

Real-time flow analysis using OSS based distributed infrastructure

OSSで実現する分散基盤を使ったリアルタイムflow分析

Shubham Saha Aakash Nand Abhik Datta Banik Oya Yu Asano Shuhei
NTT Communications

- Self Introduction
- Background
- Open Source Stack
- Anomaly detection
- Infrastructure and Architecture
- Demo
- Future Works
- Discussion

Introductions

Team Kaminari

- A team of network engineers, data engineers, and data scientists for analyzing backbone network traffic.
- We are developing the application of anomaly detection (DDoS etc) using ML and distributed infrastructure technology.
- About 1.5 years since team started

Background

DDoS attacks

- Number, scale, and severity of the impact of DDoS attacks.
- IoT devices and 5G fuel Botnets.

NTT Communications (ISP and ICT solutions provider) : monitor security threats in the network.

NEED OF THE HOUR

- Intelligent, managed DDoS protection solutions
- distributed infrastructure **X** deep learning for enhanced real-time efficiency.

Open Source Stack



pmacct is a small set of multi-purpose passive network monitoring tools



Apache Kafka is a stream-processing software platform capable of handling trillions of events a day.



Apache Hive is a data warehouse built on top of Apache Hadoop for providing data query and analysis.



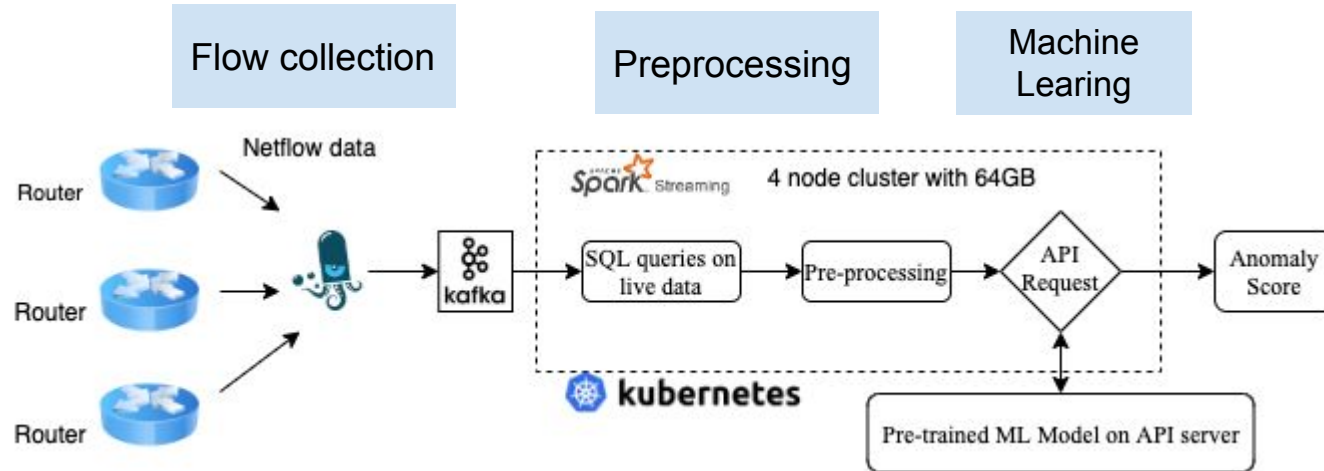
Apache Spark is an distributed general-purpose cluster-computing framework.



Kubernetes is a system for managing containerized applications across a cluster of nodes.

Infrastructure and Architecture

Architecture overview



pmacct

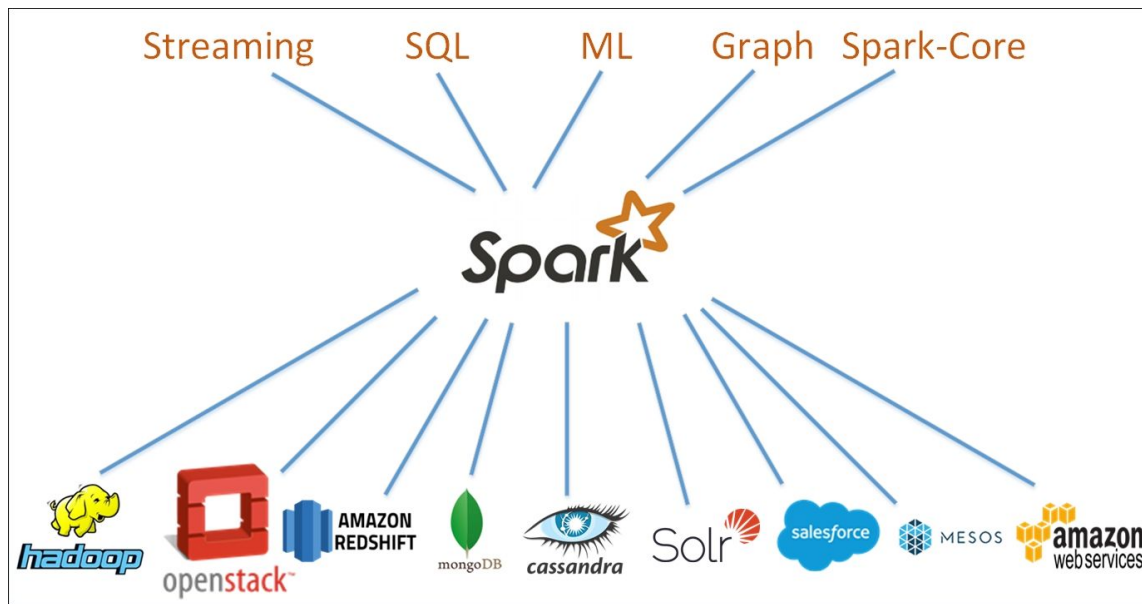
- Passive network monitoring tool for network data collection
- Streaming telemetry
- Collects data through libpcap, Netlink/NFLOG, NetFlow v1/v5/v7/v8/v9, sFlow v2/v4/v5 and IPFIX
- Saves data to a number of backends including:
 - Relational databases: MySQL, PostgreSQL and SQLite
 - noSQL databases: MongoDB and BerkeleyDB
 - AMQP message exchanges: RabbitMQ
 - Kafka message brokers
 - memory tables
 - flat files



<http://www.pmacct.net/>

What is Spark?

- Apache Spark is a open source lightning-fast unified analytics engine for big data and machine learning
- It is an optimized engine that supports general computation graphs for data analysis.



What is k8s?

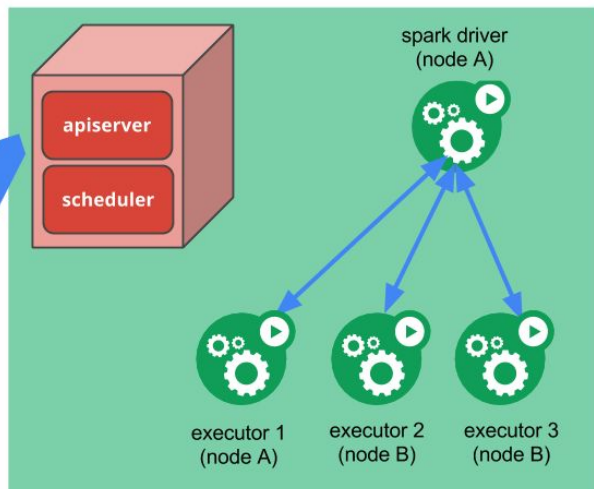
Kubernetes (K8s) is an open-source system for automating deployment, scaling, and management of containerized applications.



Why Spark?

- Speed
- Real-Time
- Deployment
- Native support for Machine Learning:
- Works well with Big Data platforms

kubernetes cluster



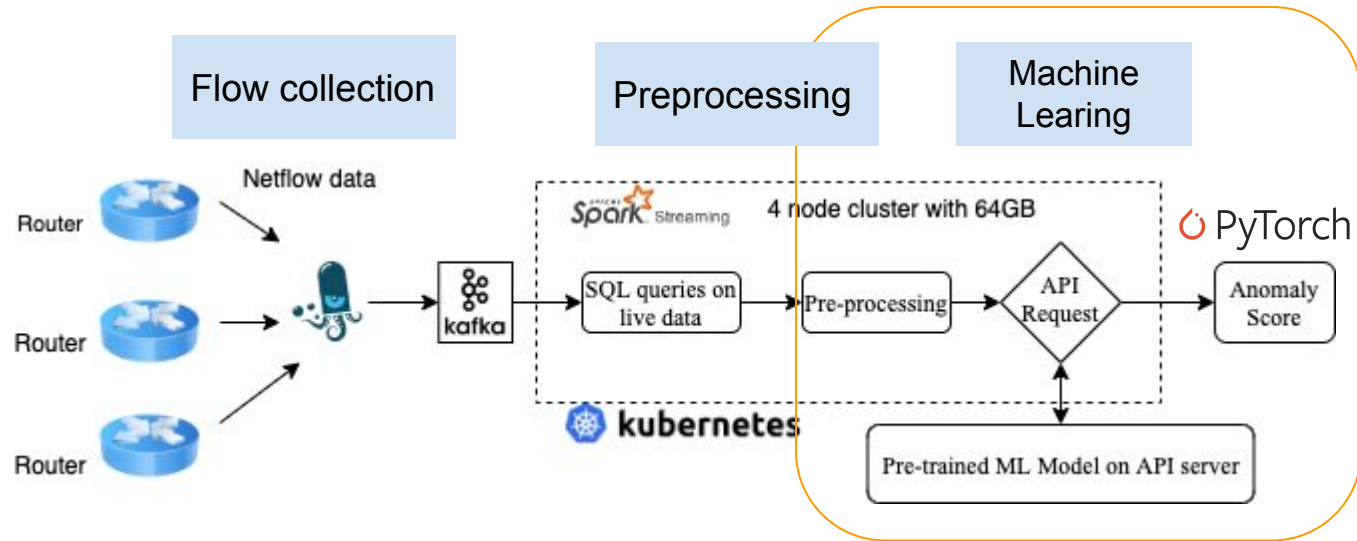
```
./bin/spark-submit \  
--master k8s://https://[REDACTED]:6443 \  
--deploy-mode cluster \  
--name predict-on-spark \  
--jars [REDACTED]/user/aakashnand/spark-streaming-kafka-  
0-8-assembly_2.11-2.4.6.jar \  
--conf spark.executor.instances=10 \  
--conf  
spark.kubernetes.container.image=[REDACTED]/test:latest \  
--conf spark.kubernetes.authenticate.driver.serviceAccountName=spark  
hdfs://[REDACTED]:8020/tmp/asano/predict-on-spark.py
```

Why Spark on Kubernetes (k8s)

	Spark on Yarn	Spark on k8s
Dependency management	Poor ✖	Good ●
Admin Overhead	High ✖	Low ●
Spark Version management	Rigid ✖	Flexible ●
Container Customization	Poor ✖	Good ●
Resource Allocator	Yarn Resource Manager	k8s Scheduler API
Spark application management	Convenient ●	Difficult ✖
Learning Curve	fast ●	slow ✖

Anomaly Detection

Architecture overview



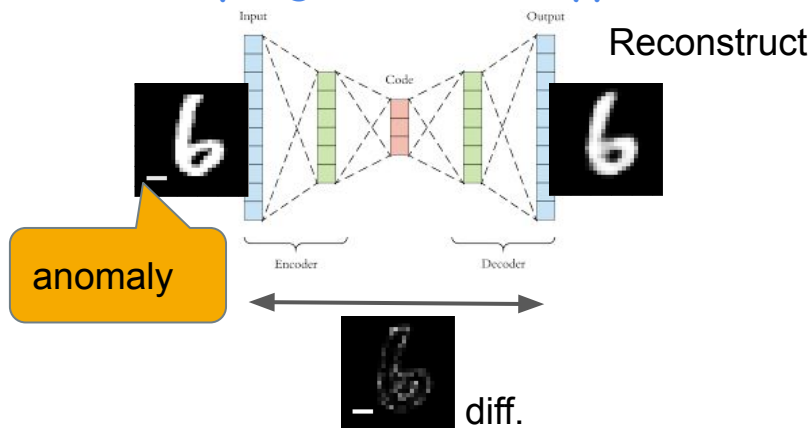
What is Torch?

- Torch is an open source machine learning library.
- A number of pieces of Deep Learning software are built on top of Torch.



Very powerful
(High accuracy)

Very useful
(Low cost of dev.)



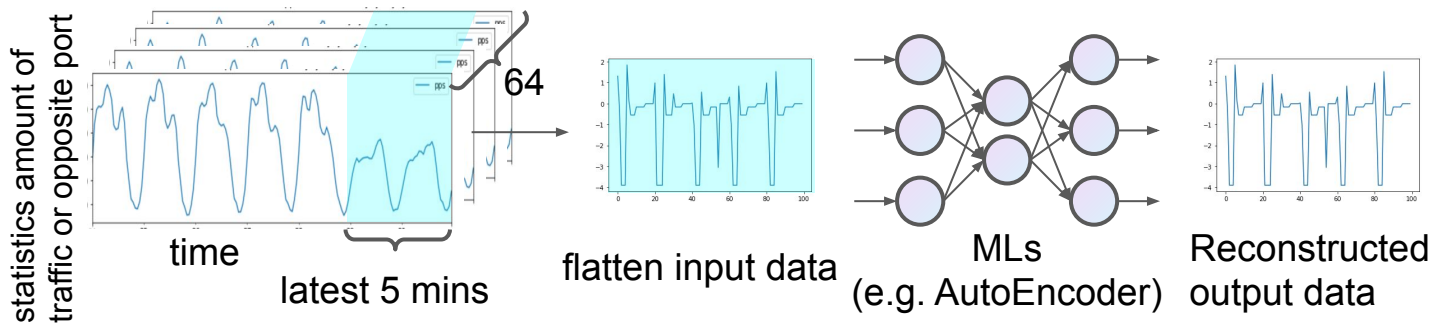
```
43 import torch
44 import torch.nn as nn
45 def preprocessing(x):
46     ... return
47 class autoencoder(nn.Module):
48     ... def __init__(self, inunit=320, unit=20, mean=0):
49         ... super(autoencoder, self).__init__()
50         ... self.preprocessing = lambda x: x.add(1).log().add(mean)
51         ... self.encoder = nn.Sequential(
52             ... nn.Linear(inunit, unit),
53             ... nn.Sigmoid(),
54             ... nn.Linear(unit, unit),
55             ... nn.Sigmoid())
56         ... self.decoder = nn.Sequential(
57             ... nn.Linear(unit, unit),
58             ... nn.Sigmoid(),
59             ... nn.Linear(unit, inunit))
60
61     ... def forward(self, x):
62         ... x = self.preprocessing(x)
63         ... z = self.encoder(x)
64         ... z = self.decoder(z)
65         ... z = nn.functional.mse_loss(x, z, reduction='none')
66         ... if z.dim() == 1:
67             ... z = torch.mean(z)
68         ... else:
69             ... z = torch.mean(z, dim=1, keepdim=True)
70         ... return z
```

- Preprocessing

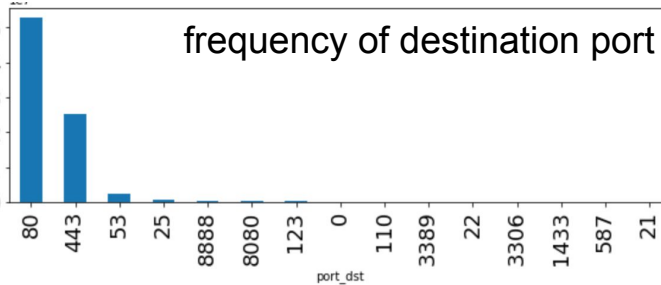
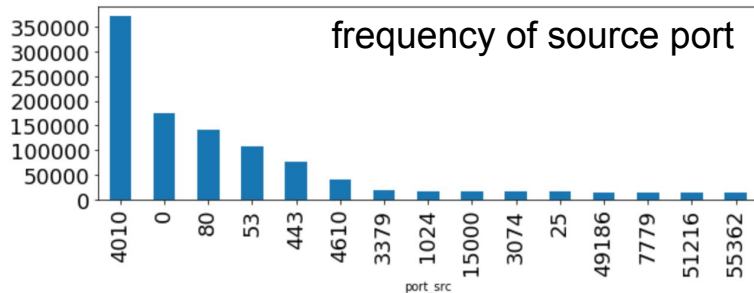
5 tuples --> counting data --> input data (vector)

- Training and Prediction (per 5 mins.)

input data

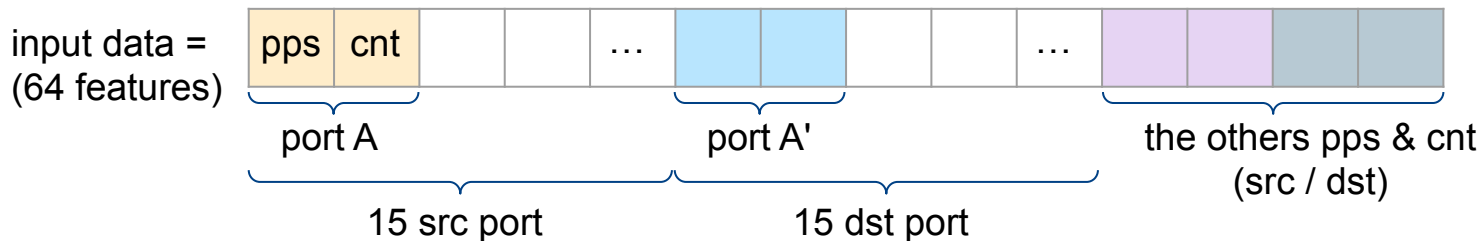


- To select 15 + 15 ports to avoid sparse array



src_port	usage	dst_port	usage
4010	Samsung Unidex ?	80	http
0	Fragment ?	443	https
80	HTTP	53	DNS
53	DNS	25	SMTP
443	HTTPS	8888	HTTP alternate
4610	QualiSystems	8080	HTTP alternate
3379	SOCORFS	123	NTP
1024	?	0	Fragment ?
15000	psyBNC	110	POP3
3074	Xbox LIVE	3389	RDP (microsoft)
25	SMTP	22	ssh
49186	cisco dpe port ?	3306	MySQL
7779	trojan ? botnet C&C	1433	MSSQL
51216	?	587	SMTP
55362	?	21	IFTTP

- Vectorization of traffic (pps) and number of opposite ports (count)



Demo

```
from pyspark.streaming.kafka import KafkaUtils
from pyspark.streaming import StreamingContext
from pyspark.sql import Row
import json
from datetime import datetime

ssc = StreamingContext(sc, 60) #60秒間隔
kafkaStream = KafkaUtils.createStream(ssc, KAFKA_ZOOKEEPER_IP ,
'test_id', {TOPIC_NAME:1})
lines = kafkaStream \
    .map(lambda x: json.loads(x[1])) \
    .map(lambda x: Row(
        timestamp=datetime.strptime(x["timestamp_arrival"],
'%Y-%m-%dT%H:%M:%S.%fz').timestamp(),
        port_src=x["port_src"],
        packets=x["packets"]
    ))
lines.pprint(num=5)
ssc.start()
ssc.awaitTermination()
```

■ Input flow from Kafka

Time: 2020-08-26 19:29:00

Row(packets=1, port_src=56108, timestamp=1598437642.0)
Row(packets=1, port_src=80, timestamp=1598437642.0)
Row(packets=1, port_src=11507, timestamp=1598437642.0)
Row(packets=1, port_src=28400, timestamp=1598437642.0)
Row(packets=1, port_src=45393, timestamp=1598437642.0)

■ Preprocessing flow

```
lines.foreachRDD(process)

def process(lines):
    ...
    df_dst = df.groupby("timestamp") \
                .pivot("port_dst") \
                .agg(F.sum("packets").alias("pps_src"),
                    F.count("port_src").alias("num_src"))
    ...
```

You can use SQL/pandas like functions to filter, aggregate, groupby etc.


```
def process(lines):  
    ...  
    AE_scores= AE_API(features)  
  
def AE_API(x):  
    ...  
    result = requests.post("http://****", data=pickle.dumps(x))  
    return result
```

```
def process(lines):  
    ...  
    outputs= np.concatenate([features,AE_scores,AR_scores], 1)  
    VIS_API(outputs)  
  
def VIS_API(lines)  
    ...  
    requests.post("http://****", data=pickle.dumps(x))
```

- Post preprocessed data to API running on flask
- AE_API returns anomaly score by Auto Encoder
- VIS_API stores and visualize result

Future Work

Our vision is to make anomaly detection not only easy to use but also universal.

Today we have demonstrated highly accurate DDoS detection using machine learning.

In future we would like to

- enhance the deployment and operation of the system using CI/CD
- extend this capability to other anomalies in networks and other infrastructures.
 - using various data sources
 - custom algorithms for different data source types

- How have you treated large volume traffic logs on your network?
- What else can be realized using
Netflow × Distributed Computing × Machine Learning?
- How to input netflow into machine learning model?
- Which features do you think should be selected (ports, protocols...)?

Questions?

