



Storm

Daniel Hagimont
hagimont@enseeiht.fr

1

This lecture is about the Storm stream processing framework.

Originally, Storm was developed by a startup which was acquired by Twitter.

Later the project was open sourced in the Apache foundation.

This lecture is a rapid presentation of Storm and a demonstration (without any labwork).

Apache Storm

- Processing of data streams
 - Real-time (on the fly)
 - Pretty much like Spark streaming
- Developed in Clojure (a dialect of Lisp)
- As Hadoop, Spark or Spark-streaming
 - Distributed
 - Reliable
 - Open-source (Apache)
 - Used by large companies (e.g. Twitter)



2


Storm is a platform for the real time processing of data streams. It has many similarities with Spark Streaming regarding the goals, but differs regarding the implementation strategy.

Storm was developed in Clojure (functional programming).

It shares many characteristics with big data platforms like Hadoop, Spark or Spark Streaming: distributed reliable, opensource and used by large companies.

Apache Storm: basic concepts

- **Spouts: data sources** (can be an application)



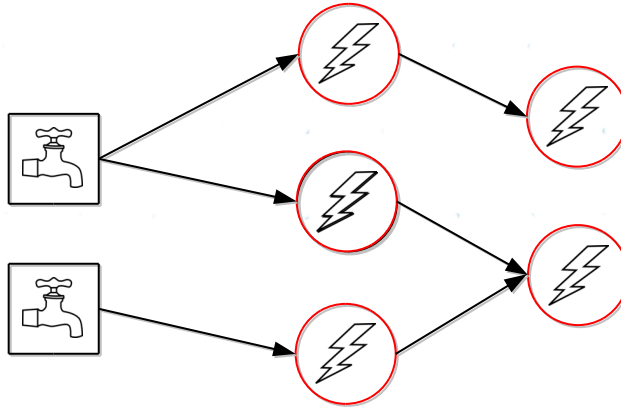
3

Its architecture relies on 2 types of nodes (nodes here do not refer to machines, but to daemons which include processing) : Spouts and Bolts.

A Spout is a data source (in red in the figure), which may acquire data from a file or a network connection. A spout can emit data in direction of several bolts.

Apache Storm: basic concepts

■ Bolts: data processing




4

A Bolt is a data processing unit which includes a particular treatment on received data. It can emit data in direction of several other bolts.

Apache Storm: basic concepts

- **Topology:** interconnection of spouts and bolts
 - Represent an Apache Storm application
- Execute indefinitely (unlike MapReduce applications)



5

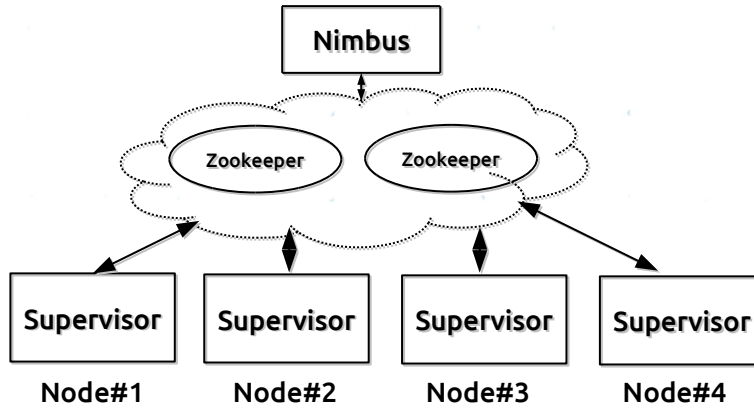
A topology is an interconnection of spouts and bolts.

It corresponds to an global application.

It executes indefinitely like Spark Streaming applications.

Apache Storm: architecture

- A Storm cluster



6

A Storm cluster is composed of a global coordinator called Nimbus and a set of execution engines (called supervisors) which hosts spouts and bolts.

In the middle, zookeeper is a distributed middleware used for coordination between Nimbus and Supervisor nodes.

Apache Storm: architecture

- **Nimbus (comparable to Yarn's ResourceManager)**
 - Entry point of a Storm cluster
 - Receive and deploy topologies (applications)
 - Receive monitoring information from nodes
 - Re-schedule a topology in case of failure
- **Supervisor (comparable to Yarn's NodeManager)**
- **Zookeeper (coordination between them)**

7


Nimbus is the coordinator of the platform. It can be compared to Yarn's ResourceManager. It is responsible of the deployment (on nodes) and management of topologies.

Supervisors are comparable to Yarn's NodeManagers. They manage resources locally.

Zookeeper is a distributed middleware which provides a distributed naming scheme (like a filesystem) and the possibility to share (read and write) information and synchronize. It is used for distributed coordination of Nimbus and Supervisors.

Apache Storm: architecture

- On a Node runs a Supervisor
 - Correspond to Yarn's NodeManager
 - Listen to reception of work
 - ◆ Worker: process
 - ◆ Executor: thread
 - ◆ Task: bolt/spout
 - Monitor resource usage
 - Can be dynamically added



A machine in a Storm cluster may run one or more worker processes for one or more topologies. Each worker process runs executors for a specific topology.

One or more executors may run within a single worker process, with each executor being a thread spawned by the worker process. Each executor runs one or more tasks of the same component (spout or bolt).

A task performs the actual data processing.

Each Node runs a Supervisor. It listens to the reception of jobs.


Jobs may be scheduled on processors. More precisely, Supervisors manage processes which host threads, a thread executing for one topology. Several spouts and bolts may be scheduled on one thread.

Supervisors monitor resource usage and report it to Nimbus.

Supervisors can be added dynamically in a cluster.

Apache Storm: data

- Spouts/bolts exchange Tuples
 - $\langle v1, v2, v3 \dots \rangle$
- A value in a Tuple can be of any serializable type
- Each spout/bolt defines the type of emitted tuples (1-tuple, 2-tuple ...) and field names
- Routing
 - A emitted tuple is sent on each of its outgoing connections




9

Spouts and bolts exchange tuples which include Java objects.

An emitted tuple (by a spout or a bolt) is sent to each outgoing connection.

Apache Storm: example

- The word count example
 - Keeps stats on words occurring in tweets or logs



10

Here is a topology which implements the Wordcount application.

`RandomSentenceSpout` is a spout which emits a sentence (1-tuple `<sentence>`) randomly from a set of sentence (to simulate the arrival of sentences from the net).

`SplitSentence` is a bolt which receive sentences, splits them into words and emits words (1-tuple `<word>`).

`WordCount` is a bolt which receives words, manages a wordcount table and emits each wordcount which is modified (2-tuple `<word, count>`).

Apache Storm: word count

```
public static class RandomSentenceSpout extends BaseRichSpout {
    SpoutOutputCollector _collector;
    Random _rand;

    public void open(Map conf, TopologyContext context, SpoutOutputCollector collector) {
        _collector = collector;
        _rand = new Random(); initialization
    }

    public void nextTuple() {
        Utils.sleep(100);
        String[] sentences = new String[]{ "the cow jumped over the moon", "an apple a day keeps the doctor
away", "four score and seven years ago", "snow white and the seven dwarfs", "i am at two with nature" };
        String sentence = sentences[_rand.nextInt(sentences.length)];
        _collector.emit(new Values(sentence)); emission of a sentence every 100ms
    }

    public void ack(Object id) {}

    public void fail(Object id) {}

    public void declareOutputFields(OutputFieldsDeclarer declarer) {
        declarer.declare(new Fields("sentence")); output: 1-tuple
    }
}
```

Here is the implementation of RandomSentenceSpout.

The spout receives with the open() method the collector reference which allows emitting data (the variable is save in a _collector variable).

Storm invokes the nextTuple() method all the time (iteratively). If there's nothing to send, it is supposed to wait a little time (not to waste too much CPU) and return.

ack() and fail() allow managing a queue of messages and to re-emit them if their processing failed.

Notice that the declareOutputFields() method allows declaring the fields of emitted tuples (here a 1-tuple with a "sentence" field).

Apache Storm: word count

```
public static class SplitSentence extends BaseBasicBolt {  
  
    public void execute(Tuple tuple, BasicOutputCollector collector) {  
        String sentence = tuple.getString(0);  
        StringTokenizer st = new StringTokenizer(sentence);  
        while (st.hasMoreTokens()) {  
            collector.emit(new Values(st.nextToken()));  
        }  
    }  
  
    public void declareOutputFields(OutputFieldsDeclarer declarer) {  
        declarer.declare(new Fields("word"));  
    }  
}
```

- A received tuple is a sentence
- The sentence is split into words
- Words are emitted

12

Here is the implementation of SplitSentence.

It defines the execute() method which includes the treatment to apply on received tuples :

- tuple is the received tuple
- collector is the variable for emitting tuples

The implementation of the method splits the received sentence into words and emits one tuple for each word.

As in the previous class, the declareOutputFields() method allows declaring the fields of emitted tuples (here a 1-tuple with a "word" field).

Apache Storm: word count

```
public static class WordCount extends BaseBasicBolt {
    Map<String, Integer> counts = new HashMap<String, Integer>();

    public void execute(Tuple tuple, BasicOutputCollector collector) {
        String word = tuple.getString(0);
        Integer count = counts.get(word); handling a received tuple
        if (count == null)
            count = 0;
        count++;
        counts.put(word, count);
        collector.emit(new Values(word, count));
    }

    public void declareOutputFields(OutputFieldsDeclarer declarer) {
        declarer.declare(new Fields("word", "count")); output: 2-tuple
    }
}
```

- A received tuple is a word
- A counter for each word in a Map
- The words and its counter are emitted

13

Here is the implementation of WordCount.

Like the previous bolt, it implements the execute() and declareOutputFields() methods.

declareOutputFields: it emits 2-tuples including "word" and "count" fields.

execute: it counts the word occurrences in a Hashtable and emits the new count for each received word.

Apache Storm: word count

```
public class WordCountTopology {  
  
    public static class RandomSentenceSpout extends BaseRichSpout {...}  
    public static class SplitSentence extends BaseBasicBolt {...}  
    public static class WordCount extends BaseBasicBolt {...}  
  
    public static void main(String[] args) throws Exception {  
  
        TopologyBuilder builder = new TopologyBuilder(); creation of the topology  
  
        builder.setSpout("spout", new RandomSentenceSpout(), 5);  
        builder.setBolt("split", new SplitSentence(), 8).shuffleGrouping("spout");  
        builder.setBolt("count", new WordCount(), 12).fieldsGrouping("split", new Fields("word"));  
  
        Config conf = new Config();  
        conf.setDebug(true); number of threads shuffle or field grouping  
  
        LocalCluster cluster = new LocalCluster();  
        cluster.submitTopology("word-count", conf, builder.createTopology());  
  
        Thread.sleep(10000);  
  
        cluster.shutdown();  
    }  
}
```

14

This is the main program. It defines the Storm topology.

It is composed of 5 instances of RandomSentenceSpout. Each instance is executed by a different thread.


It is composed of 8 instances of SplitSentence bolt. The routing of messages from RandomSentenceSpout instances to SplitSentence instances is defined with the shuffleGrouping policy.

It is composed of 12 instances of WordCount bolt. The routing of messages from SplitSentence instances to WordCount instances is defined with the fieldsGrouping policy (with field "word").

Grouping policies are explained in the next slide.

Apache Storm: routing

- Spouts/bolts execute in parallel (threads/tasks)



- When a tuple is sent from Bolt A to Bolt B, who receive it ?
 - Many modes (grouping)
 - shuffleGrouping: one tasks randomly chosen
 - fieldsGrouping: same field value=> same task
 - allGrouping, directGrouping, customGrouping ...
- For Wordcount: the same word goes to the same task

15

Spouts and bolts are instantiated several times and executed in parallel. Each instance is associated with a different thread.

The binding between a bolt A and a bolt B actually means a binding between a pool of instances of bolt A and a pool of instances of bolt B. The behavior of this binding is defined with a mode of grouping, defining which instance receives an emitted message.

- shuffleGrouping means that the target instance is randomly chosen
- fieldsGrouping with a given field means that tuples with the same value for that field should always be transmitted to the same target instance.
- allGrouping means messages are replicated and sent to all instances

There are many other grouping modes.

In the WordCount application, for the binding between SplitSentence instances and WordCount instances, the fieldsGrouping allows to always send the same word to the same instances (task), so that the count will be consistent.

Demonstration

- Install Apache Storm
 - tar xzf apache-storm-0.9.5.tar.gz
- Compile the WordCount application
 - javac -cp "../apache-storm-0.9.5/lib/*" WordCountTopology.java
 - Or in Eclipse
 - ◆ Add storm-core-0.9.5.jar in the buildpath
- Observe traces in debug mode
 - java -cp "../apache-storm-0.9.5/lib/*" WordCountTopology | grep "Emitting: count"

```
12199 [Thread-25-count] INFO backtype.storm.daemon.task - Emittng: count default [a, 45]
12199 [Thread-23-count] INFO backtype.storm.daemon.task - Emittng: count default [day, 45]
12199 [Thread-27-count] INFO backtype.storm.daemon.task - Emittng: count default [keeps, 45]
12199 [Thread-27-count] INFO backtype.storm.daemon.task - Emittng: count default [away, 45]
12199 [Thread-21-count] INFO backtype.storm.daemon.task - Emittng: count default [doctor, 45]
12199 [Thread-17-count] INFO backtype.storm.daemon.task - Emittng: count default [the, 204]
12260 [Thread-13-count] INFO backtype.storm.daemon.task - Emittng: count default [and, 113]
12261 [Thread-21-count] INFO backtype.storm.daemon.task - Emittng: count default [snow, 58]
12261 [Thread-9-count] INFO backtype.storm.daemon.task - Emittng: count default [seven, 113]
12261 [Thread-11-count] INFO backtype.storm.daemon.task - Emittng: count default [four, 56]
12261 [Thread-13-count] INFO backtype.storm.daemon.task - Emittng: count default [and, 114]
12261 [Thread-11-count] INFO backtype.storm.daemon.task - Emittng: count default [score, 56]
12261 [Thread-17-count] INFO backtype.storm.daemon.task - Emittng: count default [the, 205]
12261 [Thread-25-count] INFO backtype.storm.daemon.task - Emittng: count default [dwarfs, 58]
12261 [Thread-9-count] INFO backtype.storm.daemon.task - Emittng: count default [white, 58]
12261 [Thread-27-count] INFO backtype.storm.daemon.task - Emittng: count default [years, 56]
12261 [Thread-9-count] INFO backtype.storm.daemon.task - Emittng: count default [seven, 114]
12261 [Thread-9-count] INFO backtype.storm.daemon.task - Emittng: count default [ago, 56]
^C
hagnont@hagnont-pc:~/shared/international/vietnam/USTM/cours-cloud-bigdata/cours/bigdata/tp/5-demo-storm/wc$
```

16

Here is the procedure to execute a Storm topology locally.
Installation in a cluster is well documented.

Demonstration

- Install Apache Storm
 - tar xzf apache-storm-0.9.5.tar.gz
- Compile the WordCount application
 - javac -cp "../apache-storm-0.9.5/lib/*" WordCountTopology.java
 - Or in Eclipse
 - ◆ Add storm-core-0.9.5.jar in the buildpath
- Observe traces in debug mode
 - java -cp "../apache-storm-0.9.5/lib/*" WordCountTopology | grep "Emitting: count"

```
12199 [Thread-25-count] INFO backtype.storm.daemon.task - Emittng: count default [a, 45]
12199 [Thread-23-count] INFO backtype.storm.daemon.task - Emittng: count default [day, 45]
12199 [Thread-27-count] INFO backtype.storm.daemon.task - Emittng: count default [keeps, 45]
12199 [Thread-27-count] INFO backtype.storm.daemon.task - Emittng: count default [away, 45]
12199 [Thread-21-count] INFO backtype.storm.daemon.task - Emittng: count default [doctor, 45]
12199 [Thread-17-count] INFO backtype.storm.daemon.task - Emittng: count default [the, 204]
12260 [Thread-13-count] INFO backtype.storm.daemon.task - Emittng: count default [and, 113]
12261 [Thread-21-count] INFO backtype.storm.daemon.task - Emittng: count default [snow, 58]
12261 [Thread-9-count] INFO backtype.storm.daemon.task - Emittng: count default [seven, 113]
12261 [Thread-11-count] INFO backtype.storm.daemon.task - Emittng: count default [four, 56]
12261 [Thread-13-count] INFO backtype.storm.daemon.task - Emittng: count default [and, 114]
12261 [Thread-11-count] INFO backtype.storm.daemon.task - Emittng: count default [score, 56]
12261 [Thread-17-count] INFO backtype.storm.daemon.task - Emittng: count default [the, 205]
12261 [Thread-25-count] INFO backtype.storm.daemon.task - Emittng: count default [dwarfs, 58]
12261 [Thread-9-count] INFO backtype.storm.daemon.task - Emittng: count default [white, 58]
12261 [Thread-27-count] INFO backtype.storm.daemon.task - Emittng: count default [years, 56]
12261 [Thread-9-count] INFO backtype.storm.daemon.task - Emittng: count default [seven, 114]
12261 [Thread-9-count] INFO backtype.storm.daemon.task - Emittng: count default [ago, 56]
^C
hagnmont@hagnmont-pc:~/shared/international/vietnam/USTM/cours-cloud-bigdata/cours/bigdata/tp/5-demo-storm/wc$
```

17

Conclusion

- Distributed stream processing platform
 - Defining topologies
 - Replication for parallelism
 - Routing of data
- Recently: spark-streaming
 - Same objective : stream processing
 - Different approach
 - ◆ Spark-streaming : batch of data (RDDs) processed by Spark
 - ◆ Storm : data are routed through a topology
 - Spark benefits from a large eco-system
 - Storm is known to be more efficient

18

In conclusion, Spark-streaming and Storm share the same objective (stream processing) but they have different approaches.

Spark-streaming slices incoming data into batches (RDDs) which are processed with the Spark engine.

Storm deploys computing components (instances of spout or bolt) in the cluster and routes incoming data in the topology.

Spark is known to be part of a very rich ecosystem

Storm is known to be more efficient