



## Solr on Cassandra

COSCUP/GNOME.Asia 2010  
gasol@pixnet.tw

<http://0rz.tw/5kl2E>



關於我  
@gasolwu



不嫌Java囉唆



又喜歡Python的簡捷



且對Android有愛



開始進入正題



你的網站有內容了還不夠！





# 你的網站有內容了還不夠！

還要讓使用者找得到才行...



# 搜尋的重要性！

交給Google就行了嗎？



# Solr and Cassandra?



事情是這樣的...






- 使用者建議
- 愈來愈多使用者嵌外站服務
- 只有個人，沒有全站
- PV Up Up



那就做吧, Solution?





吃大蒜哪有不嘴臭的道理



# Solr

- created by Yonik Seeley at CNET Networks
- Contributed to Apache in January 2006
- the Lucene and Solr projects merged In March 2010
- current 1.4.1 (with lucene 2.9.3)
- Web admin interface
- many feature.





# Powerful full-text search

<http://localhost:8080/solr/select?q=title:coscup>

```
<result name="response" numFound="21" start="0" maxScore="15.267826">
  <doc>
    <float name="score">4.7711954</float>
    <str name="id">17206-24959116</str>
    <str name="title">VIM Hacks - c9s在COSCUP的講題</str>
  </doc>
  <doc>
    <float name="score">8.096988</float>
    <str name="id">1893496-27550711</str>
    <str name="title">COSCUP 09' 精簡心得, COSCUP 萬歲!</str>
  </doc>
  <doc>
    <float name="score">4.7711954</float>
    <str name="id">232580-24907067</str>
    <str name="title">COSCUP 2009開源人年會參後心得</str>
  </doc>
  <doc>
    <float name="score">4.7711954</float>
    <str name="id">232580-24906103</str>
    <str name="title">COSCUP 2009開源人年會參後心得</str>
  </doc>
  <doc>
    <float name="score">4.7711954</float>
    <str name="id">630252-29042632</str>
    <str name="title">COSCUP 2009 開源人年會小記</str>
  </doc>
</result>
```



趴xml太麻煩?

水管太小?



# JSON result

/select?q=title:coscup&wt=json

```
{ "response": { "numFound": 21,
               "start": 0,
               "maxScore": 15.267826,
               "docs": [
                 { "id": "17206-24959116", "title": "VIM Hacks - c9s在
COSCUP的講題", "score": 4.7711954},
                 { "id": "1893496-27550711", "title": "COSCUP 09' 精簡
心得, COSCUP 萬歲!", "score": 8.096988},
                 { "id": "232580-24907067", "title": "COSCUP 2009開源
人年會參後心得", "score": 4.7711954},
                 { "id": "232580-24906103", "title": "COSCUP 2009開源
人年會參後心得", "score": 4.7711954},
                 { "id": "630252-29042632", "title": "COSCUP 2009 開源
人年會小記", "score": 4.7711954}]
               }
}
```



# Multiple keyword

/select?q=title:coscup+title:心得&wt=json

```
{ "response": {  
  "numFound": 3,  
  "start": 0,  
  "maxScore": 8.46245,  
  "docs": [  
    { "id": "1893496-27550711", "title": "COSCUP 09' 精簡心得, COSCUP 萬歲! ",  
      "score": 8.46245 },  
    { "id": "232580-24907067", "title": "COSCUP 2009 開源人年會參後心得",  
      "score": 5.259093 },  
    { "id": "232580-24906103", "title": "COSCUP 2009 開源人年會參後心得",  
      "score": 5.259093 }  
  ]  
}
```



# Filter Query

`/select?q=title:coscup&fq=category:2`



# Range Query

/select?q=title:coscup+date:[\* TO NOW]

/select?q=mac+mini+price:[0 TO 19900]



# Query Boost

/select?q=title:老虎<sup>5</sup>+OR+title:老鼠

# Index Boost

```
<add>  
  <doc boost="2.5">  
    <field name="id">1234567</field>  
    <field name="title" boost="2.0">Coscup  
2010</field>  
  </add>
```



# Highlighting

/select?q=title:coscup+title:心得&hl=true&hl.fl=title

```
"highlighting":{  
  "1893496-27550711":{ "title":["<em>COSCUP</em> 09' 精簡<em>心得</em>,<br><em>COSCUP</em> 萬歲! " ]},  
  "232580-24907067":{ "title":["<em>COSCUP</em> 2009開源人年會參後<em>心得<br></em>"]},  
  "232580-24906103":{ "title":["<em>COSCUP</em> 2009開源人年會參後<em>心得<br></em>"]  
}
```





# Facet

/select?q=title:coscup+title:心得  
&facet=true&facet.fl=category

The screenshot shows a search interface with a navigation bar containing '全部', '相簿', '相片', '影音', '部落格', and '文章'. A search bar contains 'coscup' and a search button. Below the search bar, there are tabs for '文章' and 'coscup, 總共 42 筆'. A facet menu titled '文章分類' is highlighted with a red box, listing categories and their counts: '所有分類(42)', '不設分類(23)', '數位生活(14)', '心情日記(2)', '生活綜合(1)', 'KUSO(1)', and '攝影寫真(1)'. Below the facet menu, there are sorting options '依關聯度' and '依更新日期'. A blue button '快去報名 coscup 囉!' is visible, followed by a paragraph of text: '今天早上一到公司, 就想說要快點報名今年的 coscup, 我一上去報名的時候, 已經 11 點左右了...繼續閱讀 »'.



# Replication

## master

```
<requestHandler name="/replication" class="solr.ReplicationHandler"
>
  <lst name="master">
    <str name="replicateAfter">commit</str>
    <str name="confFiles">schema.xml,stopwords.txt</str>
  </lst>
</requestHandler>
```

## slave

```
<requestHandler name="/replication" class="solr.ReplicationHandler"
>
  <lst name="slave">
    <str name="masterUrl">http://foo:8080/solr/replication</str>
    <str name="pollInterval">02:30:00</str>
  </lst>
</requestHandler>
```



# Others


- Caching (filter, query, document)
- Web administration interface
- Distributed search (sharding)
- Spell Checking, More Like This




# What is Cassandra?

- Key-value store (with BigTable like structure)
- highly scalable and available
- decentralized and distributed
- Eventually consistent
- 2 famous paper
  - BigTable (data model)
  - Dynamo (distribution architecture)



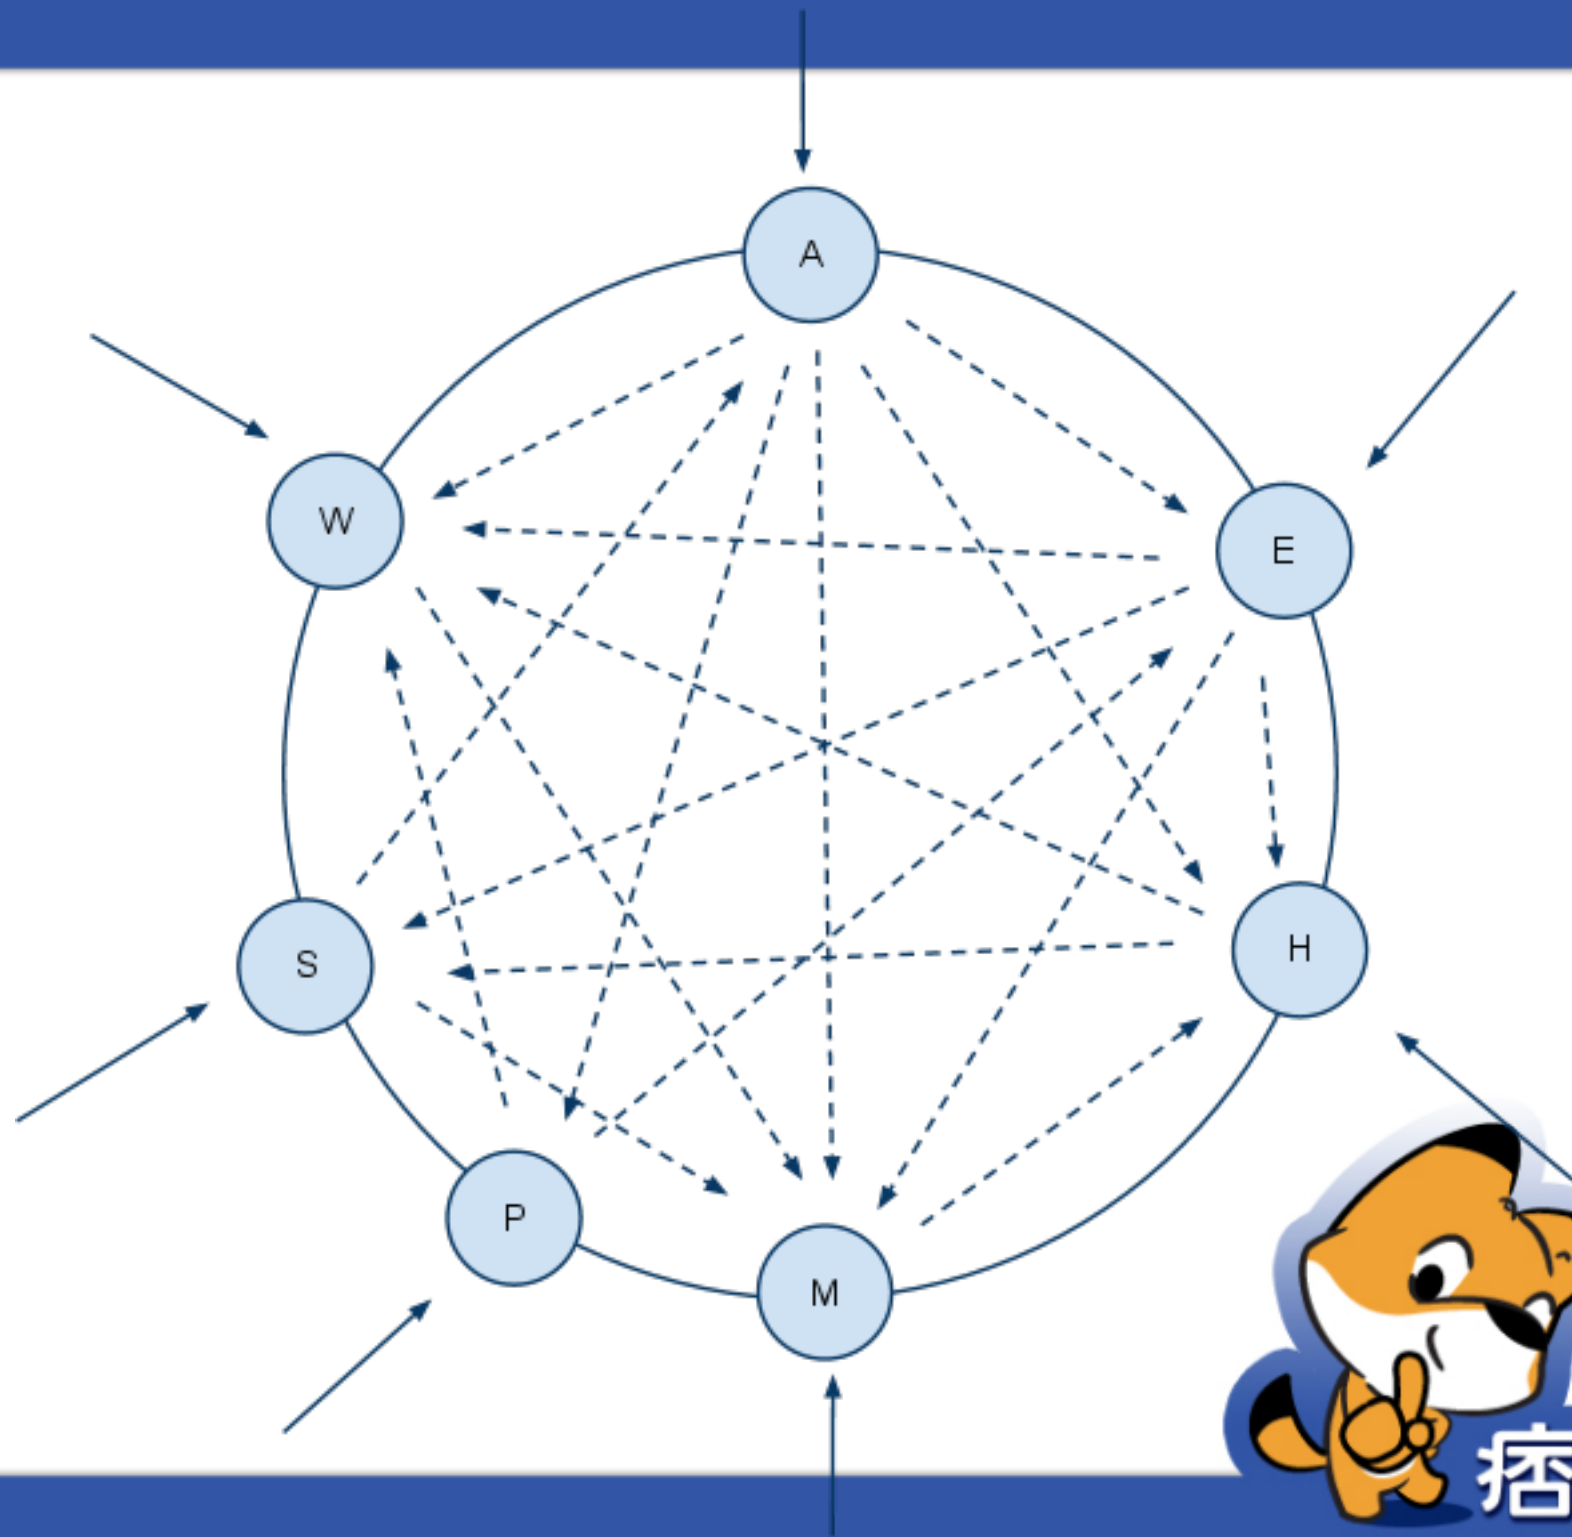


N = 3



Key "L"





# Partitioning

## RandomPartitioner

Tokens are integers in the range  $0-2^{127}$

$\text{md5}(\text{Key}) \rightarrow \text{Token}$

```
> nodetool -h localhost -p 8080 ring
Address      Status      Load          Range          Ring
10.1.2.30    Up          12.46 GB     130918853719569435579439903801217043383
10.1.2.29    Up          8.65 GB      73363270894447861520032770908021394916  |<--|
10.1.2.31    Up          4.07 GB      130918853719569435579439903801217043383  |-->|
```

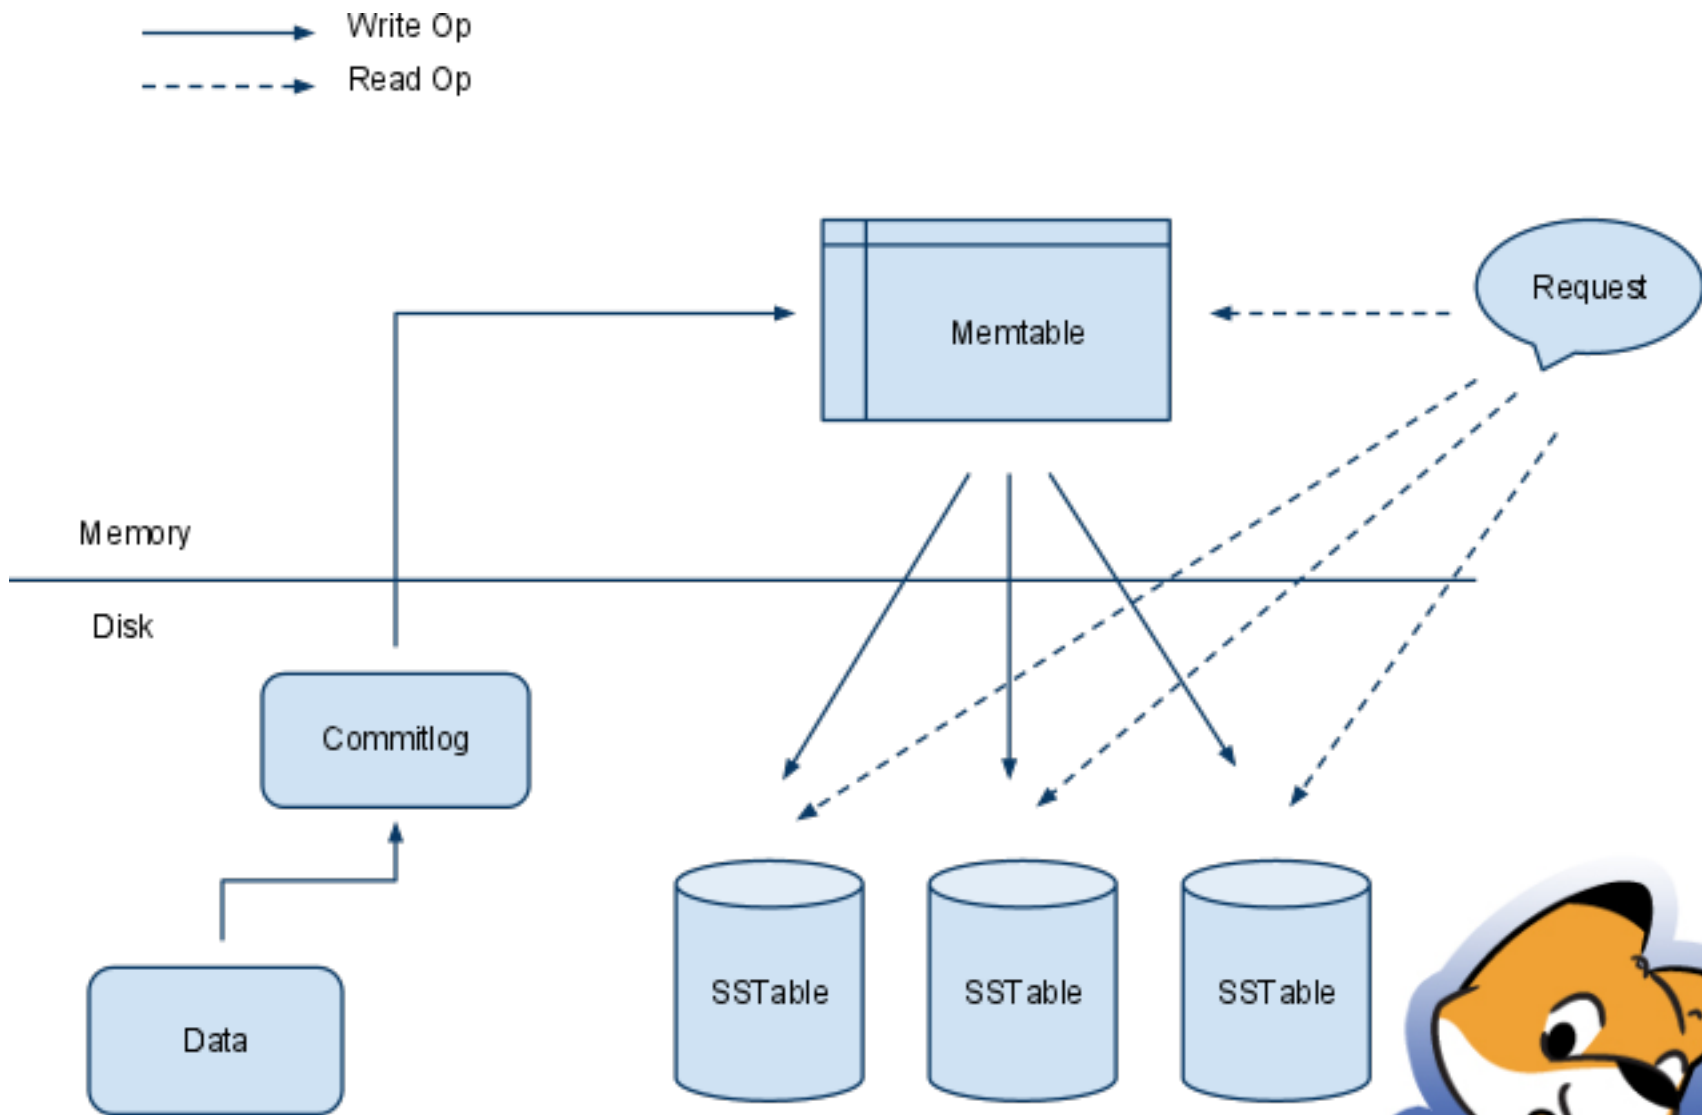
## OrderPreservingPartitioner

Tokens are UTF8 strings





# Read/Write



# Data Model

- Keyspace (like database)
- ColumnFamily (like table)
  - Standard or Super
  - two levels of indexes (key and column name)
- Column and subcolumn sorting
- Specify your own comparator
  - TimeUUID
  - LexicalUUID
  - UTF8
  - Long
  - Bytes

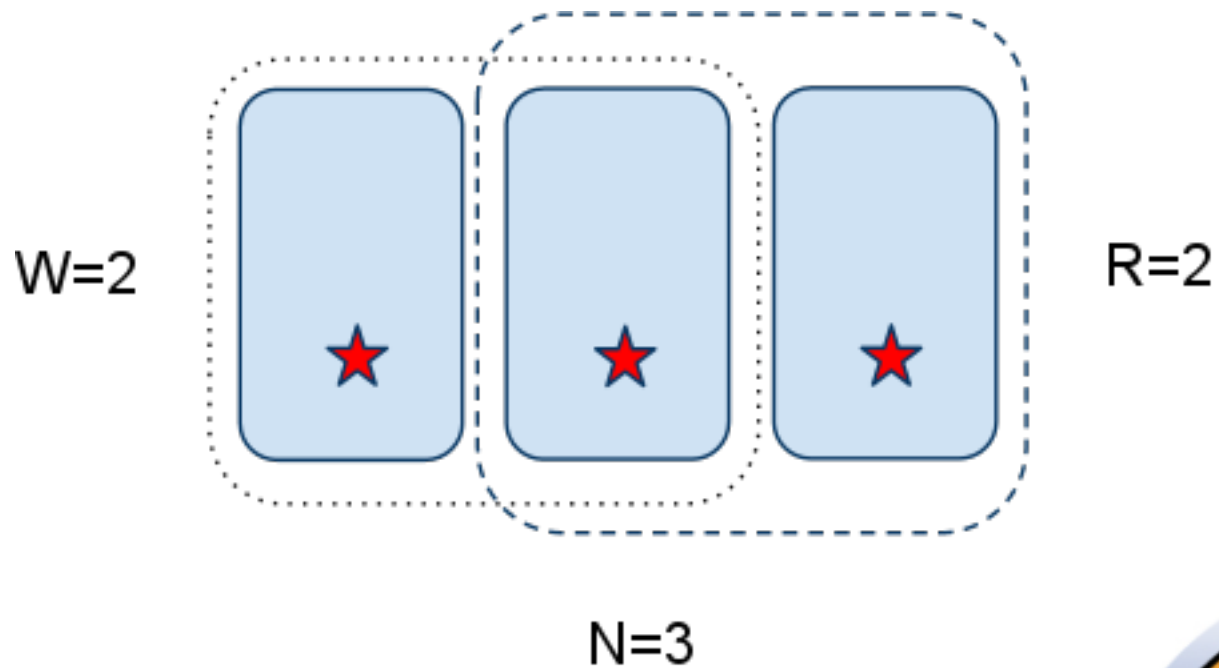


# Consistency

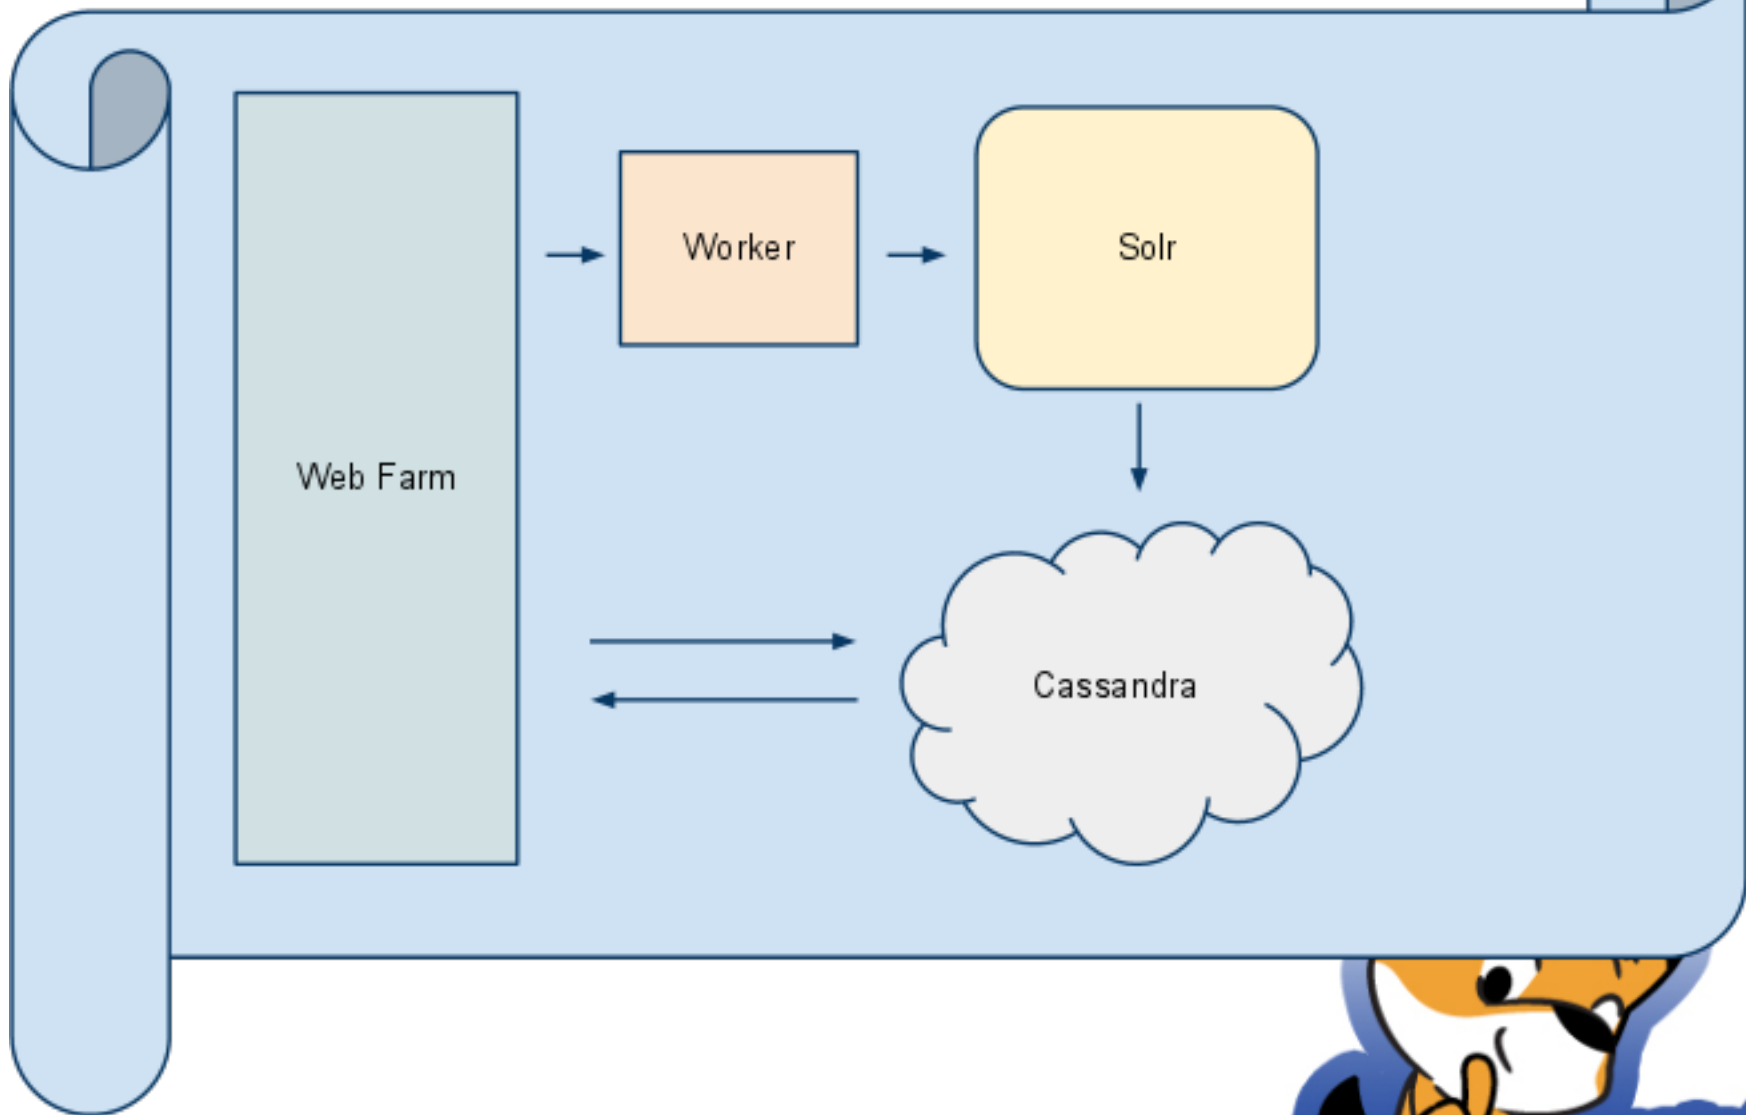
- Write
  - ZERO - asynchronously
  - ANY
  - ONE
  - QUORUM -  $N / 2 + 1$
  - ALL
- Read
  - ONE - first node
  - QUORUM - recent timestamp
- If  $W + R > N$ , you will have consistency
  - $W=1, R=N$
  - $W=N, R=1$
  - $W=Q, R=Q$  where  $Q = N / 2 + 1$



$R+W > N$  guarantees overlap of read and write quorums



# Related Post Architecture



# More Like This

/select?q=id:12345678&mlt=true&mlt.fl=title

```
<result name="match" numFound="1" start="0" maxScore="17.66708">
  <doc>
    <float name="score">17.66708</float>
    <str name="id">47374-24670476</str>
    <str name="title">開源人年會2009</str>
  </doc>
</result>
<result name="response" numFound="20261110" start="0" maxScore="2.1226282">
  <doc>
    <float name="score">2.1226282</float>
    <str name="id">1984330-26558311</str>
    <str name="title">coscup2009</str>
  </doc>
  <doc>
    <float name="score">2.0987473</float>
    <str name="id">551-9899609</str>
    <str name="title">[推廣] 2007 開源人年會</str>
  </doc>
</result>
<arr name="interestingTerms">
  <str>coscup</str>
  <str>源人</str>
  <str>2009</str>
  <str>開源</str>
  <str>人年</str>
  <str>年會</str>
  <str>org</str>
  <str>http</str>
</arr>
</response>
```



# MLT paramaters

```
<field name="title" ... termVectors="true" />
```

- ***mlt.mintf*** - minimum term frequency, default 2
- ***mlt.mindf*** - minimum document frequency, default 5
- ***max.minwl*** - minimum word length, default 0
- ***mlt.maxwl*** - maximum word length, default 0
- ***mlt.maxqt*** - maximum of query terms, default 25
- ***mlt.maxntp*** - maximum number of tokens to parse, default 5000
- ***mlt.boost*** - default false
- ***mlt.count*** - The number of similar documents to return for each result
- ***mlt.interestingTerms*** - one of "list" or "details", this will show what interesting terms are used for query.



# MLT Algorithm

compute all terms frequency.

```
/**
 * Adds terms and frequencies found in vector into the Map termFreqMap
 * @param termFreqMap a Map of terms and their frequencies
 * @param vector List of terms and their frequencies for a doc/field
 */
private void addTermFrequencies(Map termFreqMap, TermFreqVector vector)
{
    String[] terms = vector.getTerms();
    int freqs[] = vector.getTermFrequencies();
    for (int j = 0; j < terms.length; j++) {
        String term = terms[j];

        if (isNoiseWord(term)) {
            continue;
        }
        // increment frequency
        Int cnt = (Int) termFreqMap.get(term);
        if (cnt == null) {
            cnt = new Int();
            termFreqMap.put(term, cnt);
            cnt.x = freqs[j];
        }
        else {
            cnt.x += freqs[j];
        }
    }
}
```





```

/**
 * Create a PriorityQueue from a word->tf map.
 *
 * @param words a map of words keyed on the word(String) with Int objects as the values.
 */
private PriorityQueue createQueue(Map words) throws IOException {
    // have collected all words in doc and their freqs
    int numDocs = ir.numDocs();
    FreqQ res = new FreqQ(words.size()); // will order words by score

    Iterator it = words.keySet().iterator();
    while (it.hasNext()) { // for every word
        String word = (String) it.next();

        int tf = ((Int) words.get(word)).x; // term freq in the source doc
        if (minTermFreq > 0 && tf < minTermFreq) {
            continue; // filter out words that don't occur enough times in the source
        }

        // go through all the fields and find the largest document frequency
        String topField = fieldNames[0];
        int docFreq = 0;
        for (int i = 0; i < fieldNames.length; i++) {
            int freq = ir.docFreq(new Term(fieldNames[i], word));
            topField = (freq > docFreq) ? fieldNames[i] : topField;
            docFreq = (freq > docFreq) ? freq : docFreq;
        }

        if (minDocFreq > 0 && docFreq < minDocFreq) {
            continue; // filter out words that don't occur in enough docs
        }

        if (docFreq == 0) {
            continue; // index update problem?
        }

        float idf = similarity.idf(docFreq, numDocs);
        float score = tf * idf;

        // only really need 1st 3 entries, other ones are for troubleshooting
        res.insert(new Object[]{word, // the word
                                topField, // the top field
                                new Float(score), // overall score
                                new Float(idf), // idf
                                new Integer(docFreq), // freq in all docs
                                new Integer(tf)}));
    }
    return res;
}

```

sort  
by  
tf\*idf



```

/**
 * Create the More like query from a PriorityQueue
 */
private Query createQuery(PriorityQueue q) {
    BooleanQuery query = new BooleanQuery();
    Object cur;
    int qterms = 0;
    float bestScore = 0;

    while (((cur = q.pop()) != null)) {
        Object[] ar = (Object[]) cur;
        TermQuery tq = new TermQuery(new Term((String) ar[1], (String) ar[0]));

        if (boost) {
            if (qterms == 0) {
                bestScore = ((Float) ar[2]).floatValue();
            }
            float myScore = ((Float) ar[2]).floatValue();

            tq.setBoost(boostFactor * myScore / bestScore);
        }

        try {
            query.add(tq, BooleanClause.Occur.SHOULD);
        }
        catch (BooleanQuery.TooManyClauses ignore) {
            break;
        }

        qterms++;
        if (maxQueryTerms > 0 && qterms >= maxQueryTerms) {
            break;
        }
    }

    return query;
}

```

BooleanClause.  
Occur

1. MUST
2. MUST\_NOT
3. SHOULD



# Conclusion

- don't just think
- log everything

```
INFO: [] webapp=/blogarticle path=/relate params={id=2250592-7594244&mlt.fl=body&mlt.  
debug=true&mlt.maxqt=5&type=site&wt=json&fq=status:2&fq=spam:false&fq=enable:  
true&rows=20} cassandra=3 ms. terms={coscup 開源 人年 舞會 2010 } status=0  
QTime=149
```

- use \*Factory

```
<analyzer type="index" class="org.apache.lucene.analysis.cjk.CJKAnalyzer"  
>  
  <tokenizer class="org.apache.lucene.analysis.cjk.CJKTokenizer" />  
  <filter class="solr.LowerCaseFilterFactory"/>  
  ...more  
</analyzer>
```

- HTML kill you.



# cassandra-munin-plugin

<http://github.com/jamesgolick/cassandra-munin-plugins>

