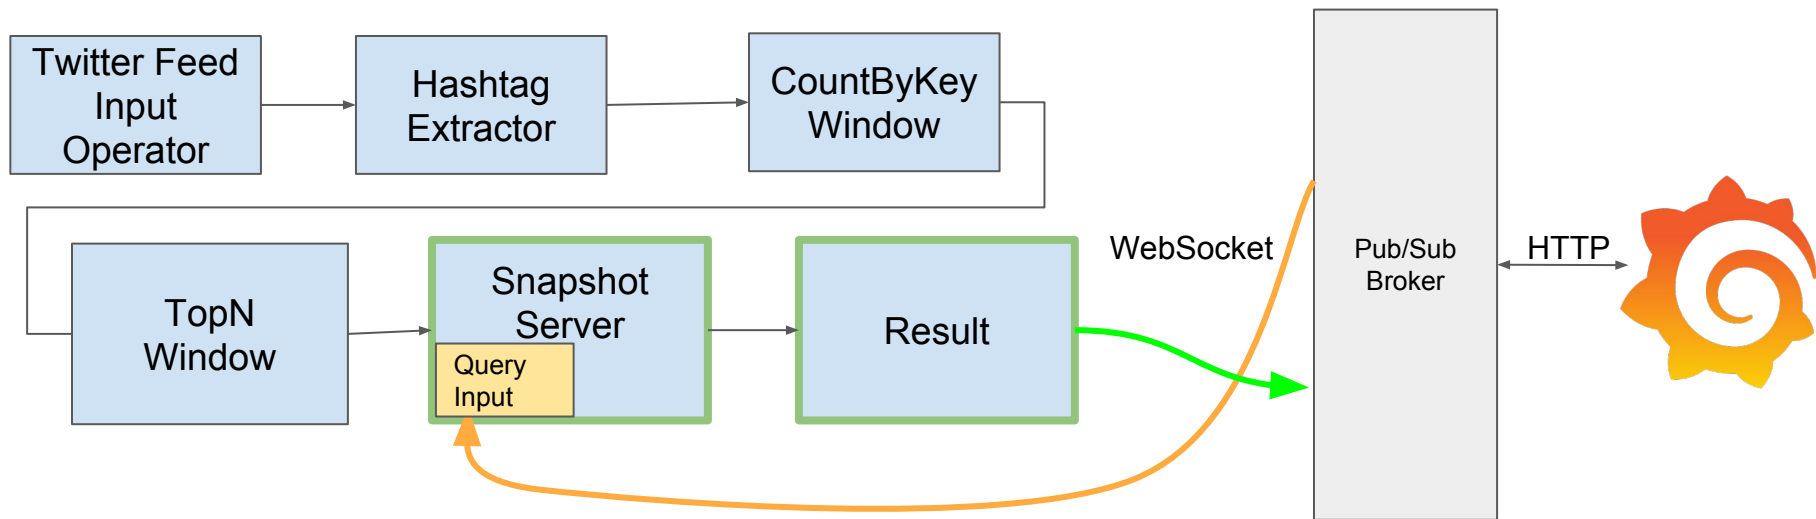
- Keyed sum accumulation (5 minute window, count trigger)
- TopN accumulation of upstream windowed counts

```
KeyedWindowedOperatorImpl<String, Long, MutableLong, Long> countByKey =
    new KeyedWindowedOperatorImpl<>();
countByKey.setAccumulation(new SumLong());
countByKey.setDataStorage(new InMemoryWindowedKeyedStorage<String, MutableLong>());
countByKey.setWindowOption(new WindowOption.TimeWindows(Duration.standardMinutes(5)));
countByKey.setWindowStateStorage(new InMemoryWindowedStorage<WindowState>());
countByKey.setTriggerOption(TriggerOption.AtWatermark().
    withEarlyFiringsAtEvery(25).accumulatingFiredPanels());

TopNByKey<String, Long> topNByKey = new TopNByKey<>();
topNByKey.setN(10);
WindowedOperatorImpl<KeyValPair<String, Long>, Map<String, Long>, List<KeyValPair<String, Long>>>
    topN = new WindowedOperatorImpl<>();
topN.setAccumulation(topNByKey);
topN.setDataStorage(new InMemoryWindowedStorage<Map<String, Long>>());
topN.setWindowStateStorage(new InMemoryWindowedStorage<WindowState>());
topN.setTriggerOption(TriggerOption.AtWatermark().
    .withEarlyFiringsAtEvery(1).accumulatingFiredPanels());
```

Queryable State

A set of operators in the library that support real-time queries of operator state.



- Pub/Sub server: <https://github.com/atrato/pubsub-server>
- Grafana data source: <https://github.com/atrato/apex-grafana-datasource-server>

Queryable State

- Snapshot server
 - Stateful operator that holds last received list of objects
 - Receives query and emits the list as JSON formatted query result
- Source schema configured, result fields via query
- Predefined schemas (Apex library): “Snapshot”, “Dimensional”

```
TopNSnapshotServer topTagsSnapshot = new TopNSnapshotServer();
String JSON = SchemaUtils.jarResourceFileToString(SCHEMA); {
topTagsSnapshot.setSnapshotSchemaJSON(JSON);
    "values": [
        {"name": "hashtag", "type": "string"},
        {"name": "count", "type": "long"},
        {"name": "label", "type": "string"}
    ]
}
```

```
PubSubWebSocketAppDataResult wsResult = dag.addOperator("topNResult",
    new PubSubWebSocketAppDataResult());
wsResult.setUri(uri);
wsResult.setTopic(topic);
```

Demo

```
Thomass-MBP:twitter thomas$  
Thomass-MBP:twitter thomas$  
Thomass-MBP:twitter thomas$ mvn test -Dtest=TwitterStatsAppTest
```

I

Apex - Recent Additions & Roadmap



- Apex runner in Apache Beam
 - Iterative processing
 - Integrated with Apache Samoa, opens up ML
 - Integrated with Apache Calcite, enables SQL
 - Scalable, incremental state management
 - User defined control tuples (watermarks, batch control, ...)
-

- Enhanced support for Batch Processing
- Support for Mesos and Kubernetes
- Encrypted Streams
- Support for Python

Resources



- <http://apex.apache.org/>
- Powered by Apex - <http://apex.apache.org/powered-by-apex.html>
- Learn more - <http://apex.apache.org/docs.html>
- Getting involved - <http://apex.apache.org/community.html>
- Download - <http://apex.apache.org/downloads.html>
- Follow @ApacheApex - <https://twitter.com/apacheapex>
- Meetups - <https://www.meetup.com/topics/apache-apex/>
- Examples - <https://github.com/apache/apex-malhar/tree/master/examples>
- Slideshare - <http://www.slideshare.net/ApacheApex/presentations>