

SPM: Large Scale Performance Monitoring for ElasticSearch HBase Solr & Friends

#spmbuzz

#bbuzz


Otis Gospodnetić – Sematext International
[@otisg](#) ◦ [@sematext](#) ◦ [sematext.com](#)

Agenda

- Introductions
- SPM Architecture
- SPM Client Side
- SPM Receiver & Flume
- SPM & HBase
- SPM Front End
- Lessons
- Interesting Metrics

About Otis Gospodnetić

- **ASF Member:** Lucene, Solr, Nutch, Mahout
- **Author:** Lucene in Action 1 & 2
- **Entrepreneur:** Sematext, Simpy



About Sematext

Products & Services

Consulting, Development, Tech Support:

- **Search**; Lucene, Solr, ElasticSearch...
- **Big Data Analytics**; Hadoop, HBase, Voldemort...
- **Web Crawling**; Nutch, Droids
- **Machine Learning**; Mahout

Sematext Metrics

- 100% **organic**: keine GMO / VC
- 5 years old
- < 10 people
- 7 countries
- 3 timezones
- 3 continents
- > 100 customers



Instagram
Fast beautiful photo sharing

tumblr.



THOMSON REUTERS

Salesforce

Etsy



LIBRARY OF
CONGRESS

Scalable Performance Monitoring

- Multi-tenant APM SaaS
- **With:** HBase, Hadoop, Flume, Jetty, Esper...
- **For:** ElasticSearch, HBase, Solr, Sensei...
- Amazon Cloud

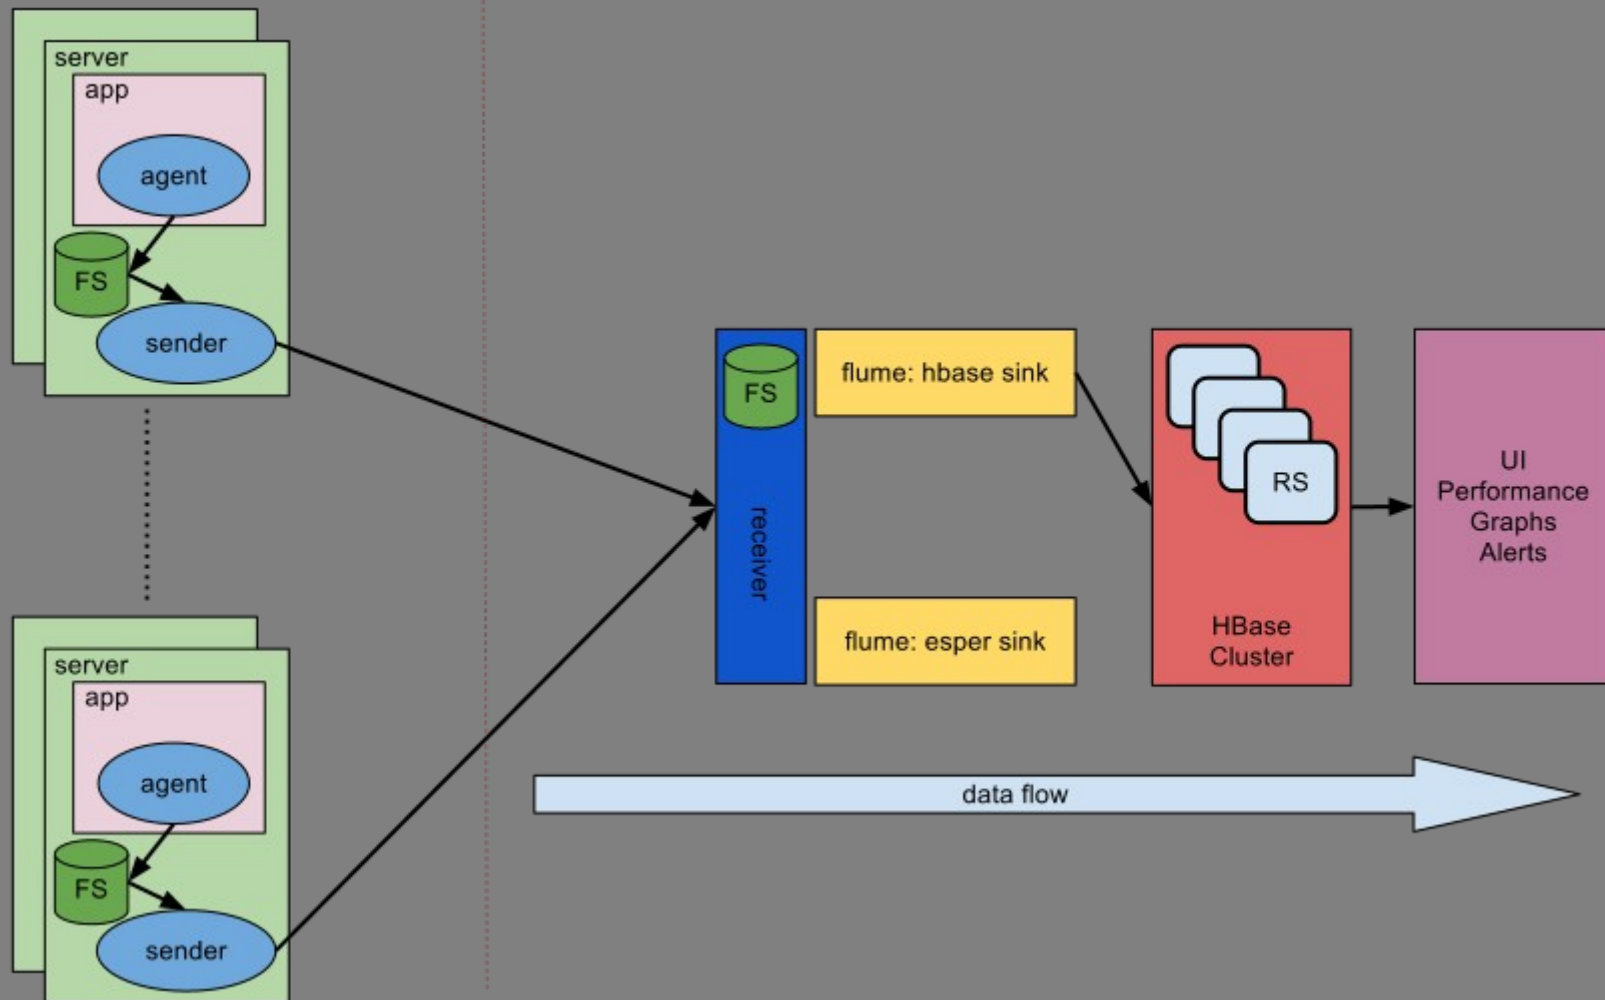
apps.semamtext.com

ganz frei

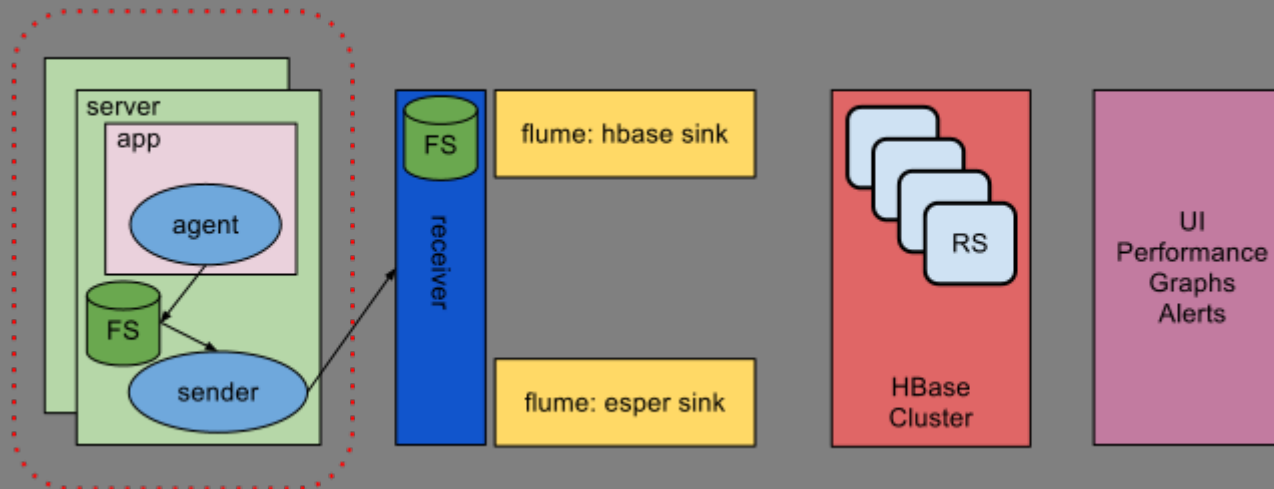
SPM Functionality

- Performance graphs for **any time period**
- No loss of **metric precision/granularity**
- **Slice & dice metrics** with filters
- **Compare** two time periods
- Integrated **alerts**
- Integrated **email reports**
- Monitor **multiple** (types of) **systems**
- Metric **arithmetic** (ganz soon!)

Bird's Eye View



SPM: Client Side

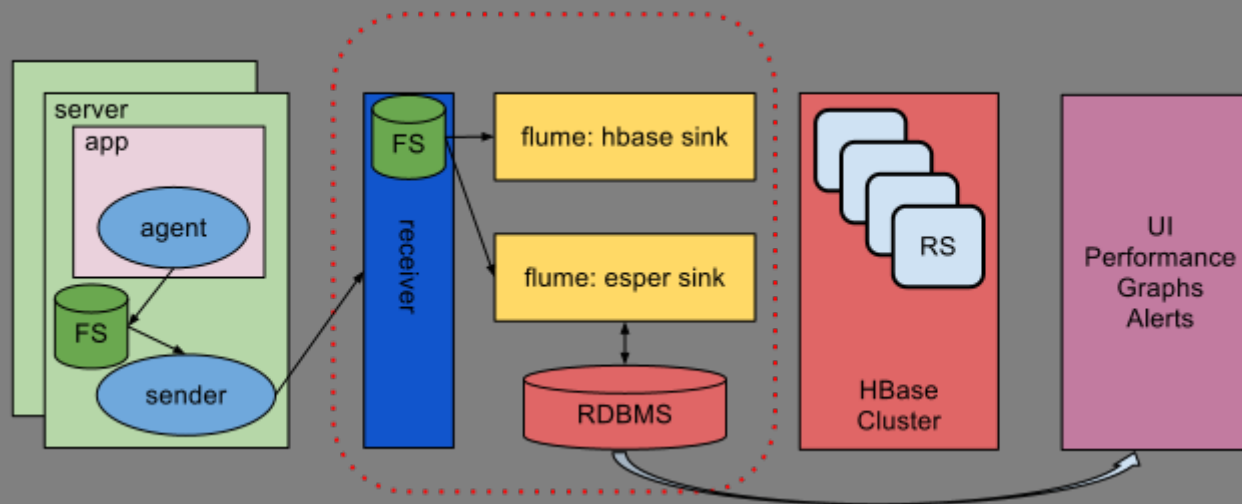


- In-**JVM Agent** – app stats collector
- **Collectd** – system stats collector
- Out-of-JVM **Sender** – collector & sender
- JMX, bytecode instrumentation, Flume, HttpClient, local buffering

SPM: Client Side

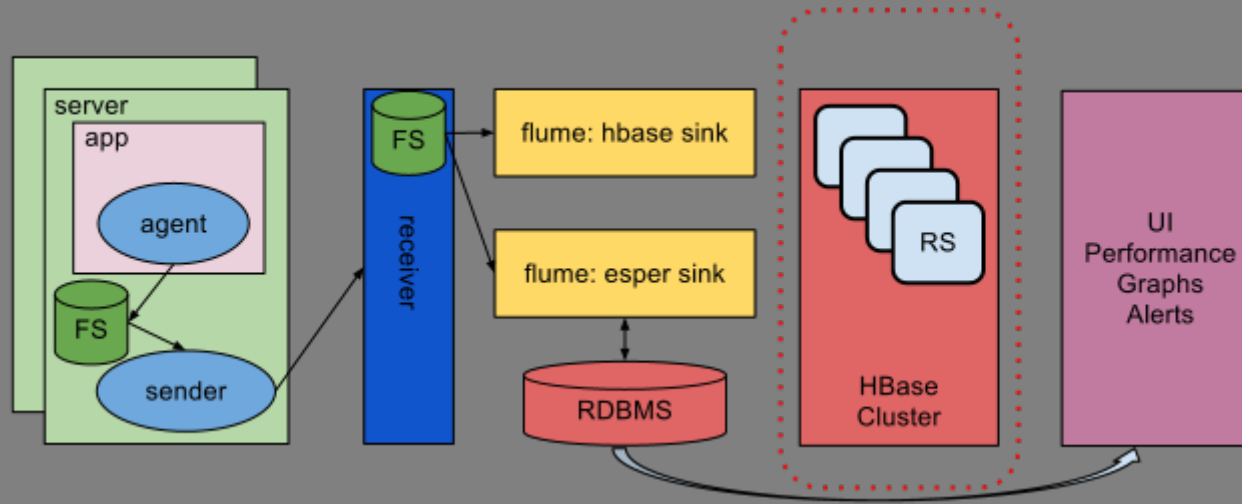
- apt-get install **collectd**
- run SPM client **installer**
- `-Dcom.sun.management.jmxremote`
`-javaagent:/spm/spm-monitor/lib/spm-monitor-es-1.5.0-withdeps.jar=/spm/spm-monitor/conf/spm-monitor-config-token-here-default.xml`
- `/etc/init.d/spm-sender start`

SPM: Receiver & Flume



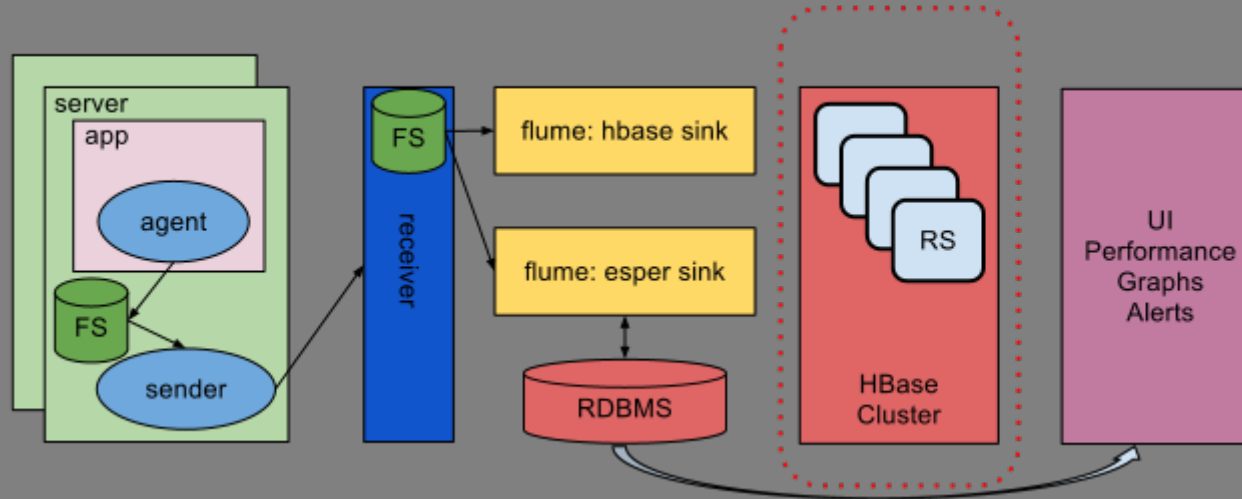
- Apache + Jetty + KeepAlive
- Flume tail + HBase sink
- In-memory metric pre-aggregation - HBaseHUT
- Flume tail + Esper sink = Performance Alerts

SPM: HBase



- Raw metrics table
- Aggregated metrics table
- Aggregations via MapReduce
- Aggregations on the fly - HBaseHUT

SPM: HBase Aggregations



- Generic aggregations framework
- Large dataset report building (HDFS/HBase)

SPM: Front End


- JavaScript UI, jQuery, Flot
- Flot → HighCharts → d3
- Filters:
 - Time, Host, JVM
 - Core, Request Handler (Solr)
 - Node, Index, and Shard (ElasticSearch)
 - Network interface
 - Disk
 - ...

ElasticSearch Indexing Perf.

Indexing - 2012.06.05 03:00 to 2012.06.05 09:00

hide extended chart stacked

indexed docs deleted docs indexing rate deleting rate



SPM Numbers

- >1000 data points/second
- 6 HBase nodes
- BB raw rows & MM aggregate rows

Lessons

- EC2
- Linux
- KeepAlives
- HBase

Lessons: EC2



Otis Gospodnetić @otisg

16 May

EC2 is the best app tuning environment ever made by the man kind. If your app is semi-fast on EC2, it will fly everywhere else!

[Collapse](#) [Reply](#) [Delete](#) [Favorite](#)


38
RETWEETS

6
FAVORITES



10:53 AM - 16 May 12 via HootSuite · [Details](#)

Agreement?



- Benchmarking impossible
- Spot instance prices spike
- Are you feeling lucky?

Lessons: swappiness

```
>sudo sysctl -A | grep vm.swappiness  
vm.swappiness = 60
```

```
>sudo sysctl -w vm.swappiness = 0  
vm.swappiness = 0
```

```
>sudo sysctl -A | grep vm.swappiness  
vm.swappiness = 0
```

```
>free -m
```

	total	used	free	shared
Mem:	1717	1669	48	0
Swap:	895	0	895	

```
>sudo /sbin/swapoff -a && sudo /sbin/swapon -a
```

Lessons: nofile

> ulimit -n

1024

Default – often way too low



> sudo vim /etc/security/limits.conf

spmsender - nofile 32000

Bump it up!



> sudo vim /etc/pam.d/su

session required pam_limits.so

Lessons: KeepAlives

```
#  
# KeepAlive: Whether or not to allow persistent connections (more than  
# one request per connection). Set to "Off" to deactivate.  
#  
KeepAlive On  
  
#  
# MaxKeepAliveRequests: The maximum number of requests to allow  
# during a persistent connection. Set to 0 to allow an unlimited amount.  
# We recommend you leave this number high, for maximum performance.  
#  
MaxKeepAliveRequests 100  
  
#  
# KeepAliveTimeout: Number of seconds to wait for the next request from  
# same client on the same connection.  
#  
KeepAliveTimeout 70
```

- KeepAlives matter at high request rates

Lessons: HBase Keys

- Bad keys can hotspot RegionServers
- Solution: **HBaseWD**
<https://github.com/sematext/HBaseWD>

Lessons: Append-only Updates

- Read-Modify-Write can be slow
- Solution: **HBaseHUT**
<http://github.com/sematext/HBaseHUT>

Lessons: HTable vs. asynchbase

- HTable is blocking and serial
- Solution:
<https://github.com/stumbleupon/asynchbase>
- See:
<https://issues.apache.org/jira/browse/HBASE-5539>
<http://tsunanet.net/~tsuna/asynchbase/benchmark/viz.html>

Poll: Search Monitoring

What do you use for Solr performance monitoring?

Ganglia 2.33% (1 votes)



Munin 16.28% (7 votes)



Zabbix 9.3% (4 votes)



Statsd & Graphite 4.65% (2 votes)



Home-grown/Other (pls share in comments) 25.58% (11 votes)



NewRelic 16.28% (7 votes)



SPM for Solr 4.65% (2 votes)



Nothing 20.93% (9 votes)



Total Votes: 43

What do you use for ElasticSearch performance monitoring?

Ganglia or Munin or Zabbix or Cacti 15.79% (6 votes)



BigDesk 31.58% (12 votes)



SPM for ElasticSearch 2.63% (1 votes)



Home-grown/Other (pls share in comments) 18.42% (7 votes)



Nothing 31.58% (12 votes)



Total Votes: 38

Interesting Metrics

- Solr
- ElasticSearch
- HBase

Buzzworthiness

- Big Data
- Search
- Analytics
- Real-time
- Large Scale
- Multi-tenant
- SaaS
- Cloud

We're Hiring

Dig **Search**?

Dig **Analytics**?

Dig **Big Data**?

Dig **Performance**?

Dig working with and in **open-source**?

We're hiring world-wide!

<http://sematext.com/about/jobs.html>

Contact

sematext.com

blog.sematext.com

github.com/sematext

[@sematext](#)

[@otisg](#)

otis@sematext.com

Stop by our booth