



OpenSOC

The Open Security Operations Center

for

Analyzing 1.2 Million Network Packets per Second in
Real Time

James Sirota,
Big Data Architect
Cisco Security Solutions Practice
jsirota@cisco.com

Sheetal Dolas
Principal Architect
Hortonworks
sheetal@hortonworks.com

June 3, 2014

Over Next Few Minutes

- Problem Statement & Business Case for OpenSOC
- Solution Architecture and Design
- Best Practices and Lessons Learned
- Q & A

Business Case

“There's now a growing sense of fatalism:
It's no longer if or when you get hacked,
but the assumption is that you've already
been hacked,
with a focus on minimizing the damage.”

Source: Dark Reading / Security's New
Reality: Assume The Worst

Breaches Happen in Hours...

But Go Undetected for Months or Even Years

In 60% of breaches, data is stolen in hours

Initial Attack to Initial Compromise



10%



75%



12%



2%



0%



1%



1%

Initial Compromise to Data Exfiltration



8%



38%



14%



25%



8%



8%



0%

Initial Compromise to Discovery



0%



0%



2%



13%



29%



54%



2%

Discovery to Containment/ Restoration



0%



1%



9%



32%



38%



17%



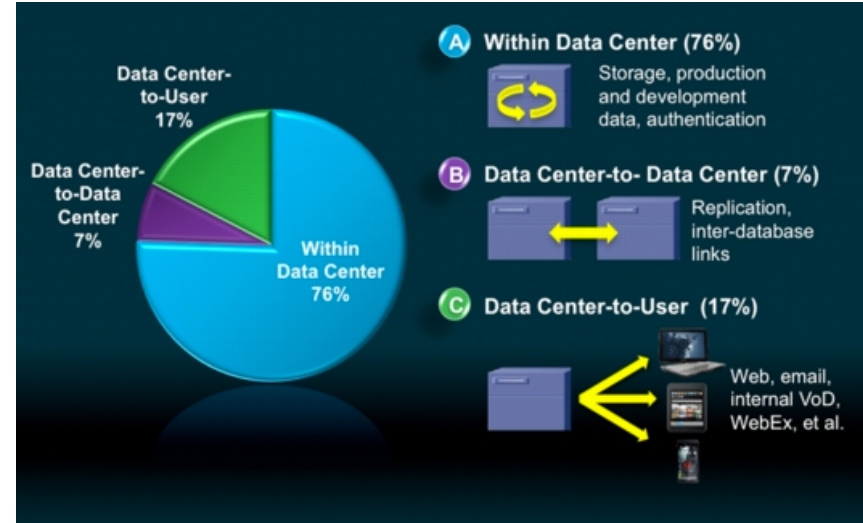

4%

54% of breaches are not discovered for months

Timespan of events by percent of breaches

Source: 2013 Data Breach Investigations Report

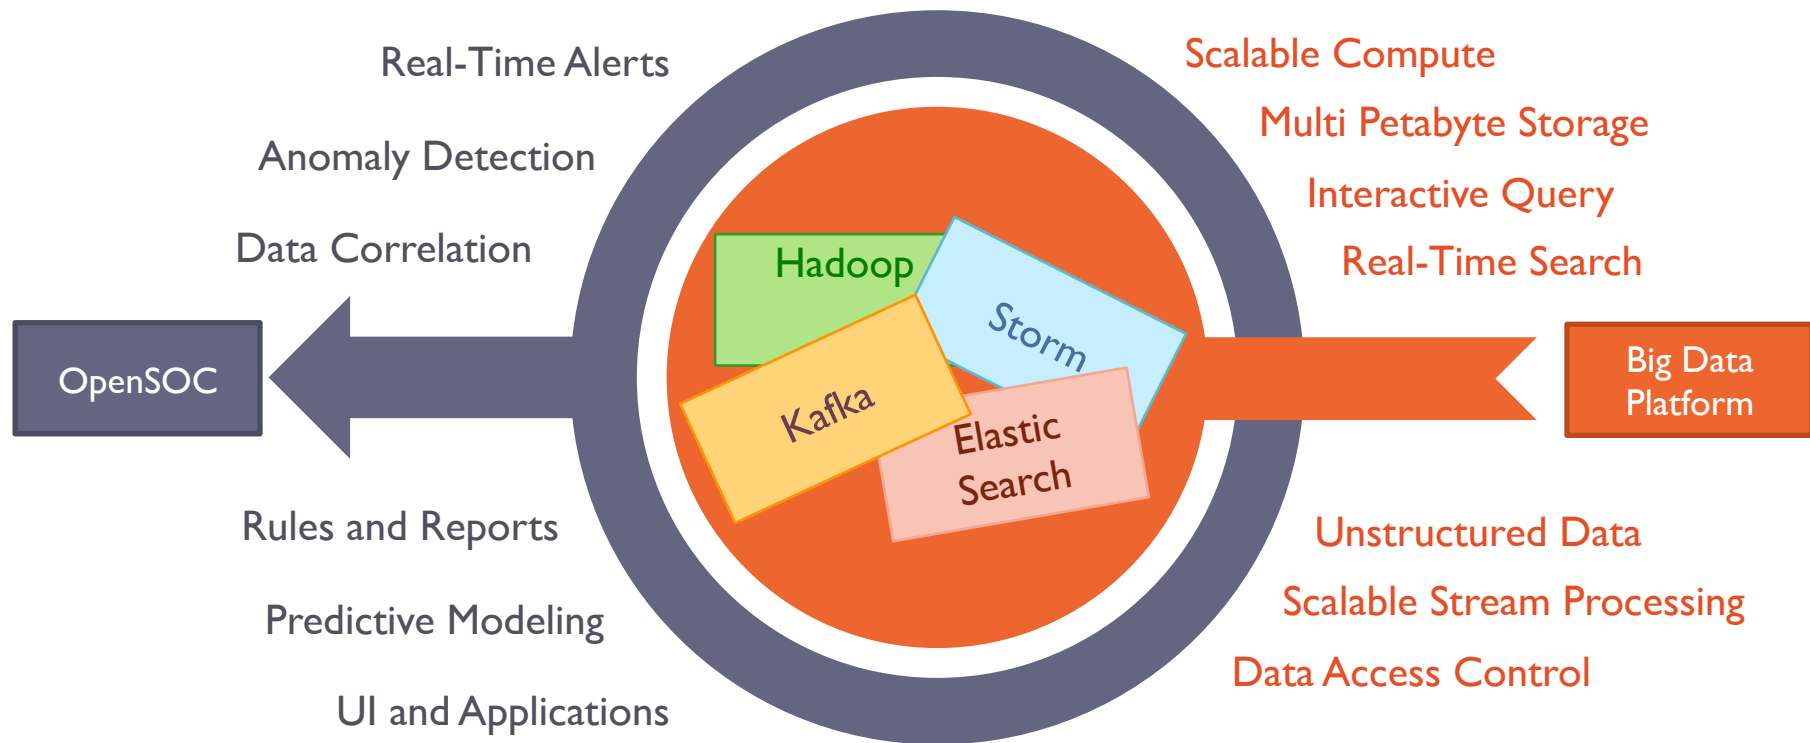
Cisco Global Cloud Index



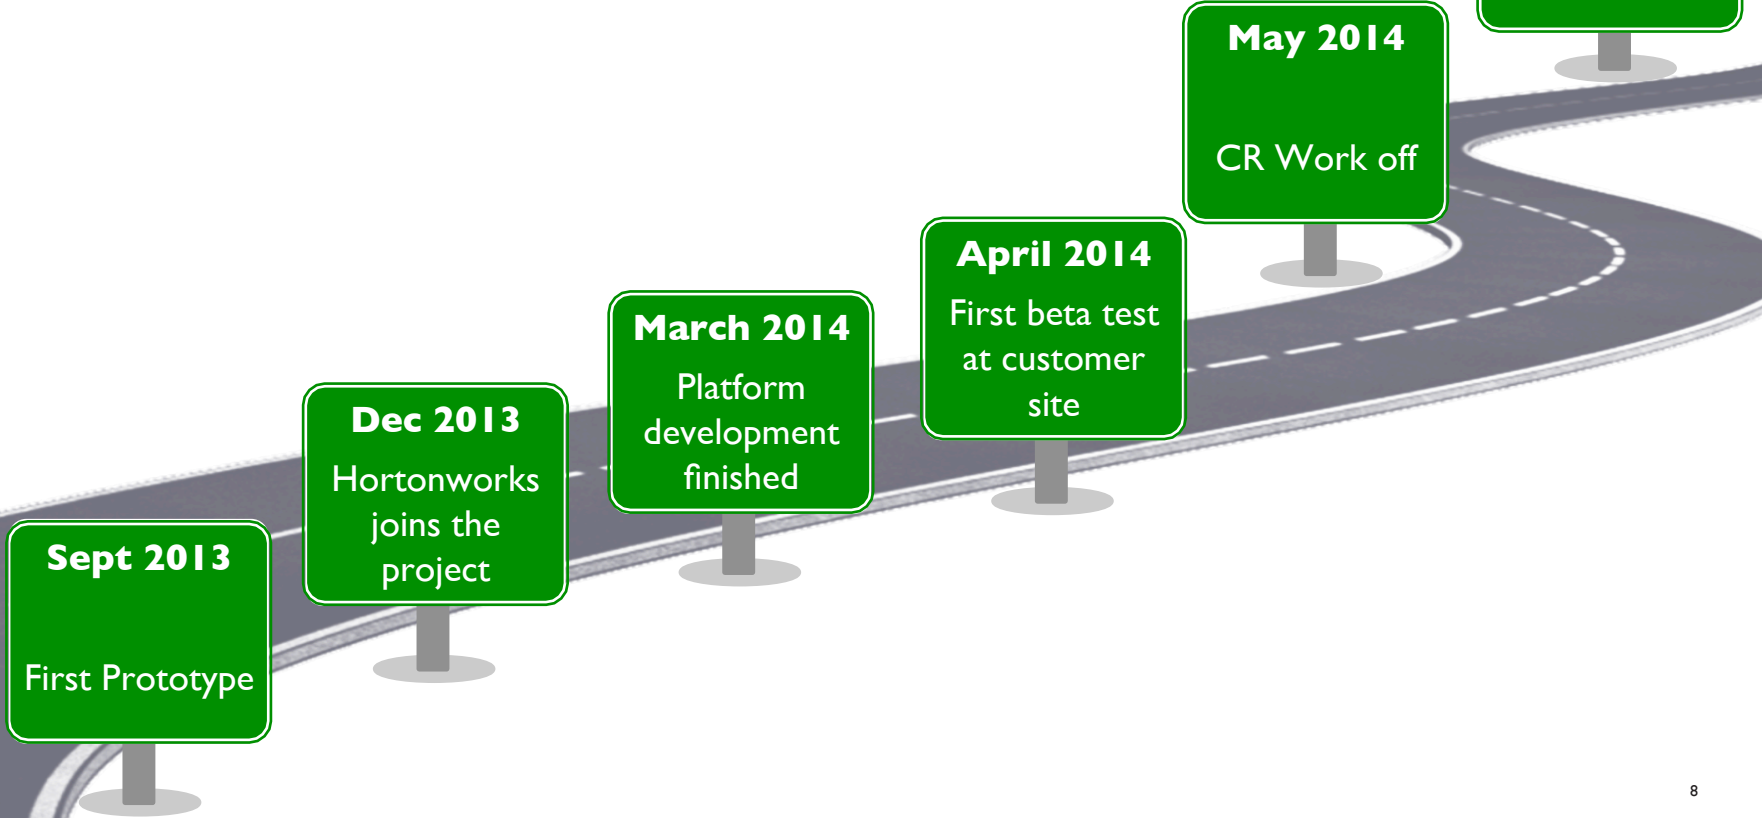
Source: 2014 Cisco Global Cloud Index

Introducing OpenSOC

Intersection of Big Data and Security Analytics

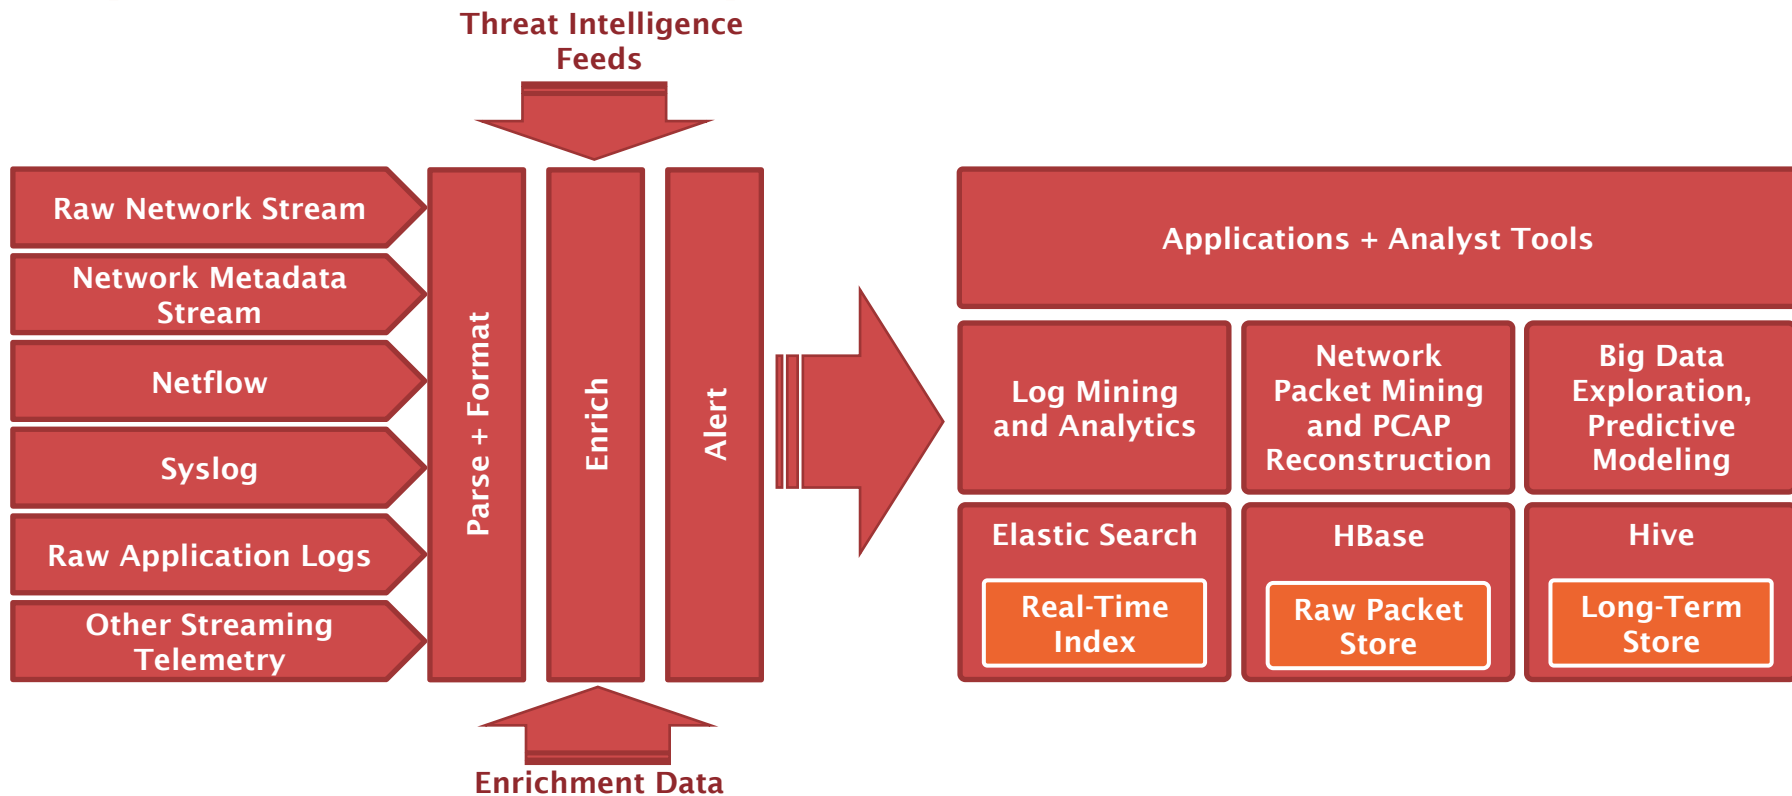


OpenSOC Journey



Solution Architecture & Design

OpenSOC Conceptual Architecture



Key Functional Capabilities

- Raw Network Packet Capture, Store, Traffic Reconstruction
- Telemetry Ingest, Enrichment and Real-Time Rules-Based Alerts
- Real-Time Telemetry Search and Cross-Telemetry Matching
- Automated Reports, Anomaly Detection and Anomaly Alerts
- Rich Analytics Apps and Integration with Existing Analytics Tools

The OpenSOC Advantage

- Fully-Backed by Cisco and Used Internally for Multiple Customers
- Free, Open Source and Apache Licensed
- Built on Highly-Scalable and Proven Platforms (Hadoop, Kafka, Storm)
- Extensible and Pluggable Design
- Flexible Deployment Model (On-Premise or Cloud)
- Centralize your processes, people and data

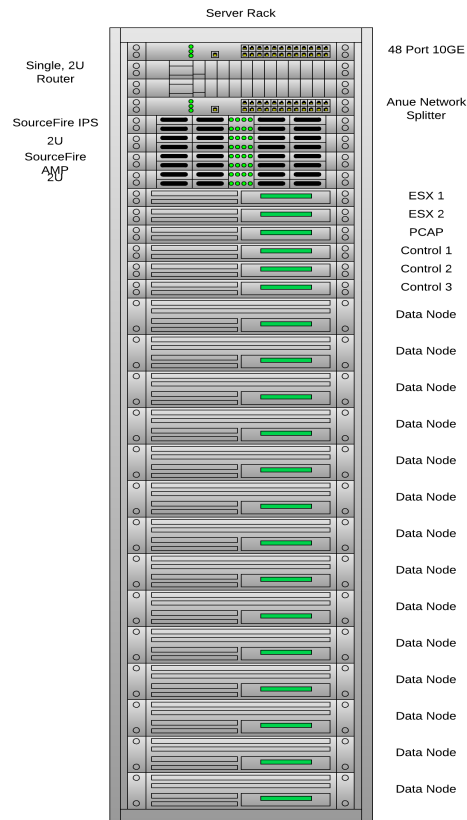
OpenSOC Deployment at Cisco

Hardware footprint (40u)

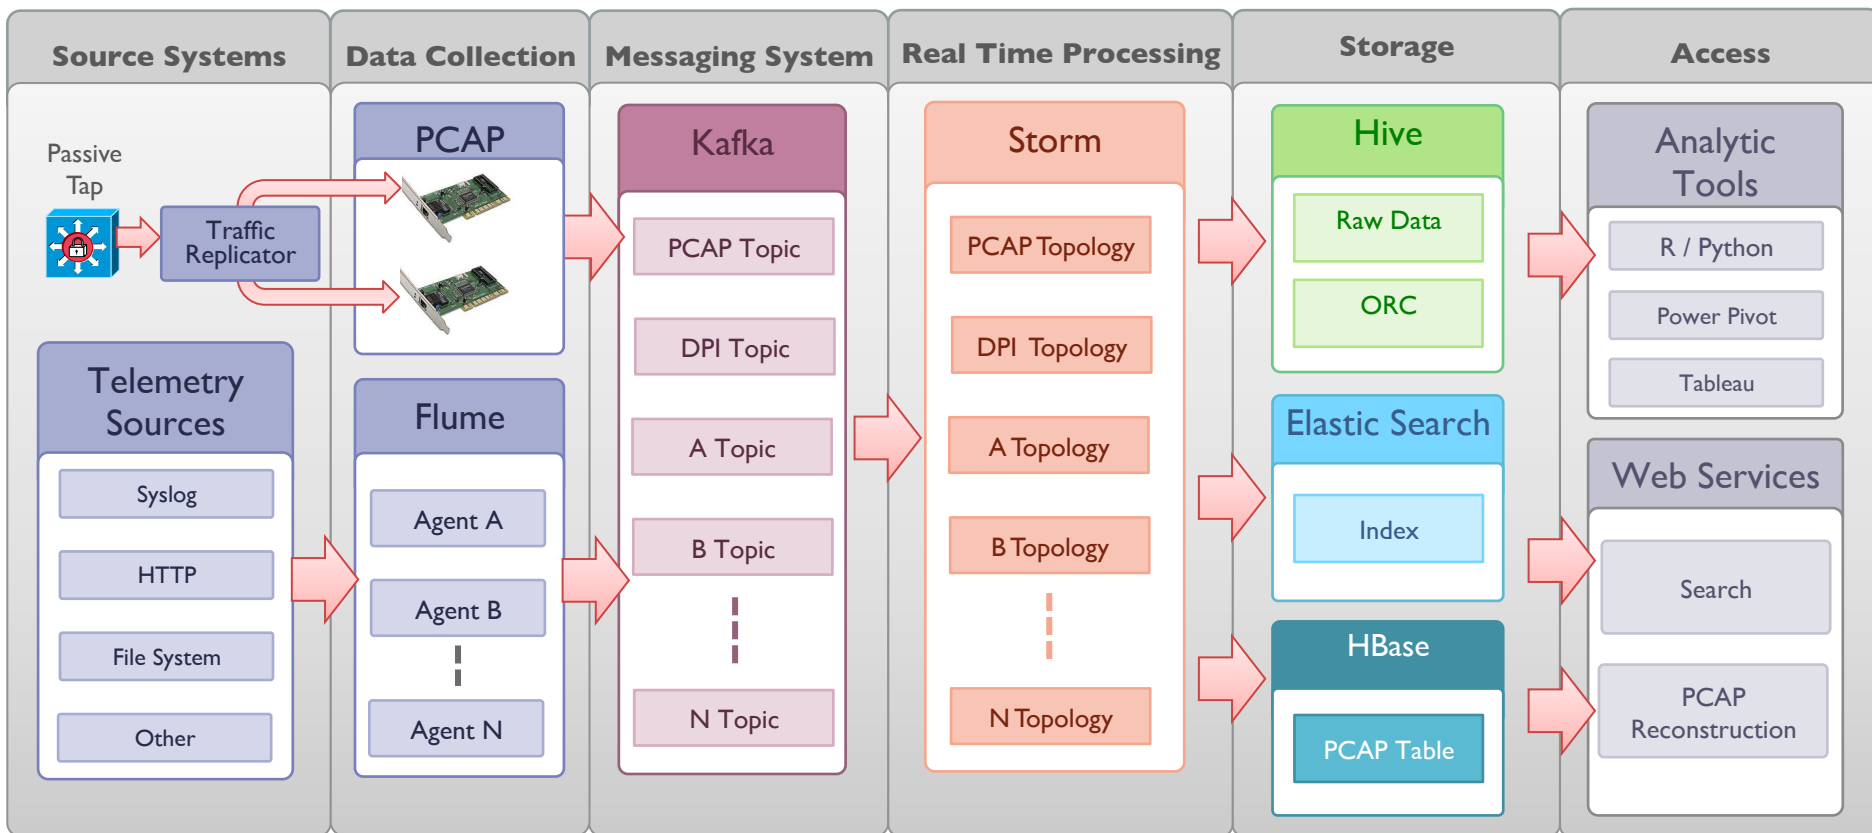
- **14 Data Nodes** (UCS C240 M3)
- 3 Cluster Control Nodes (UCS C220 M3)
- 2 ESX Hypervisor Hosts (UCS C220 M3)
- **1 PCAP Processor** (UCS C220 M3 + Napatech NIC)
- 2 SourceFire Threat alert processors
- 1 Anue Network Traffic splitter
- 1 Router
- 1 48 Port **10GE Switch**

Software Stack

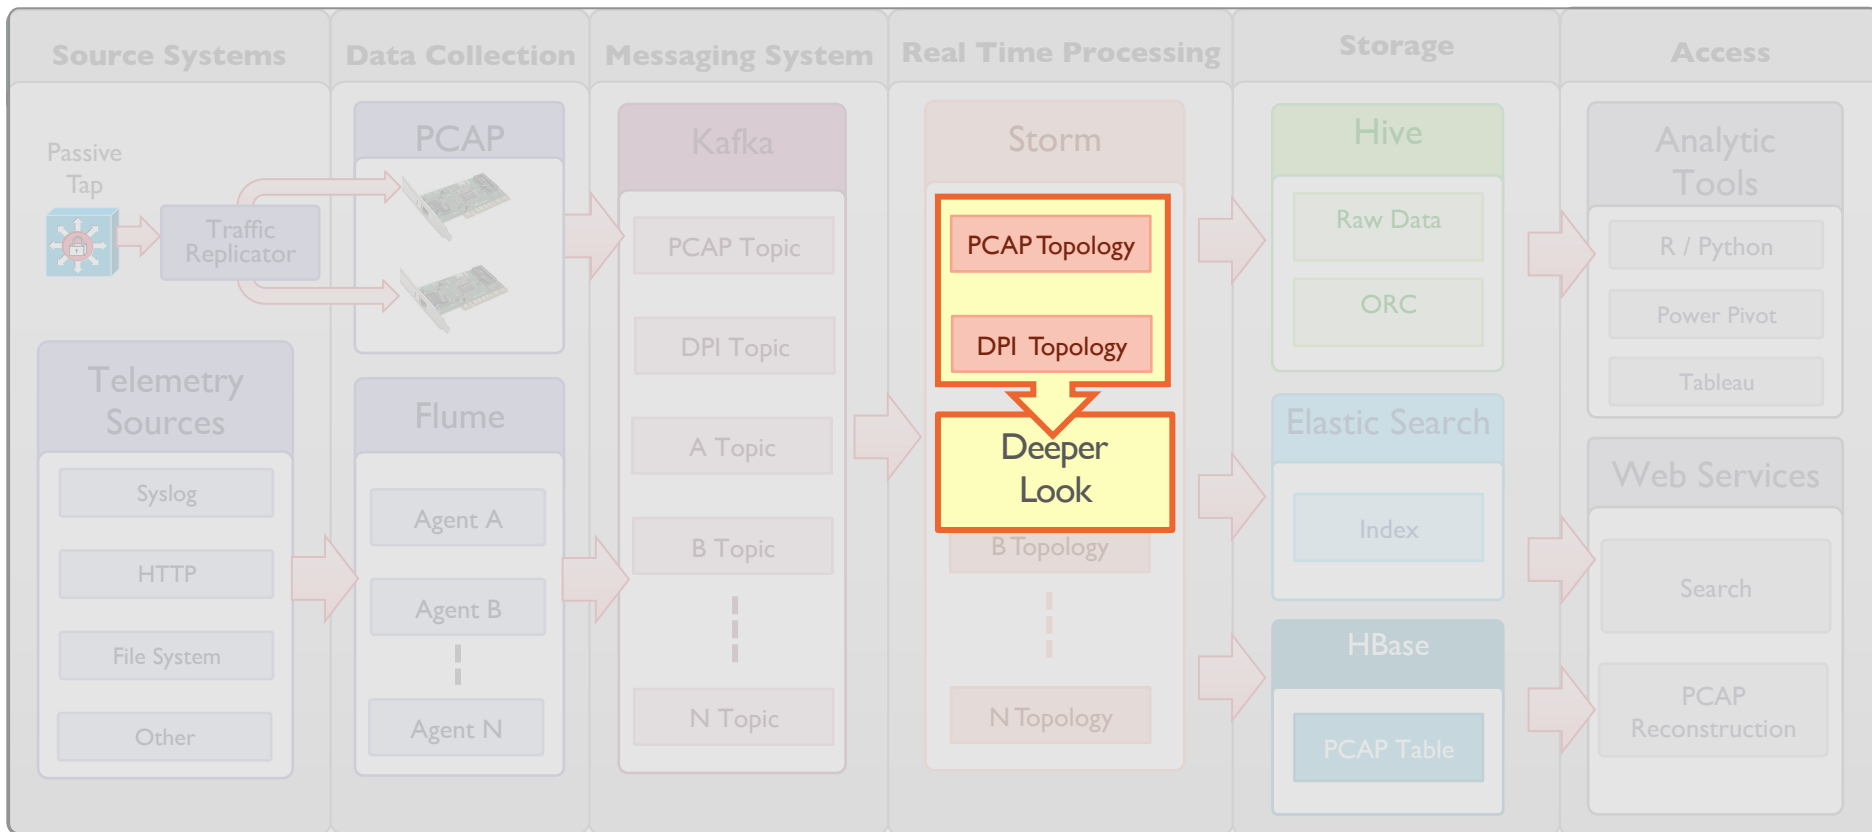
- HDP 2.1
- Kafka 0.8
- Elastic Search 1.1
- MySQL 5.5



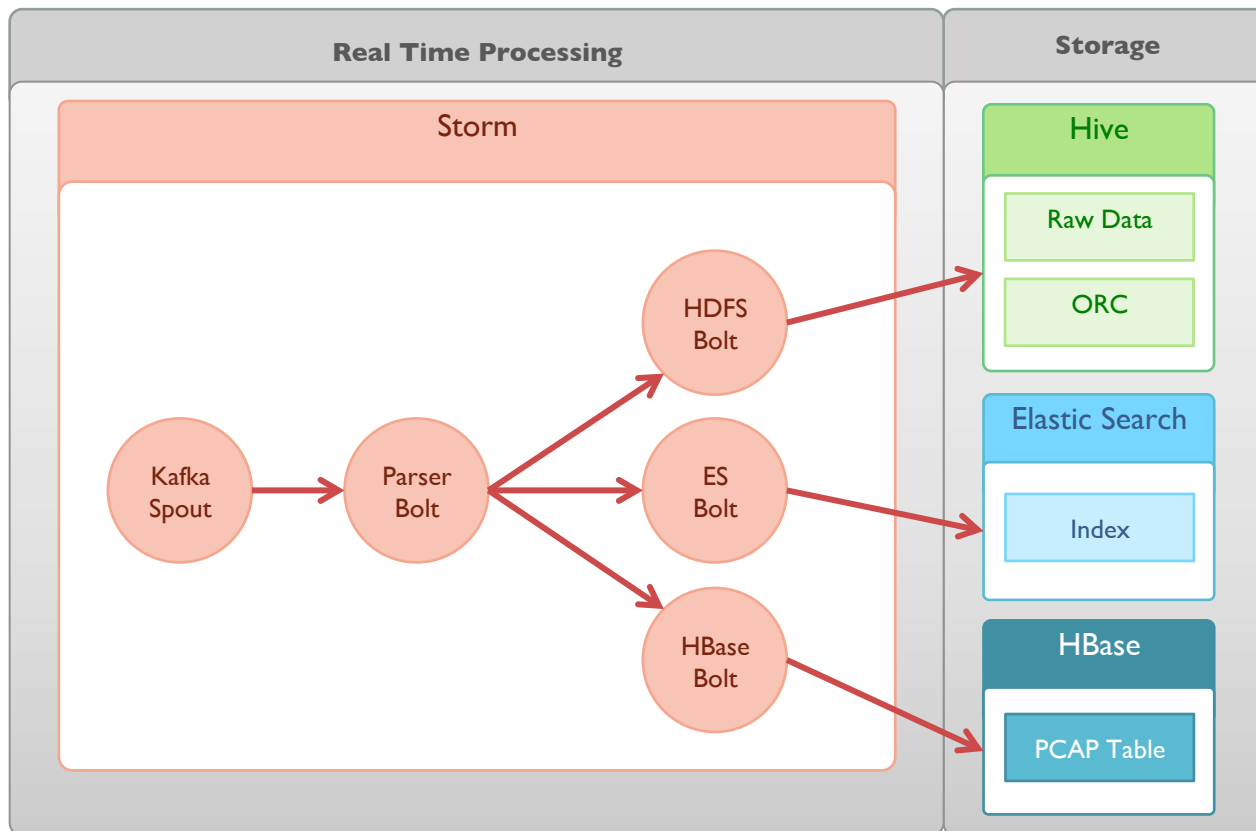
OpenSOC - Stitching Things Together



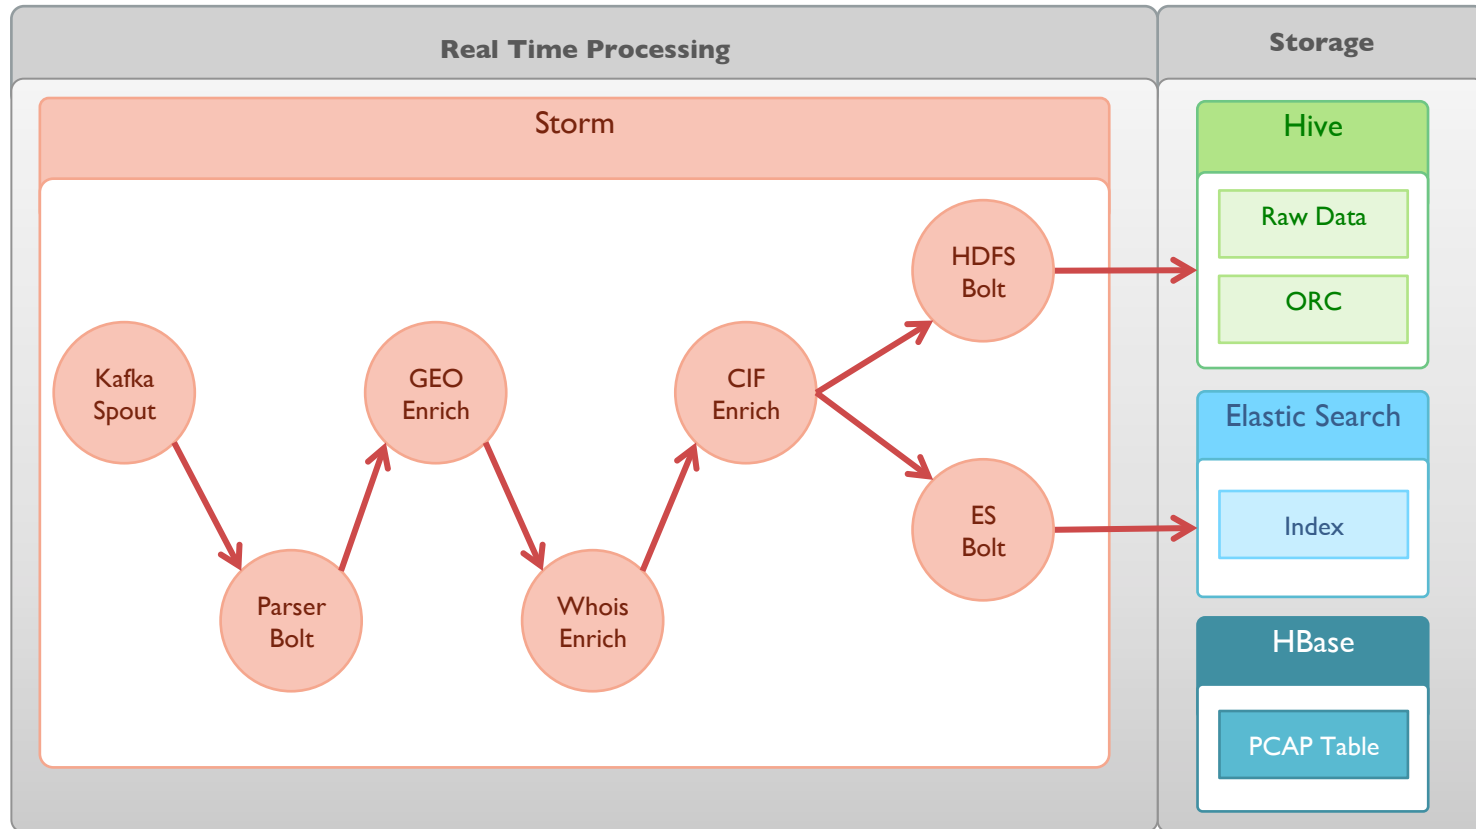
OpenSOC - Stitching Things Together



PCAP Topology



DPI Topology & Telemetry Enrichment

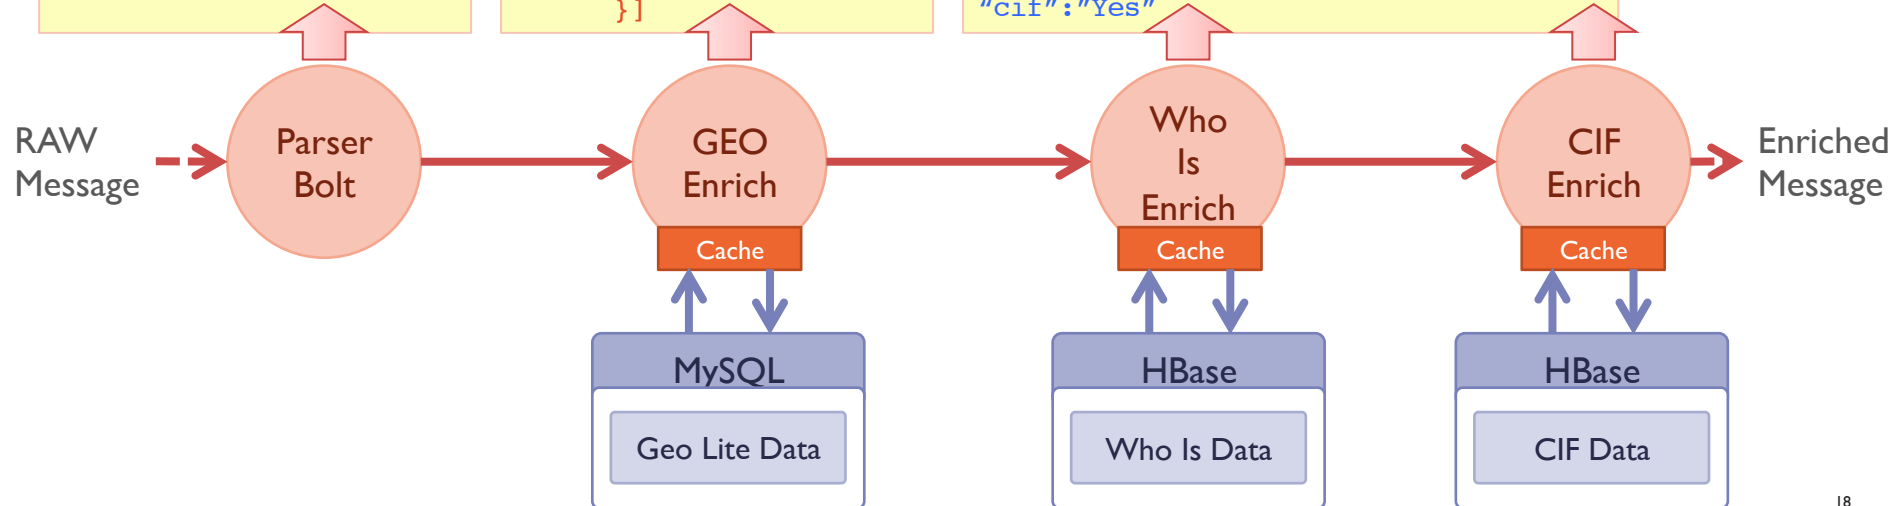


Enrichments

```
{  
  "msg_key1": "msg value1",  
  "src_ip": "10.20.30.40",  
  "dest_ip": "20.30.40.50",  
  "domain": "mydomain.com"  
}
```

```
"geo":[ { "region": "CA",  
  "postalCode": "95134",  
  "areaCode": "408",  
  "metroCode": "807",  
  "longitude": -121.946,  
  "latitude": 37.425,  
  "locId": 4522,  
  "city": "San Jose",  
  "country": "US"  
} ]
```

```
"whois":[ {  
  "OrgId": "CISCOS",  
  "Parent": "NET-144-0-0-0-0",  
  "OrgAbuseName": "Cisco Systems Inc",  
  "RegDate": "1991-01-17",  
  "OrgName": "Cisco Systems",  
  "Address": "170 West Tasman Drive",  
  "NetType": "Direct Assignment"  
} ],  
"cif": "Yes"
```



Applications: Telemetry Matching and DPI

The screenshot displays the Cisco Security Analytics interface with several key components and annotations:

- Navigation Bar:** Includes Cisco logo, DemoUser, Settings, Help, Knowledge Base, Contact Us, and Logout.
- Search Criteria:** A form for searching alerts, with "SFIMS Alert" entered and a time range from 5/14/2014 11:45 to 5/14/2014 12:00. An orange box labeled "Step 1: Search" points to this section.
- Session View:** A map showing the location of the session, with a blue pin indicating the location. An orange box labeled "Step 2: Match" points to the "Session Id" field in the "Message" section.
- Message Section:** Displays a list of messages. The "Session Id" field is highlighted with an orange box labeled "Step 3: Analyze".
- PCAP File:** A section on the right showing the "PCAP File" as "sample.pcap". An orange box labeled "Step 4: Build PCAP" points to this section.

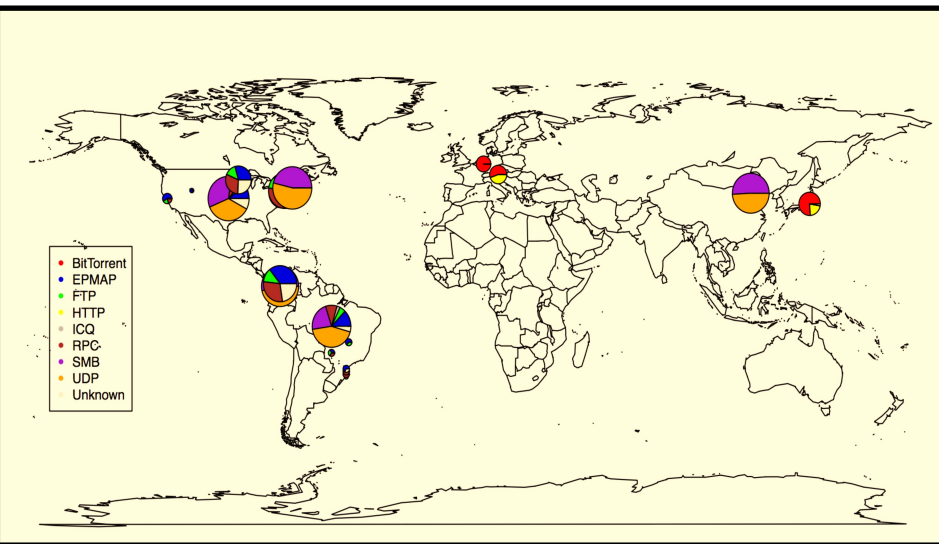
Session View Details:

- Session Id : 9495989310
- Geo Latitude : 40.5525
- Source : [Redacted]
- Message : SFIMS: Primary Base Malware P... User: Unknown... UnknownInterfa... Passive... Secur...
- Tags : [Redacted]
- PCAP File : sample.pcap

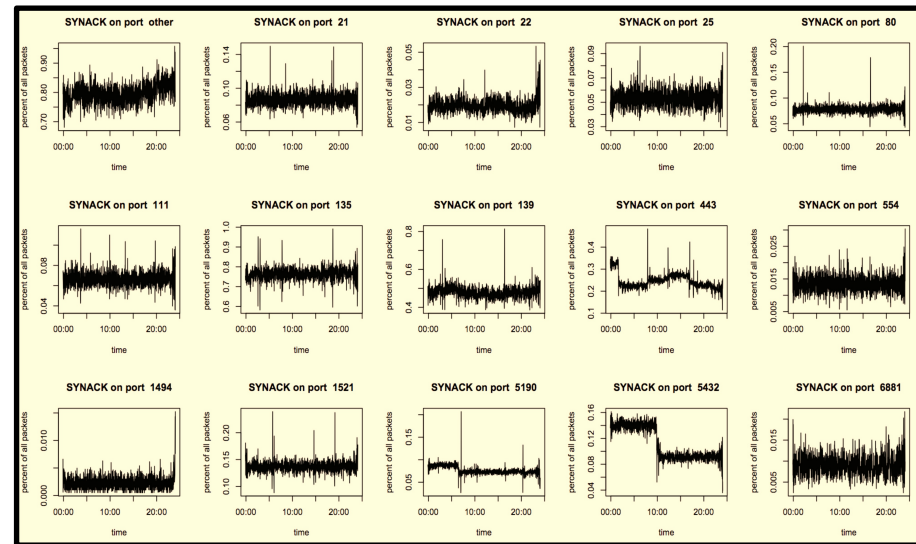
Table:

Session Id	Source IP	Destination IP	Protocol
9495989310	[Redacted]	[Redacted]	6
9495989310	[Redacted]	[Redacted]	6
9495989310	[Redacted]	[Redacted]	6
9495989310	[Redacted]	[Redacted]	6
9495989310	[Redacted]	[Redacted]	6

Integration with Analytics Tools



Dashboards




Reports

Best Practices and Lessons Learned

Journey Towards Highly Scalable Application

Kafka Tuning

This is where we began



Some code optimizations and increased parallelism

February 7, 2014

Sheetal Dolas

250K ps now with more workers

9:27 AM



Ron Lee

congrats! you figured it out?

12:06 PM

Sheetal Dolas

yes a little bit. had to dig into code and change some implementations

12:07 PM

but still long way to go

12:07 PM

we want to be millionaires

▼ 12:08 PM

Kafka Tuning

- Is Disk I/O heavy
- Kafka 0.8+ supports replication and JBOD
 - Better performance compared to RAID
- Parallelism is largely driven by number of disks and partitions per topic
- Key configuration parameters:
 - `num.io.threads` - Keep it at least equal to number of disks provided to Kafka
 - `num.network.threads` - adjust it based on number of concurrent producers, consumers and replication factor

After Kafka Tuning

February 7, 2014

Sheetal Dolas

700K ps on kafka now

4:14 PM



Ron Lee

cpus peaked out at 50%, network peaked at 1G, plenty of ram, need to hit the servers harder, I will add one more node tonight

▼ 10:13 PM

Bottleneck Isolation, Resource Profiling, Load Balancing

February 10, 2014

Sheetal Dolas

and we are millionaire on kafka spout

10:06 AM

1.3 mn ps

10:07 AM



Ron Lee

hurray! so we broke a million!, time to celebrate

10:07 AM

Sheetal Dolas

kafka is knocked. now it is time for hbase

10:08 AM

HBase Tuning

This is where we began

February 10, 2014

Sheetal Dolas

ok now 5.8K per sec on hbase

2:43 PM

8 K with increased parallelism

2:50 PM

I think we have some configuration issues on HDFS

3:03 PM

Row Key Design

- Row Key design is critical (gets or scans or both?)
 - Keys with IP Addresses
 - Standard IP addresses have only two variations of the first character : 1 & 2
 - Minimum key length will be 7 characters and max 15 with a typical average of 12
 - Subnet range scans become difficult – range of 90 to 220 excludes 112
 - IP converted to hex (10.20.30.40 => 0a141e28)
 - gives 16 variations of first key character
 - consistently 8 character key
 - Easy to search for subnet ranges

Experiments with Row Key

February 10, 2014

3:45 PM

20K in HBase now - 30K message size

3:59 PM

February 11, 2014

Sheetal Dolas

base needle not moving beyond 55K

12:21 AM

*hbase

12:30 AM

moved from 20K to 55K ps but stuck there

12:32 AM

Region Splits

- Know your data
 - Auto split under high workload can result into hotspots and split storms
 - Understand your data and presplit the regions
 - Identify how many regions a RS can have to perform optimally. Use the formula below

$$(\text{RS memory}) * (\text{total memstore fraction}) / ((\text{memstore size}) * (\# \text{ column families}))$$

With Region Pre-Splits

February 12, 2014

Sheetal Dolas

I feel like laughing a monster laugh. base jumped from 55K to 350K

9:59 AM

*hbase

10:07 AM

Skype is always auto correcting hbase to base

10:08 AM

i wish it could auto correct the performance rather

10:08 AM

Know Your Application

- Enable Micro Batching (client side buffer)
- Smart shuffle/grouping in storm
- Understand your data and situationally exploit various WAL options
- Watch for many minor compactions
 - For heavy 'write' workload Increase `hbase.hstore.blockingStoreFiles` (we used 200)

And Finally

February 12, 2014

Sheetal Dolas

what the hell did I do wrong

5:03 PM

HBase hit 1 million

5:04 PM

Kafka Spout

Kafka Spout

- Parallelism is controlled by number of partitions per topic
 - Set Kafka spout parallelism equal to number of partitions in topic
 - Other key parameters that drive performance
 - `fetchSizeBytes`
 - `bufferSizeBytes`

Mysteriously Missing Data

March 31, 2014

Sheetal Dolas

my search for MH370 seems to be getting concluded

4:58 PM

On march 8th it was reported that data does not match between two systems, we came up with all weird theories, looked at every damn system involved, did 100 of tests and what not

▼ 5:01 PM

finally it turned out that a bug in kafka spout I am using is not reading all data and just dropping it

5:01 PM

for two weeks our speculations where changing like Malaysian officials

5:03 PM

Mysteriously Missing Data Root Cause

- A bug in Kafka spout that used to miss out some partitions and loose data
 - It is now fixed and available from Hortonworks repository (<http://repo.hortonworks.com/content/repositories/releases/org/apache/storm/storm-Kafka>)

Storm

Storm

- Every small thing counts at scale
 - Even simple string operations can slowdown throughput when executed on millions of Tuples



Sheetal Dolas

okay

so just sourcefire

Time in milliseconds to parse 100000 events

Original Transformation : 732

New Transformation - Removal of ':' from keys : 6143

New Transformation - Removal of ':' from keys + trim : 6763



jamsiro .

wow. really?

replaceall takes that long?

instead of trim you can just substring

Storm

- Error handling is critical
 - Poorly handled errors can lead to topology failure and eventually loss of data (or data duplication)

Spouts (All time)

Id ▲	Executors	Tasks	Emitted	Transferred	Complete latency (ms)	Acked	Failed	Last error
kafka	64	64	1326180	1326180	218	1319300	37723	storm.kafka.FailedFetchException: Error fetching data from [Partition{host=node09:9092, partition=3}] for topic [pcap]: [OFFSET_OUT_OF_RANGE] at storm.kafka.KafkaUtils.fetchMessages(KafkaUtils.java:1

Storm

- Tune & Scale individual spout and bolts before performance testing/tuning entire topology
 - Write your own simple data generator spouts and no-op bolts
- Making as many things configurable as possible helps a lot

Lessons Learned

- When it comes to Hadoop...partner up
- Separate the hype from the opportunity
- Start small then scale up
- Design Iteratively
- It doesn't work unless you have proven it at scale
- Keep an eye on ROI

Looking for Community Partners

Cisco + Hortonworks + Community Support for OpenSOC

How can you contribute?

- Technology Partner Program – contribute developers to join the Cisco and Hortonworks team

Thank you!

We are hiring:

jsirota@cisco.com

sheetal@hortonworks.com



CISCO



Hortonworks